

# Cloud Container Engine

## API Reference

**Issue** 01  
**Date** 2025-03-04



**Copyright © Huawei Cloud Computing Technologies Co., Ltd. 2025. All rights reserved.**

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Cloud Computing Technologies Co., Ltd.

## **Trademarks and Permissions**



HUAWEI and other Huawei trademarks are the property of Huawei Technologies Co., Ltd.

All other trademarks and trade names mentioned in this document are the property of their respective holders.

## **Notice**

The purchased products, services and features are stipulated by the contract made between Huawei Cloud and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

## **Huawei Cloud Computing Technologies Co., Ltd.**

Address: Huawei Cloud Data Center Jiaoxinggong Road  
Qianzhong Avenue  
Gui'an New District  
Gui Zhou 550029  
People's Republic of China

Website: <https://www.huaweicloud.com/intl/en-us/>

---

# Contents

---

<b>1 Before You Start.....</b>	<b>1</b>
<b>2 API Overview.....</b>	<b>5</b>
<b>3 Calling APIs.....</b>	<b>24</b>
3.1 Making an API Request.....	24
3.2 Authentication.....	28
3.3 Response.....	29
<b>4 APIs.....</b>	<b>32</b>
4.1 API URL.....	32
4.2 Cluster Management.....	33
4.2.1 Creating a Cluster.....	33
4.2.2 Reading a Specified Cluster.....	79
4.2.3 Listing Clusters in a Specified Project.....	98
4.2.4 Updating a Specified Cluster.....	118
4.2.5 Deleting a Cluster.....	140
4.2.6 Hibernating a Cluster.....	163
4.2.7 Waking Up a Cluster.....	164
4.2.8 Obtaining a Cluster Certificate.....	165
4.2.9 Modifying Cluster Specifications.....	170
4.2.10 Querying a Job.....	174
4.2.11 Binding/Unbinding Public API Server Address.....	178
4.3 Node Management.....	183
4.3.1 Creating a Node.....	183
4.3.2 Reading a Specified Node.....	237
4.3.3 Listing All Nodes in a Cluster.....	262
4.3.4 Updating a Specified Node.....	287
4.3.5 Deleting a Node.....	312
4.3.6 Enabling Scale-In Protection for a Node.....	337
4.3.7 Disabling Scale-In Protection for a Node.....	339
4.3.8 Synchronizing Nodes.....	340
4.3.9 Accepting a Node.....	342
4.3.10 Managing a Node in a Customized Node Pool.....	361
4.3.11 Resetting a Node.....	364

4.3.12 Removing a Node.....	384
4.3.13 Migrating a Node.....	389
4.4 Node Pool Management.....	401
4.4.1 Creating a Node Pool.....	401
4.4.2 Reading a Specified Node Pool.....	468
4.4.3 Listing All Node Pools in a Specified Cluster.....	498
4.4.4 Updating a Specified Node Pool.....	529
4.4.5 Deleting a Node Pool.....	583
4.4.6 Scaling a Node Pool.....	613
4.4.7 Synchronizing Node Pools.....	616
4.5 Storage Management.....	621
4.5.1 Creating a PVC (to be discarded).....	621
4.5.2 Deleting a PVC (to be discarded).....	629
4.6 Add-on Management.....	633
4.6.1 Installing an Add-on Instance.....	633
4.6.2 Listing Add-on Templates.....	643
4.6.3 Updating an Add-on Instance.....	650
4.6.4 Rolling Back an Add-on Instance.....	661
4.6.5 Deleting an Add-on Instance.....	670
4.6.6 Querying an Add-on Instance.....	671
4.6.7 Listing Add-on Instances.....	677
4.7 Tag Management.....	683
4.7.1 Adding Resource Tags to a Specified Cluster in Batches.....	683
4.7.2 Deleting Resource Tags of a Specified Cluster in Batches.....	685
4.8 Configuration Management.....	688
4.8.1 Obtaining the Parameters That Can Be Configured for a Node Pool.....	688
4.8.2 Obtaining the List of Parameters That Can Be Configured for a Cluster.....	689
4.8.3 Obtaining the Parameters That Can Be Configured for a Node Pool.....	691
4.8.4 Changing the Values of Configuration Parameters of a Node Pool.....	694
4.9 Chart Management.....	700
4.9.1 Uploading a Chart.....	700
4.9.2 Obtaining a Chart List.....	702
4.9.3 Obtaining a Release List.....	704
4.9.4 Updating a Chart.....	708
4.9.5 Creating a Release.....	710
4.9.6 Deleting a Chart.....	714
4.9.7 Updating a Release.....	715
4.9.8 Obtaining a Chart.....	719
4.9.9 Deleting a Release.....	721
4.9.10 Downloading a Chart.....	723
4.9.11 Obtaining a Release.....	724
4.9.12 Obtaining Chart Values.....	727

4.9.13 Obtaining Historical Records of a Release.....	728
4.9.14 Obtaining the Quota of a User Chart.....	731
4.10 Add-on Instance Parameters.....	733
4.10.1 CoreDNS.....	733
4.10.2 CCE Container Storage (Everest).....	738
4.10.3 CCE Node Problem Detector.....	742
4.10.4 Kubernetes Dashboard.....	746
4.10.5 CCE Cluster Autoscaler.....	748
4.10.6 NGINX Ingress Controller.....	753
4.10.7 Kubernetes Metrics Server.....	758
4.10.8 CCE Advanced HPA.....	761
4.10.9 CCE AI Suite (NVIDIA GPU).....	765
4.10.10 Volcano Scheduler.....	768
4.10.11 CCE Secrets Manager for DEW.....	773
4.10.12 CCE Network Metrics Exporter.....	774
4.10.13 NodeLocal DNSCache.....	776
4.10.14 Cloud Native Cluster Monitoring.....	780
4.10.15 Cloud Native Logging.....	783
<b>5 Kubernetes APIs.....</b>	<b>789</b>
<b>6 Permissions and Supported Actions.....</b>	<b>794</b>
<b>7 Appendix.....</b>	<b>801</b>
7.1 Status Code.....	801
7.2 Error Codes.....	805
7.3 Obtaining a Project ID.....	811
7.4 Obtaining an Account ID.....	812
7.5 Specifying Add-ons to Be Installed During Cluster Creation.....	813
7.6 How to Obtain Parameters in the API URI.....	815
7.7 Creating a VPC and Subnet.....	816
7.8 Creating a Key Pair.....	817
7.9 Node Flavor Description.....	817
7.10 Adding a Salt in the password Field When Creating a Node.....	817
7.11 Maximum Number of Pods That Can Be Created on a Node.....	820
7.12 Node OS.....	822
7.13 Space Allocation of a Data Disk.....	822
7.14 Attaching Disks to a Node.....	826

# 1 Before You Start

---

Cloud Container Engine (CCE) is a container service that allows you to run containers efficiently in the cloud. CCE provides highly scalable, high-performance, enterprise-class Kubernetes clusters and supports Docker containers. With CCE, you can easily deploy, manage, and scale containerized applications in the cloud.

This document describes how to use APIs for performing operations on CCE, such as creating or deleting CCE resources, modifying resource specifications, or adding NICs. For details about all supported operations, see [API Overview](#).

If you plan to access CCE resources through an API, ensure that you are familiar with CCE concepts.

CCE supports both Kubernetes-native APIs and proprietary APIs. With these APIs, you can use all functions of CCE.

- CCE has opened APIs through API gateways to support operations on cloud service infrastructures (for example, creating a node). Operations on cluster resources (such as [creating a workload](#)) are also supported.
- Kubernetes-native APIs: You can perform operations on cluster resources (such as [creating a workload](#)) using the Kubernetes-native API server. However, operations on cloud service infrastructures (such as creating a node) are not supported.

For details about Kubernetes-native API versions, see <https://kubernetes.io/docs/concepts/overview/kubernetes-api/>.

## NOTE

- The Kubernetes-native APIs called in the current version do not support HTTP persistent connections.
- The Kubernetes-native APIs in the current version include Beta APIs, whose version names include **beta**, for example, **v1beta1**. This type of APIs varies according to Kubernetes-native APIs. Therefore, you are advised to use this type of APIs in unimportant scenarios, for example, short-term test clusters.

CCE supports Representational State Transfer (REST) APIs, allowing you to call APIs using HTTPS. For details about API calling, see [3 Calling APIs](#).

## Endpoints

An endpoint is the **request address** for calling an API. Endpoints vary depending on services and regions. An endpoint can be obtained from [Regions and Endpoints](#).

You need to select an endpoint based on your service requirements.

- The URL format for cluster, node, node pool, add-on, and quota management is **https://Endpoint/uri**. *uri* indicates the resource path, that is, the API access path.
- The URL format for Kubernetes APIs, storage management, and add-on management is **https://{clusterid}.Endpoint/uri**. In the URL, *{clusterid}* indicates the cluster ID, and *uri* indicates the resource path, that is, the path for API access.

### NOTE

- The format of the URL called by the add-on management APIs is **https://{clusterid}.Endpoint/uri**. However, *{clusterid}* is used only for the domain name and is not verified or used by the APIs. Set *{clusterid}* in the query or body. For details about *{clusterid}*, see the add-on management sections.
- *{clusterid}* is required for Kubernetes APIs and storage management, which indicates the cluster that needs to be accessed by calling the API.

**Table 1-1** URL parameters

Parameter	Description
{clusterid}	Cluster ID. After a cluster is created, call the <a href="#">API for obtaining a cluster in a specified project</a> to obtain the cluster ID.
Endpoint	Entry (URL) for a web service. Endpoints vary depending on services and regions.
uri	Access path of an API for performing an operation. Obtain the path from the URI of an API. For example, the <b>resource-path</b> of the API used to obtain a user token is <b>v3/auth/tokens</b> .

## Notes and Constraints

- CCE imposes a quota on the number and capacity of resources that a user can access. By default, you can create a maximum of five clusters in each region and a cluster can have a maximum of 50 nodes.
- For more constraints, see API description.

## Basic Concepts

- Domain  
A domain is created upon successful registration. The domain has full access permissions for all of its cloud services and resources. It can be used to reset user passwords and grant user permissions. The domain user is a payment

entity, which should not be used directly to perform routine management. For security purposes, create users and grant them permissions for routine management.

- User

An IAM user is created using an account to use cloud services. Each IAM user has their own identity credentials (password and access keys).

The account name, username, and password will be required for API authentication.

- Region

A region is a geographic area in which cloud resources are deployed. Availability zones (AZs) in the same region can communicate with each other over an intranet, while AZs in different regions are isolated from each other. Deploying cloud resources in different regions can better suit certain user requirements or comply with local laws or regulations.


- AZ

An AZ comprises of one or more physical data centers equipped with independent ventilation, fire, water, and electricity facilities. Computing, network, storage, and other resources in an AZ are logically divided into multiple clusters. AZs within a region are interconnected using high-speed optical fibers to allow you to build cross-AZ high-availability systems.

- Project

A project corresponds to a region. Default projects are defined to group and physically isolate resources (including compute, storage, and network resources) across regions. Users can be granted permissions in a default project to access all resources under their accounts in the region associated with the project. If you need more refined access control, create subprojects under a default project and create resources in subprojects. Then you can assign users the permissions required to access only the resources in the specific subprojects.

**Figure 1-1** Project isolation model



- Enterprise project

Enterprise projects group and manage resources across regions. Resources in different enterprise projects are logically isolated. An enterprise project can contain resources of multiple regions, and resources can be added to or removed from enterprise projects.



For details about enterprise projects and about how to obtain enterprise project IDs, see *Enterprise Management User Guide*.

# 2 API Overview

APIs provided by CCE are classified into two types: proprietary APIs and Kubernetes-native APIs. By using these two types of APIs, you can use all functions provided by CCE, including creating clusters and nodes, using the Kubernetes-native APIs to create workloads, and using the proprietary CCE APIs to monitor application data.

Type	Subtype	Description
Proprietary CCE APIs	<b>Cluster-related APIs</b>	Manage clusters, including creating and deleting clusters. You can use APIs in this category to create clusters and obtain information about created clusters.
	<b>Node-related APIs</b>	Manage nodes, including creating and deleting nodes. You can use these APIs in this category to add nodes to clusters and obtain information about created nodes.
	<b>Node Pool-related APIs</b>	Manage node pools, including creating and deleting node pools. You can use APIs in this category to create node pools and obtain information about created node pools.
	<b>Add-on-related APIs</b>	Manage add-ons, including querying AddonTemplates and creating, updating, deleting, and obtaining AddonInstances.
Kubernetes-native APIs	N/A	For details about how to call Kubernetes-native APIs, see <a href="#">Kubernetes APIs</a> .

## Cluster-related APIs

**Table 2-1** Cluster-related APIs

API	Description
<a href="#">Creating a Cluster</a>	Create an empty cluster, which has only master nodes but do not have worker nodes.
<a href="#">Reading a Specified Cluster</a>	Obtain details about a specified cluster.
<a href="#">Listing Clusters in a Specified Project</a>	Obtain details about all clusters in a specified project.
<a href="#">Updating a Specified Cluster</a>	Update information about a specified cluster.
<a href="#">Deleting a Cluster</a>	Delete a specified cluster.
<a href="#">Hibernate a Cluster</a>	Hibernate a specified cluster.
<a href="#">Waking Up a Cluster</a>	Wake up a hibernated cluster.
<a href="#">Obtaining Cluster Certificates</a>	Obtain certificates of a specified cluster.
<a href="#">Obtaining Job Information</a>	Obtain the progress of a job with a specified job ID returned after a job request is issued.

## Node-related APIs

**Table 2-2** Node-related APIs

API	Description
<a href="#">Creating a Node</a>	Create a node in a specified cluster.
<a href="#">Reading a Specified Node</a>	Obtain details about a node with a specified node ID.
<a href="#">Reading all Nodes in a Cluster</a>	Obtain details about all nodes in a cluster with a specified cluster ID.
<a href="#">Updating a Specified Node</a>	Update information about a specified node.
<a href="#">Deleting a Node</a>	Delete a specified node
<a href="#">Accepting a Node</a>	Accept a node into a specified cluster.
<a href="#">Resetting a Node</a>	Reset a node in a specified cluster.

API	Description
<a href="#">Removing a Node</a>	Remove a node from a specified cluster.
<a href="#">Migrating a Node</a>	Migrate a node from a specified cluster to another cluster.

## Node Pool-related APIs

Table 2-3 Node pool-related APIs

API	Description
<a href="#">Creating a Node Pool</a>	Create a node pool in a specified cluster.
<a href="#">Reading a Specified Node Pool</a>	Obtain details about a node with a specified node ID.
<a href="#">Listing All Node Pools in a Specified Cluster</a>	Obtain details about all node pools in a cluster with a specified cluster ID.
<a href="#">Updating a Specified Node Pool</a>	Update information about a specified node pool.
<a href="#">Deleting a Node Pool</a>	Delete a specified node pool.

## Add-on-related APIs

Table 2-4 Add-on-related APIs

API	Description
<a href="#">Installing an Add-on Instance</a>	Install an add-on by using the add-on template.
<a href="#">Listing Add-on Templates</a>	Query add-on information.
<a href="#">Updating an Add-on Instance</a>	Update an add-on instance.
<a href="#">Deleting an Add-on Instance</a>	Delete an add-on instance.
<a href="#">Reading an Add-on Instance</a>	Obtain details about an add-on instance.
<a href="#">Listing Add-on Instances</a>	List all add-on instances in the cluster.

## Kubernetes APIs

API	Function	URI
Node	Reading a specified node	GET /api/v1/nodes/{name}
	Listing all nodes	GET /api/v1/nodes
	Updating a specified node	PATCH /api/v1/nodes/{name}
Namespace	Creating a namespace	POST /api/v1/namespaces
	Deleting a namespace	DELETE /api/v1/namespaces/{name}
	Querying a specified namespace	GET /api/v1/namespaces/{name}
	Replacing a specified namespace	PUT /api/v1/namespaces/{name}
	Replacing the status of a specified namespace	PUT /api/v1/namespaces/{name}/status
	Replacing the finalize values of a specified namespace	PUT /api/v1/namespaces/{name}/finalize
	Listing namespaces	GET /api/v1/namespaces
	Updating a specified namespace	PATCH /api/v1/namespaces/{name}
Resource quotas	Querying resource quotas	GET /api/v1/resourcequotas
	Creating a resource quota	POST /api/v1/namespaces/{namespace}/resourcequotas
	Updating a resource quota	PUT /api/v1/namespaces/{namespace}/resourcequotas/{name}
	Deleting a resource quota	DELETE /api/v1/namespaces/{namespace}/resourcequotas/{name}
Pod	Creating a pod	POST /api/v1/namespaces/{namespace}/pods
	Deleting a pod	DELETE /api/v1/namespaces/{namespace}/pods/{name}
	Deleting all pods	DELETE /api/v1/namespaces/{namespace}/pods
	Reading a specified pod	GET /api/v1/namespaces/{namespace}/pods/{name}

API	Function	URI
	Replacing a specified pod	PUT /api/v1/namespaces/{namespace}/pods/{name}
	Replacing the status of a specified pod	PUT /api/v1/namespaces/{namespace}/pods/{name}/status
	Listing all pods in a specified namespace	GET /api/v1/namespaces/{namespace}/pods
	Listing pods	GET /api/v1/pods
	Updating a specified pod	PATCH /api/v1/namespaces/{namespace}/pods/{name}
Deploy ment	Creating a Deployment	POST /apis/apps/v1/namespaces/{namespace}/deployments
	Rolling back a Deployment	PATCH /apis/apps/v1/namespaces/{namespace}/deployments/{name} (for clusters of v1.17 or later) POST /apis/apps/v1beta1/namespaces/{namespace}/deployments/{name}/rollback (for clusters of v1.15 and earlier) POST /apis/extensions/v1beta1/namespaces/{namespace}/deployments/{name}/rollback (for clusters of v1.15 and earlier)
	Deleting a Deployment	DELETE /apis/apps/v1/namespaces/{namespace}/deployments/{name}
	Deleting all Deployments	DELETE /apis/apps/v1/namespaces/{namespace}/deployments
	Reading a specified Deployment	GET /apis/apps/v1/namespaces/{namespace}/deployments/{name}
	Reading the status of a specified Deployment	GET /apis/apps/v1/namespaces/{namespace}/deployments/{name}/status
	Reading the scaling operation of a specified Deployment	GET /apis/apps/v1/namespaces/{namespace}/deployments/{name}/scale
	Replacing a specified Deployment	PUT /apis/apps/v1/namespaces/{namespace}/deployments/{name}
	Replacing the status of a specified Deployment	PUT /apis/apps/v1/namespaces/{namespace}/deployments/{name}/status
	Replacing the scaling operation of a specified Deployment	PUT /apis/apps/v1/namespaces/{namespace}/deployments/{name}/scale

API	Function	URI
	Listing Deployments in a specified namespace	GET /apis/apps/v1/namespaces/{namespace}/deployments
	Listing all Deployments	GET /apis/apps/v1/deployments
	Updating a specified Deployment	PATCH /apis/apps/v1/namespaces/{namespace}/deployments/{name}
	Updating the status of a specified Deployment	PATCH /apis/apps/v1/namespaces/{namespace}/deployments/{name}/status
	Updating the scaling operation of a specified Deployment	PATCH /apis/apps/v1/namespaces/{namespace}/deployments/{name}/scale
StatefulSet	Creating a StatefulSet	POST /apis/apps/v1/namespaces/{namespace}/statefulsets
	Deleting a specified StatefulSet	DELETE /apis/apps/v1/namespaces/{namespace}/statefulsets/{name}
	Deleting all StatefulSets	DELETE /apis/apps/v1/namespaces/{namespace}/statefulsets
	Reading a specified StatefulSet	GET /apis/apps/v1/namespaces/{namespace}/statefulsets/{name}
	Reading the status of a specified StatefulSet	GET /apis/apps/v1/namespaces/{namespace}/statefulsets/{name}/status
	Replacing a specified StatefulSet	PUT /apis/apps/v1/namespaces/{namespace}/statefulsets/{name}
	Replacing the status of a specified StatefulSet	PUT /apis/apps/v1/namespaces/{namespace}/statefulsets/{name}/status
	Listing StatefulSets in a specified namespace	GET /apis/apps/v1/namespaces/{namespace}/statefulsets
	Listing all StatefulSets	GET /apis/apps/v1/statefulsets
	Updating a specified StatefulSet	PATCH /apis/apps/v1/namespaces/{namespace}/statefulsets/{name}
	Updating the status of a specified StatefulSet	PATCH /apis/apps/v1/namespaces/{namespace}/statefulsets/{name}/status
	DaemonSet	Creating a DaemonSet
Deleting a specified DaemonSet		DELETE /apis/apps/v1/namespaces/{namespace}/daemonsets/{name}
Deleting all DaemonSets		DELETE /apis/apps/v1/namespaces/{namespace}/daemonsets

API	Function	URI
	Reading a specified DaemonSet	GET /apis/apps/v1/namespaces/{namespace}/daemonsets/{name}
	Reading the status of a specified DaemonSet	GET /apis/apps/v1/namespaces/{namespace}/daemonsets/{name}/status
	Updating a specified DaemonSet	PATCH /apis/apps/v1/namespaces/{namespace}/daemonsets/{name}
	Updating the status of a specified DaemonSet	PATCH /apis/apps/v1/namespaces/{namespace}/daemonsets/{name}/status
	Listing all DaemonSets	GET /apis/apps/v1/daemonsets
	Listing DaemonSets in a specified namespace	GET /apis/apps/v1/namespaces/{namespace}/daemonsets
	Replacing a specified DaemonSet	PUT /apis/apps/v1/namespaces/{namespace}/daemonsets/{name}
	Replacing the status of a specified DaemonSet	PUT /apis/apps/v1/namespaces/{namespace}/daemonsets/{name}/status
Job	Creating a job	POST /apis/batch/v1/namespaces/{namespace}/jobs
	Deleting a job	DELETE /apis/batch/v1/namespaces/{namespace}/jobs/{name}
	Deleting all jobs	DELETE /apis/batch/v1/namespaces/{namespace}/jobs
	Reading a specified job	GET /apis/batch/v1/namespaces/{namespace}/jobs/{name}
	Reading the status of a specified job	GET /apis/batch/v1/namespaces/{namespace}/jobs/{name}/status
	Replacing a specified job	PUT /apis/batch/v1/namespaces/{namespace}/jobs/{name}
	Replacing the status of a specified job	PUT /apis/batch/v1/namespaces/{namespace}/jobs/{name}/status
	Listing jobs in a specified namespace	GET /apis/batch/v1/namespaces/{namespace}/jobs
	Listing all jobs	GET /apis/batch/v1/jobs
	Updating the status of a specified job	PATCH /apis/batch/v1/namespaces/{namespace}/jobs/{name}/status
	Updating a specified job	PATCH /apis/batch/v1/namespaces/{namespace}/jobs/{name}



API	Function	URI
CronJob	Creating a cron job	POST /apis/batch/v1/namespaces/{namespace}/cronjobs (only for clusters of v1.25 or later) POST /apis/batch/v1beta1/namespaces/{namespace}/cronjobs (only for clusters of v1.23 or earlier)
	Deleting a cron job	DELETE /apis/batch/v1/namespaces/{namespace}/cronjobs/{name} (only to clusters of v1.25 or later) DELETE /apis/batch/v1beta1/namespaces/{namespace}/cronjobs/{name} (only for clusters of v1.23 or earlier)
	Deleting all cron jobs	DELETE /apis/batch/v1/namespaces/{namespace}/cronjobs (only for clusters of v1.25 or later) DELETE /apis/batch/v1beta1/namespaces/{namespace}/cronjobs (only for clusters of v1.23 or earlier)
	Reading a specified cron job	GET /apis/batch/v1/namespaces/{namespace}/cronjobs/{name} (only for clusters of v1.25 or later) GET /apis/batch/v1beta1/namespaces/{namespace}/cronjobs/{name} (only for clusters of v1.23 or earlier)
	Reading the status of a specified cron job	GET /apis/batch/v1/namespaces/{namespace}/cronjobs/{name}/status (only for clusters of v1.25 or later) GET /apis/batch/v1beta1/namespaces/{namespace}/cronjobs/{name}/status (only for clusters of v1.23 or earlier)
	Replacing a specified cron job	PUT /apis/batch/v1/namespaces/{namespace}/cronjobs/{name} (only for clusters of v1.25 or later) PUT /apis/batch/v1beta1/namespaces/{namespace}/cronjobs/{name} (only for clusters of v1.23 or earlier)
	Replacing the status of a specified cron job	PUT /apis/batch/v1/namespaces/{namespace}/cronjobs/{name}/status (only for clusters of v1.25 or later) PUT /apis/batch/v1beta1/namespaces/{namespace}/cronjobs/{name}/status (only for clusters of v1.23 or earlier)

API	Function	URI
	Listing cron jobs under a specified namespace	GET /apis/batch/v1/namespaces/{namespace}/cronjobs (only for clusters of v1.25 or later) GET /apis/batch/v1beta1/namespaces/{namespace}/cronjobs (only for clusters of v1.23 or earlier)
	Listing all cron jobs	GET /apis/batch/v1/cronjobs (only for clusters of v1.25 or later) GET /apis/batch/v1beta1/cronjobs (only for clusters of v1.23 or earlier)
	Updating the status of a specified cron job	PATCH /apis/batch/v1/namespaces/{namespace}/cronjobs/{name}/status (only for clusters of v1.25 or later) PATCH /apis/batch/v1beta1/namespaces/{namespace}/cronjobs/{name}/status (only for clusters of v1.23 or earlier)
	Updating a specified cron job	PATCH /apis/batch/v1/namespaces/{namespace}/cronjobs/{name} (only for clusters of v1.25 or later) PATCH /apis/batch/v1beta1/namespaces/{namespace}/cronjobs/{name} (only for clusters of v1.23 or earlier)
ReplicaSet	Listing ReplicaSets	GET /apis/apps/v1/namespaces/{namespace}/replicasets
	Reading a specified ReplicaSet	GET /apis/apps/v1/namespaces/{namespace}/replicasets/{name}
	Listing all ReplicaSets	GET /apis/apps/v1/replicasets
ReplicationController	Creating a ReplicationController	POST /api/v1/namespaces/{namespace}/replicationcontrollers
	Deleting a ReplicationController	DELETE /api/v1/namespaces/{namespace}/replicationcontrollers/{name}
	Deleting all ReplicationControllers	DELETE /api/v1/namespaces/{namespace}/replicationcontrollers
	Reading a ReplicationController under a specified namespace	GET /api/v1/namespaces/{namespace}/replicationcontrollers/{name}
	Replacing a ReplicationController under a specified namespace	PUT /api/v1/namespaces/{namespace}/replicationcontrollers/{name}

API	Function	URI
	Replacing the status of a ReplicationController under a specified namespace	PUT /api/v1/namespaces/{namespace}/replicationcontrollers/{name}/status
	Listing ReplicationControllers in a specified namespace	GET /api/v1/namespaces/{namespace}/replicationcontrollers
	Listing ReplicationControllers	GET /api/v1/replicationcontrollers
	Updating a specified ReplicationController	PATCH /api/v1/namespaces/{namespace}/replicationcontrollers/{name}
Endpoints	Creating an endpoint	POST /api/v1/namespaces/{namespace}/endpoints
	Deleting an endpoint	DELETE /api/v1/namespaces/{namespace}/endpoints/{name}
	Deleting all endpoints	DELETE /api/v1/namespaces/{namespace}/endpoints
	Querying a specified endpoint	GET /api/v1/namespaces/{namespace}/endpoints/{name}
	Replacing a specified endpoint	PUT /api/v1/namespaces/{namespace}/endpoints/{name}
	Listing endpoints	GET /api/v1/endpoints
	Listing endpoints in a specified namespace	GET /api/v1/namespaces/{namespace}/endpoints
	Updating a specified endpoint	PATCH /api/v1/namespaces/{namespace}/endpoints/{name}
Service	Creating a Service	POST /api/v1/namespaces/{namespace}/services
	Deleting a specified Service	DELETE /api/v1/namespaces/{namespace}/services/{name}
	Obtaining a specified Service	GET /api/v1/namespaces/{namespace}/services/{name}
	Replacing a specified Service	PUT /api/v1/namespaces/{namespace}/services/{name}
	Listing Services in a specified namespace	GET /api/v1/namespaces/{namespace}/services
	Listing Services	GET /api/v1/services

API	Function	URI
	Updating a specified Service	PATCH /api/v1/namespaces/{namespace}/services/{name}
Ingress	Creating an ingress	<p>POST /apis/networking.k8s.io/v1/namespaces/{namespace}/ingresses (for clusters of v1.21 and later)</p> <p>POST /apis/networking.k8s.io/v1beta1/namespaces/{namespace}/ingresses (for clusters from v1.15 to v1.21)</p> <p>POST /apis/extensions/v1beta1/namespaces/{namespace}/ingresses (for clusters earlier than v1.15)</p>
	Updating a specified ingress	<p>PATCH /apis/networking.k8s.io/v1/namespaces/{namespace}/ingresses/{name} (for clusters of v1.21 and later)</p> <p>PATCH /apis/networking.k8s.io/v1beta1/namespaces/{namespace}/ingresses/{name} (for clusters from v1.15 to v1.21)</p> <p>PATCH /apis/extensions/v1beta1/namespaces/{namespace}/ingresses/{name} (for clusters earlier than v1.15)</p>
	Replacing a specified ingress	<p>PUT /apis/networking.k8s.io/v1/namespaces/{namespace}/ingresses/{name} (for clusters of v1.21 and later)</p> <p>PUT /apis/networking.k8s.io/v1beta1/namespaces/{namespace}/ingresses/{name} (for clusters from v1.15 to v1.21)</p> <p>PUT /apis/extensions/v1beta1/namespaces/{namespace}/ingresses/{name} (for clusters earlier than v1.15)</p>
	Deleting an ingress	<p>DELETE /apis/networking.k8s.io/v1/namespaces/{namespace}/ingresses/{name} (for clusters of v1.21 and later)</p> <p>DELETE /apis/networking.k8s.io/v1beta1/namespaces/{namespace}/ingresses/{name} (for clusters from v1.15 to v1.21)</p> <p>DELETE /apis/extensions/v1beta1/namespaces/{namespace}/ingresses/{name} (for clusters earlier than v1.15)</p>

API	Function	URI
	Deleting all ingresses	DELETE /apis/networking.k8s.io/v1/namespaces/{namespace}/ingresses (for clusters of v1.21 and later) DELETE /apis/networking.k8s.io/v1beta1/namespaces/{namespace}/ingresses (for clusters from v1.15 to v1.21) DELETE /apis/extensions/v1beta1/namespaces/{namespace}/ingresses (for clusters earlier than v1.15)
	Obtaining a specified ingress	GET /apis/networking.k8s.io/v1/namespaces/{namespace}/ingresses/{name} (for clusters of v1.21 and later) GET /apis/networking.k8s.io/v1beta1/namespaces/{namespace}/ingresses/{name} (for clusters from v1.15 to v1.21) GET /apis/extensions/v1beta1/namespaces/{namespace}/ingresses/{name} (for clusters earlier than v1.15)
	Listing ingresses in a specified namespace	GET /apis/networking.k8s.io/v1/namespaces/{namespace}/ingresses (for clusters of v1.21 and later) GET /apis/networking.k8s.io/v1beta1/namespaces/{namespace}/ingresses (for clusters from v1.15 to v1.21) GET /apis/extensions/v1beta1/namespaces/{namespace}/ingresses (for clusters earlier than v1.15)
	Listing ingresses	GET /apis/networking.k8s.io/v1/ingresses (for clusters of v1.21 and later) GET /apis/networking.k8s.io/v1beta1/ingresses (for clusters from v1.15 to v1.21) GET /apis/extensions/v1beta1/ingresses (for clusters earlier than v1.15)
	Obtaining the status of an ingress in a specified namespace	GET /apis/networking.k8s.io/v1/namespaces/{namespace}/ingresses/{name}/status (for clusters of v1.21 and later) GET /apis/networking.k8s.io/v1beta1/namespaces/{namespace}/ingresses/{name}/status (for clusters from v1.15 to v1.21) GET /apis/extensions/v1beta1/namespaces/{namespace}/ingresses/{name}/status (for clusters earlier than v1.15)

API	Function	URI
	Replacing the status of an ingress in a specified namespace	<p>PUT /apis/networking.k8s.io/v1/namespaces/{namespace}/ingresses/{name}/status (for clusters of v1.21 and later)</p> <p>PUT /apis/networking.k8s.io/v1beta1/namespaces/{namespace}/ingresses/{name}/status (for clusters from v1.15 to v1.21)</p> <p>PUT /apis/extensions/v1beta1/namespaces/{namespace}/ingresses/{name}/status (for clusters earlier than v1.15)</p>
	Updating the status of an ingress in a specified namespace	<p>PATCH /apis/networking.k8s.io/v1/namespaces/{namespace}/ingresses/{name}/status (for clusters of v1.21 and later)</p> <p>PATCH /apis/networking.k8s.io/v1beta1/namespaces/{namespace}/ingresses/{name}/status (for clusters from v1.15 to v1.21)</p> <p>PATCH /apis/extensions/v1beta1/namespaces/{namespace}/ingresses/{name}/status (for clusters earlier than v1.15)</p>
Network Policy	Creating a network policy	POST /apis/networking.k8s.io/v1/namespaces/{namespace}/networkpolicies
	Updating a specified network policy	PATCH /apis/networking.k8s.io/v1/namespaces/{namespace}/networkpolicies/{name}
	Replacing a specified network policy	PUT /apis/networking.k8s.io/v1/namespaces/{namespace}/networkpolicies/{name}
	Deleting a specified network policy	DELETE /apis/networking.k8s.io/v1/namespaces/{namespace}/networkpolicies/{name}
	Deleting network policies in batches	DELETE /apis/networking.k8s.io/v1/namespaces/{namespace}/networkpolicies
	Reading a specified network policy	GET /apis/networking.k8s.io/v1/namespaces/{namespace}/networkpolicies/{name}
	Listing network policies in a specified namespace	GET /apis/networking.k8s.io/v1/namespaces/{namespace}/networkpolicies

API	Function	URI
	Listing all network policies	GET /apis/networking.k8s.io/v1/networkpolicies
PersistentVolume	Creating a PersistentVolume	POST /api/v1/persistentvolumes
	Deleting a specified PersistentVolume	DELETE /api/v1/persistentvolumes/{name}
	Deleting all PersistentVolumes	DELETE /api/v1/persistentvolumes
	Reading a specified PersistentVolume	GET /api/v1/persistentvolumes/{name}
	Replacing a specified PersistentVolume	PUT /api/v1/persistentvolumes/{name}
	Replacing the status of a specified PersistentVolume	PUT /api/v1/persistentvolumes/{name}/status
	Listing all PersistentVolumes	GET /api/v1/persistentvolumes
	Updating a specified PersistentVolume	PATCH /api/v1/persistentvolumes/{name}
PersistentVolumeClaim	Creating a PersistentVolumeClaim	POST /api/v1/namespaces/{namespace}/persistentvolumeclaims
	Deleting a specified PersistentVolumeClaim	DELETE /api/v1/namespaces/{namespace}/persistentvolumeclaims/{name}
	Deleting all PersistentVolume-Claims	DELETE /api/v1/namespaces/{namespace}/persistentvolumeclaims
	Reading a specified PersistentVolumeClaim	GET /api/v1/namespaces/{namespace}/persistentvolumeclaims/{name}
	Replacing a specified PersistentVolumeClaim	PUT /api/v1/namespaces/{namespace}/persistentvolumeclaims/{name}
	Replacing the status of a specified PersistentVolumeClaim	PUT /api/v1/namespaces/{namespace}/persistentvolumeclaims/{name}/status
	Listing PersistentVolume-Claims in a specified namespace	GET /api/v1/namespaces/{namespace}/persistentvolumeclaims

API	Function	URI
	Listing all PersistentVolume-Claims	GET /api/v1/persistentvolumeclaims
	Updating a specified PersistentVolumeClaim	PATCH /api/v1/namespaces/{namespace}/persistentvolumeclaims/{name}
ConfigMap	Creating a ConfigMap	POST /api/v1/namespaces/{namespace}/configmaps
	Deleting a ConfigMap	DELETE /api/v1/namespaces/{namespace}/configmaps/{name}
	Deleting all ConfigMaps	DELETE /api/v1/namespaces/{namespace}/configmaps
	Reading a specified ConfigMap	GET /api/v1/namespaces/{namespace}/configmaps/{name}
	Replacing a specified ConfigMap	PUT /api/v1/namespaces/{namespace}/configmaps/{name}
	Listing ConfigMaps in a specified namespace	GET /api/v1/namespaces/{namespace}/configmaps
	Listing all ConfigMaps	GET /api/v1/configmaps
	Updating a specified ConfigMap	PATCH /api/v1/namespaces/{namespace}/configmaps/{name}
Secret	Creating a secret	POST /api/v1/namespaces/{namespace}/secrets
	Deleting a secret	DELETE /api/v1/namespaces/{namespace}/secrets/{name}
	Deleting all secrets in a specified namespace	DELETE /api/v1/namespaces/{namespace}/secrets
	Reading a specified secret	GET /api/v1/namespaces/{namespace}/secrets/{name}
	Replacing a specified secret	PUT /api/v1/namespaces/{namespace}/secrets/{name}
	Listing the secrets in a namespace	GET /api/v1/namespaces/{namespace}/secrets
	Listing secrets in a cluster	GET /api/v1/secrets
RBAC/ClusterRole	Creating a ClusterRole	POST /apis/rbac.authorization.k8s.io/v1/clusterroles
	Updating a specified ClusterRole	PATCH /apis/rbac.authorization.k8s.io/v1/clusterroles/{name}



API	Function	URI
	Replacing a specified ClusterRole	PUT /apis/rbac.authorization.k8s.io/v1/clusterroles/{name}
	Deleting a specified ClusterRole	DELETE /apis/rbac.authorization.k8s.io/v1/clusterroles/{name}
	Deleting ClusterRoles in batches	DELETE /apis/rbac.authorization.k8s.io/v1/clusterroles
	Reading a specified ClusterRole	GET /apis/rbac.authorization.k8s.io/v1/clusterroles/{name}
	Listing ClusterRoles	GET /apis/rbac.authorization.k8s.io/v1/clusterroles
RBAC/ ClusterRoleBinding	Creating a ClusterRoleBinding	POST /apis/rbac.authorization.k8s.io/v1/clusterrolebindings
	Updating a specified ClusterRoleBinding	PATCH /apis/rbac.authorization.k8s.io/v1/clusterrolebindings/{name}
	Replacing a specified ClusterRoleBinding	PUT /apis/rbac.authorization.k8s.io/v1/clusterrolebindings/{name}
	Deleting a specified ClusterRoleBinding	DELETE /apis/rbac.authorization.k8s.io/v1/clusterrolebindings/{name}
	Deleting ClusterRoleBindings in batches	DELETE /apis/rbac.authorization.k8s.io/v1/clusterrolebindings
	Reading a specified ClusterRoleBinding	GET /apis/rbac.authorization.k8s.io/v1/clusterrolebindings/{name}
	Listing ClusterRoleBindings	GET /apis/rbac.authorization.k8s.io/v1/clusterrolebindings
RBAC/ Role	Creating a Role	POST /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/roles
	Updating a specified Role	PATCH /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/roles/{name}
	Replacing a specified Role	PUT /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/roles/{name}
	Deleting a specified Role	DELETE /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/roles/{name}
	Deleting Roles in batches	DELETE /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/roles
	Reading a specified Role	GET /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/roles/{name}

API	Function	URI
	Listing Roles in a specified namespace	GET /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/roles
	Listing all Roles	GET /apis/rbac.authorization.k8s.io/v1/roles
RBAC/ RoleBinding	Creating a RoleBinding	POST /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/rolebindings
	Updating a specified RoleBinding	PATCH /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/rolebindings/{name}
	Replacing a specified RoleBinding	PUT /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/rolebindings/{name}
	Deleting a specified RoleBinding	DELETE /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/rolebindings/{name}
	Deleting RoleBindings in batches	DELETE /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/rolebindings
	Reading a specified RoleBinding	GET /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/rolebindings/{name}
	Listing RoleBindings in a specified namespace	GET /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/rolebindings
	Listing all RoleBindings	GET /apis/rbac.authorization.k8s.io/v1/rolebindings
API groups	Listing APIVersions	GET /api
	Listing APIGroups	GET /apis
	listing APIResources of GroupVersion apiregistration.k8s.io/v1beta1	GET /apis/apiregistration.k8s.io/v1beta1
	listing APIResources of GroupVersion extensions/v1beta1	GET /apis/extensions/v1beta1
	listing APIResources of GroupVersion apps/v1&apps/v1beta1	GET /apis/apps/v1 (for clusters later than v1.15) GET /apis/apps/v1beta1 (for clusters of v1.15 and earlier)
	listing APIResources of GroupVersion authentication.k8s.io/v1	GET /apis/authentication.k8s.io/v1

API	Function	URI
	listing APIResources of GroupVersion authentication.k8s.io/v1beta1	GET /apis/authentication.k8s.io/v1beta1
	listing APIResources of GroupVersion authorization.k8s.io/v1	GET /apis/authorization.k8s.io/v1
	listing APIResources of GroupVersion authorization.k8s.io/v1beta1	GET /apis/authorization.k8s.io/v1beta1
	listing APIResources of GroupVersion autoscaling/v1	GET /apis/autoscaling/v1
	listing APIResources of GroupVersion batch/v1	GET /apis/batch/v1
	listing APIResources of GroupVersion certificates.k8s.io/v1beta1	GET /apis/certificates.k8s.io/v1beta1
	listing APIResources of GroupVersion networking.k8s.io/v1	GET /apis/networking.k8s.io/v1
	listing APIResources of GroupVersion policy/v1beta1	GET /apis/policy/v1beta1
	listing APIResources of GroupVersion rbac.authorization.k8s.io/v1beta1	GET /apis/rbac.authorization.k8s.io/v1beta1
	listing APIResources of GroupVersion storage.k8s.io/v1	GET /apis/storage.k8s.io/v1
	listing APIResources of GroupVersion storage.k8s.io/v1beta1	GET /apis/storage.k8s.io/v1beta1
	listing APIResources of GroupVersion apiextensions.k8s.io/v1beta1	GET /apis/apiextensions.k8s.io/v1beta1
	listing APIResources of GroupVersion v1	GET /api/v1

API	Function	URI
Event	Reading events	GET /api/v1/events
	Listing events in a specified namespace	GET /api/v1/namespaces/{namespace}/events

# 3 Calling APIs

## 3.1 Making an API Request

This section describes the structure of a REST API request, and uses the IAM API for **obtaining a user token** as an example to demonstrate how to call an API. The obtained token can then be used to authenticate the calling of other APIs.

### Request URI

A request URI is in the following format:

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

Although a request URI is included in the request header, most programming languages or frameworks require the request URI to be transmitted separately.


**Table 3-1** URI parameters

Parameter	Description
URI-scheme	Protocol used to transmit requests. All APIs use HTTPS.
Endpoint	Domain name or IP address of the server bearing the REST service. The endpoint varies between services in different regions. It can be obtained from <b>Regions and Endpoints</b> . For example, the endpoint of IAM in the <b>my-kualalumpur-1</b> region is <b>iam.my-kualalumpur-1.myhuaweicloud.com</b> .
resource-path	Access path of an API for performing an operation. Obtain the path from the URI of an API. For example, the <b>resource-path</b> of the API used to obtain a user token is <b>/v3/auth/tokens</b> .
query-string	Query parameter, which is optional. Ensure that a question mark (?) is included before each query parameter that is in the format of <i>Parameter name=Parameter value</i> . For example, <b>?limit=10</b> indicates that a maximum of 10 data records will be displayed.

For example, to obtain an IAM token in the **AP-Kuala Lumpur-OP6** region, obtain the endpoint of IAM (**iam.my-kualalumpur-1.myhuaweicloud.com**) for this region and the **resource-path** (**/v3/auth/tokens**) in the URI of the API used to **obtain a user token**. Then, construct the URI as follows:

```
https://iam.my-kualalumpur-1.myhuaweicloud.com/v3/auth/tokens
```

**Figure 3-1** Example URI



**NOTE**

To simplify the URI display in this document, each API is provided only with a **resource-path** and a request method. The **URI-scheme** of all APIs is **HTTPS**, and the endpoints of all APIs in the same region are identical.

## Request Methods

The HTTP protocol defines the following request methods that can be used to send a request to the server.

**Table 3-2** HTTP methods

Method	Description
GET	Requests the server to return specified resources.
PUT	Requests the server to update specified resources.
POST	Requests the server to add resources or perform special operations.
DELETE	Requests the server to delete specified resources, for example, an object.
HEAD	Requests the server to return the response header.
PATCH	Requests the server to update partial content of a specified resource. If the resource does not exist, a new resource will be created.

For example, in the case of the API used to **obtain a user token**, the request method is **POST**. The request is as follows:

```
POST https://iam.my-kualalumpur-1.myhuaweicloud.com/v3/auth/tokens
```

## Request Header

You can also add additional header fields to a request, such as the fields required by a specified URI or HTTP method. For example, to request for the authentication information, add **Content-Type**, which specifies the request body type.

Common request header fields are as follows.

**Table 3-3** Common request header fields

Parameter	Description	Mandatory	Example Value
Host	Specifies the server domain name and port number of the resources being requested. The value can be obtained from the URL of the service API. The value is in the format of <i>Hostname:Port number</i> . If the port number is not specified, the default port is used. The default port number for <b>https</b> is <b>443</b> .	No This field is mandatory for AK/SK authentication.	code.test.com or code.test.com:443
Content-Type	Specifies the type (or format) of the message body. The default value <b>application/json</b> is recommended. Other values of this field will be provided for specific APIs if any.	Yes	application/json
Content-Length	Specifies the length of the request body. The unit is byte.	No	3495
X-Project-Id	Specifies the project ID. Obtain the project ID by following the instructions in <a href="#">Obtaining a Project ID</a> .	No	e9993fc787d94b6c886cb aa340f9c0f4

Parameter	Description	Mandatory	Example Value
X-Auth-Token	Specifies the user token. It is a response to the API for <b>obtaining a user token</b> (This is the only API that does not require authentication). After the request is processed, the value of <b>X-Subject-Token</b> in the response header is the token value.	No This field is mandatory for token authentication.	The following is part of an example token: MIIPAgYJKoZlhvcNAQc-Co...ggg1BBIINPXsidG9rZ

 **NOTE**

In addition to supporting authentication using tokens, APIs support authentication using AK/SK, which uses SDKs to sign a request. During the signature, the **Authorization** (signature authentication) and **X-Sdk-Date** (time when a request is sent) headers are automatically added in the request.

For more details, see "Authentication Using AK/SK" in [Authentication](#).

The API used to **obtain a user token** does not require authentication. Therefore, only the **Content-Type** field needs to be added to requests for calling the API. An example of such requests is as follows:

```
POST https://iam.my-kualalumpur-1.myhuaweicloud.com/v3/auth/tokens
Content-Type: application/json
```

### (Optional) Request Body

This part is optional. The body of a request is often sent in a structured format as specified in the **Content-Type** header field. The request body transfers content except the request header.

The request body varies between APIs. Some APIs do not require the request body, such as the APIs requested using the GET and DELETE methods.

In the case of the API used to **obtain a user token**, the request parameters and parameter description can be obtained from the API request. The following provides an example request with a body included. Replace **username**, **domainname**, **\*\*\*\*\*** (login password), and **xxxxxxxxxxxxxxxxxxxx** (project name) with the actual values. Obtain a project name from [Regions and Endpoints](#).

 **NOTE**

The **scope** parameter specifies where a token takes effect. You can set **scope** to an account or a project under an account. In the following example, the token takes effect only for the resources in a specified project. For more information about this API, see [Obtaining a User Token](#).



```

POST https://iam.my-kualalumpur-1.myhuaweicloud.com/v3/auth/tokens
Content-Type: application/json

{
  "auth": {
    "identity": {
      "methods": [
        "password"
      ],
      "password": {
        "user": {
          "name": "username",
          "password": "*****",
          "domain": {
            "name": "domainname"
          }
        }
      }
    },
    "scope": {
      "project": {
        "name": "xxxxxxxxxxxxxxxxxxxxx"
      }
    }
  }
}

```

If all data required for the API request is available, you can send the request to call the API through [curl](#), [Postman](#), or coding. In the response to the API used to obtain a user token, **x-subject-token** is the desired user token. This token can then be used to authenticate the calling of other APIs.

## 3.2 Authentication

Requests for calling an API can be authenticated using either of the following methods:

- Token authentication: Requests are authenticated using tokens.
- AK/SK authentication: Requests are encrypted using an AK/SK. AK/SK-based authentication is recommended because it is more secure than token-based authentication.

### Token-based Authentication

#### NOTE

The validity period of a token is 24 hours. When using a token for authentication, cache it to prevent frequently calling the IAM API used to obtain a user token.

A token specifies temporary permissions in a computer system. During API authentication using a token, the token is added to requests to get permissions for calling the API.

You can obtain a token by calling the [Obtaining a User Token](#) API. When you call the API, set **auth.scope** in the request body to **project**.

```

{
  "auth": {
    "identity": {
      "methods": [

```

```

    "password"
  ],
  "password": {
    "user": {
      "name": "username",
      "password": "*****",
      "domain": {
        "name": "domainname"
      }
    }
  }
},
"scope": {
  "project": {
    "name": "xxxxxxxx"
  }
}
}
}

```

After a token is obtained, the **X-Auth-Token** header field must be added to requests to specify the token when calling other APIs. For example, if the token is **ABCDEFJ....**, **X-Auth-Token: ABCDEFJ....** can be added to a request as follows:

```

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

```

## AK/SK Authentication

### NOTE

AK/SK authentication supports API requests with a body not larger than 12 MB. For API requests with a larger body, token authentication is recommended.

In AK/SK authentication, AK/SK is used to sign requests and the signature is then added to the requests for authentication.

- AK: access key ID, which is a unique identifier used in conjunction with a secret access key to sign requests cryptographically.
- SK: secret access key, which is used in conjunction with an AK to sign requests cryptographically. It identifies a request sender and prevents the request from being modified.

In AK/SK authentication, you can use an AK/SK to sign requests based on the signature algorithm or using the signing SDK. For details about how to sign requests and use the signing SDK, see [API Request Signing Guide](#).

### NOTE

The signing SDK is only used for signing requests and is different from the SDKs provided by services.

## 3.3 Response

### Status Code

After sending a request, you will receive a response, including a status code, response header, and response body.

A status code is a group of digits, ranging from 1xx to 5xx. It indicates the status of a request. For more information, see [Status Code](#).

For example, if status code **201** is returned for calling the API used to [obtain a user token](#), the request is successful.

## Response Header

Similar to a request, a response also has a header, for example, **Content-Type**.

[Figure 3-2](#) shows the response header fields for the API used to [obtain a user token](#). The **x-subject-token** header field is the desired user token. This token can then be used to authenticate the calling of other APIs.

**Figure 3-2** Header fields of the response to the request for obtaining a user token

```

connection → keep-alive
content-type → application/json
date → Tue, 12 Feb 2019 06:52:13 GMT
server → Web Server
strict-transport-security → max-age=31536000; includeSubdomains;
transfer-encoding → chunked
via → proxy A
x-content-type-options → nosniff
x-download-options → noopen
x-frame-options → SAMEORIGIN
x-lam-trace-id → 218d45ab-d674-4995-af3a-2d0255ba41b5
x-subject-token → [REDACTED]
x-xss-protection → 1; mode=block;

```

## Response Body

The body of a response is often returned in structured format as specified in the **Content-Type** header field. The response body transfers content except the response header.

The following is part of the response body for the API used to [obtain a user token](#).

```

{
  "token": {
    "expires_at": "2019-02-13T06:52:13.855000Z",
    "methods": [
      "password"
    ],
    "catalog": [
      {
        "endpoints": [
          {
            "region_id": "",

```

If an error occurs during API calling, an error code and a message will be displayed. The following shows an error response body.

```
{
  "error_msg": "The format of message is error",
  "error_code": "AS.0001"
}
```

In the response body, **error\_code** is an error code, and **error\_msg** provides information about the error.

# 4 APIs

## 4.1 API URL

- The URL format for cluster, node, node pool, add-on, and quota management is **https://Endpoint/uri**. *uri* indicates the resource path, that is, the API access path.
- The URL format for Kubernetes APIs, storage management, and add-on management is **https://{clusterid}.Endpoint/uri**. In the URL, *{clusterid}* indicates the cluster ID, and *uri* indicates the resource path, that is, the path for API access.

 NOTE

- The format of the URL called by the add-on management APIs is **https://{clusterid}.Endpoint/uri**. However, *{clusterid}* is used only for the domain name and is not verified or used by the APIs. Set *{clusterid}* in the query or body. For details about *{clusterid}*, see the add-on management sections.
- *{clusterid}* is required for Kubernetes APIs and storage management, which indicates the cluster that needs to be accessed by calling the API.

**Table 4-1** URL parameters

Parameter	Description
{clusterid}	Cluster ID. After a cluster is created, call the <a href="#">API for obtaining a cluster in a specified project</a> to obtain the cluster ID.
Endpoint	URL that is the entry point for a web service. You can obtain it from <a href="#">Endpoints</a> .
uri	Access path of an API for performing an operation. Obtain the path from the URI of an API. For example, the <b>resource-path</b> of the API used to obtain a user token is <b>v3/auth/tokens</b> .

## 4.2 Cluster Management

### 4.2.1 Creating a Cluster

#### Function

This API is used to create an empty cluster, which has only master nodes but no worker nodes. After creating a cluster by calling this API, you can add nodes by [creating nodes](#).

#### NOTE

- The URL for cluster management is in the format of **https://Endpoint/uri**, in which *uri* indicates the resource path, that is, the path for API access.
- By default, ICAgent is not installed when you call this API to create a cluster. If you need to install ICAgent, add "cluster.install.addons.external/install":[{"addonTemplateName":"icagent"}]" to **annotations** in the request body. ICAgent will be automatically installed during cluster creation. ICAgent is an O&M data collection agent used by Application Performance Management (APM). It runs on each server to collect data from probes in real time. ICAgent is the prerequisite for achieving application O&M. If ICAgent is not installed, the application O&M functions cannot be used.

#### Constraints

Before calling the CCE API to create a cluster, ensure that the following conditions are met:

- A VPC is available. Otherwise, the cluster cannot be created. If a VPC is available, you do not need to create a new one. A VPC provides an isolated, configurable, manageable virtual network environment for CCE clusters.
- Before creating a cluster, you should properly plan the container and Service CIDR blocks. After a cluster using a container tunnel network is created, you cannot modify the CIDR blocks. After a cluster using a VPC network or Cloud Native 2.0 network is created, you can only add CIDR blocks to the cluster but cannot modify the existing CIDR blocks. To modify the existing CIDR blocks, you need to create a new cluster. Exercise caution when performing this operation.
- An agency has been correctly created and is not deleted. If the agency verification fails, the cluster fails to be created. You can log in to the CCE console to see if there is any agency. If no agency is created, the system prompts you to create one. If an agency has been created, no message is displayed.
- By default, an account can create up to five clusters in each region. If you need to create more clusters, you can submit an application to increase the quota.

#### URI

POST /api/v3/projects/{project\_id}/clusters

**Table 4-2** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

## Request Parameters

**Table 4-3** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

**Table 4-4** Request body parameters

Parameter	Mandatory	Type	Description
kind	Yes	String	API type. The value is fixed at <b>Cluster</b> or <b>cluster</b> and cannot be changed.
apiVersion	Yes	String	API version. The value is fixed at <b>v3</b> and cannot be changed.
metadata	Yes	<a href="#">ClusterMetadata</a> object	Basic information about a cluster. Metadata is a collection of attributes.
spec	Yes	<a href="#">ClusterSpec</a> object	Detailed description of the cluster. CCE creates or updates objects by defining or updating <b>spec</b> .
status	No	<a href="#">ClusterStatus</a> object	Cluster status and job ID of the cluster creation job.

**Table 4-5** ClusterMetadata

Parameter	Mandatory	Type	Description
name	Yes	String	Cluster name. Enter 4 to 128 characters, starting with a lowercase letter and not ending with a hyphen (-). Only lowercase letters, digits, and hyphens (-) are allowed.
uid	No	String	Cluster ID, which uniquely identifies a cluster. This ID is automatically generated after a cluster is created. Only the automatically generated ID will take effect.
alias	No	String	Alias of a cluster name displayed on the CCE console, and the name can be changed. Enter 4 to 128 characters starting with a letter and not ending with a hyphen (-). Only digits, letters, and hyphens (-) are allowed. A cluster alias must be unique. In the request body for creating or updating a cluster, if the cluster alias is not specified or is left blank, the alias of the cluster is the same as the cluster name. In the response body for obtaining a cluster, if the cluster alias is not configured, the cluster alias will not be returned.



Parameter	Mandatory	Type	Description
annotations	No	Map<String,String>	<p>Cluster annotations, in the format of key-value pairs.</p> <pre>"annotations": {   "key1" : "value1",   "key2" : "value2" }</pre> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• <b>annotations:</b> Does not label or select objects. The metadata in <b>annotations</b> may be small or large, structured or unstructured, and may include characters that are not allowed in labels.</li> <li>• This field is not stored in the database and is used only to specify the add-ons to be installed in the cluster.</li> <li>• Install ICAgent during cluster creation by adding the key-value pair "cluster.install.addons.external / install":{"addonTemplateName":"icagent"}".</li> </ul>
labels	No	Map<String,String>	<p>Cluster labels, in the format of key-value pairs.</p> <p><b>NOTE</b></p> <p>The value of this field is automatically generated by the system and is used by the frontend to identify the features supported by the cluster during the upgrade. Customized values are invalid.</p>
creationTimestamp	No	String	Time when the cluster was created.
updateTimestamp	No	String	Time when the cluster was updated.

**Table 4-6** ClusterSpec

Parameter	Mandatory	Type	Description
category	No	String	Cluster type. Options: <ul style="list-style-type: none"> <li>• <b>CCE</b>: CCE cluster CCE cluster supports hybrid deployment of VMs and BMSs, and heterogeneous nodes such as GPU and NPU nodes, allowing you to run your containers in a secure and stable container runtime environment based on a high-performance network model.</li> </ul>
type	No	String	Master node architecture: <ul style="list-style-type: none"> <li>• <b>VirtualMachine</b>: x86</li> </ul>

Parameter	Mandatory	Type	Description
flavor	Yes	String	<p>Cluster specifications. Specifications of clusters 1.15 or later versions can be changed after they are created. For details, see <a href="#">Modifying Cluster Specifications</a>. Options:</p> <ul style="list-style-type: none"> <li>• <b>cce.s1.small</b>: a small-scale CCE cluster with one master node and a maximum of 50 worker nodes</li> <li>• <b>cce.s1.medium</b>: a medium-scale CCE cluster with one master node and a maximum of 200 worker nodes</li> <li>• <b>cce.s2.small</b>: a small-scale CCE cluster with three master nodes and a maximum of 50 worker nodes</li> <li>• <b>cce.s2.medium</b>: a medium-scale CCE cluster with three master nodes and a maximum of 200 worker nodes</li> <li>• <b>cce.s2.large</b>: a large-scale CCE cluster with three master nodes and a maximum of 1,000 worker nodes</li> <li>• <b>cce.s2.xlarge</b>: an ultra-large-scale CCE cluster with three master nodes and a maximum of 2,000 worker nodes</li> </ul>

Parameter	Mandatory	Type	Description
			<p><b>NOTE</b> The fields in the parameters are described as follows:</p> <ul style="list-style-type: none"> <li>• <b>s1</b>: specifies a cluster with one master node. If the master node is faulty, the cluster will become unavailable, but running workloads in the cluster are not affected.</li> <li>• <b>s2</b>: specifies an HA cluster with three master nodes. If one of the master nodes is faulty, the cluster is still available.</li> <li>• <b>small</b>: specifies that a cluster can manage a maximum of 50 worker nodes.</li> <li>• <b>medium</b>: specifies that a cluster can manage a maximum of 200 worker nodes.</li> <li>• <b>large</b>: specifies that a cluster can manage a maximum of 1,000 worker nodes.</li> <li>• <b>xlarge</b>: specifies that a cluster can manage a maximum of 2,000 worker nodes.</li> </ul>

Parameter	Mandatory	Type	Description
version	No	String	<p>Version of a cluster, which mirrors the baseline version of the Kubernetes community. The latest version is recommended.</p> <p>You can create clusters of the latest three versions on the CCE console. To learn which cluster versions are available, log in to the CCE console, create a cluster, and check supported cluster versions. You can call APIs to create clusters of other versions. However, these cluster versions will be gradually terminated. For details about the support policy, see the CCE announcement.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• If not specified, a cluster of the latest version will be created.</li> <li>• If a baseline cluster version is specified but the R version is not specified, a cluster of the latest R version will be created by default. It is a good practice not to specify the R version.</li> </ul>

Parameter	Mandatory	Type	Description
platformVersion	No	String	<p>CCE cluster platform version, indicating the internal version under the cluster version (<b>version</b>). Platform versions are used to trace iterations in a major cluster version. They are unique within a major cluster version and recounted when the major cluster version changes. This parameter cannot be customized. When you create a cluster, the latest corresponding platform version is automatically selected.</p> <p>The format of <b>platformVersion</b> is <b>cce.X.Y</b>.</p> <ul style="list-style-type: none"> <li>• <b>X</b>: internal feature version, indicating changes in features, patches, or OS support in the cluster version. The value starts from <b>1</b> and increases monotonically.</li> <li>• <b>Y</b>: patch version of an internal feature version. It is used only for software package update after the feature version is brought online. No other modification is involved. The value starts from <b>0</b> and increases monotonically.</li> </ul>
legacyVersion	No	String	<p>(Discarded) Outdated version of CCE clusters, which serves no practical purpose and is only used to showcase the cluster version and platform version combination. The version number is globally unique. For example, if the cluster version is va.b and the platformVersion is cce.X.Y, the value of legacyVersion is va.b.X-rY.</p>

Parameter	Mandatory	Type	Description
description	No	String	Cluster description, for example, which purpose the cluster is intended to serve. By default, this field is left unspecified. To modify cluster description after the cluster is created, call the <a href="#">API for updating cluster information</a> or go to the cluster details page on the CCE console. Only UTF-8 encoding is supported.
customSan	No	Array of strings	Custom SAN field in the server certificate of the cluster API server, which must comply with the SSL and X509 format specifications. <ol style="list-style-type: none"> <li>1. Duplicate names are not allowed.</li> <li>2. Must comply with the IP address and domain name formats.</li> </ol> Example: SAN 1: DNS Name=example.com SAN 2: DNS Name=www.example.com SAN 3: DNS Name=example.net SAN 4: IP Address=93.184.216.34
ipv6enable	No	Boolean	Whether a cluster supports IPv6 addresses. It is supported by clusters 1.15 and later.  <b>NOTE</b> After IPv6 is enabled, iptables is not supported. The VPC network model does not support IPv4/IPv6 dual-stack.
hostNetwork	Yes	<a href="#">HostNetwork</a> object	Node networking parameters, including VPC and subnet ID. This field is mandatory because nodes in a cluster communicate with each other by using a VPC.
containerNetwork	Yes	<a href="#">ContainerNetwork</a> object	Container networking parameters, including the container network model and container CIDR block.
serviceNetwork	No	<a href="#">ServiceNetwork</a> object	Service CIDR block, including IPv4 CIDR blocks.

Parameter	Mandatory	Type	Description
publicAccess	No	<b>PublicAccess</b> object	Cluster API access control.
authentication	No	<b>Authentication</b> object	Configurations of the cluster authentication mode.
billingMode	No	Integer	Billing mode of a cluster. <ul style="list-style-type: none"> <li>• <b>0</b>: pay-per-use</li> </ul> Defaults to pay-per-use.
masters	No	Array of <b>MasterSpec</b> objects	Advanced configurations of the master nodes. If this parameter is not configured, no value is returned.
kubernetesSvcIpRange	No	String	Service CIDR blocks for Kubernetes clusterIPs. This field is available only for clusters of v1.11.7 and later. If this parameter is not specified during cluster creation, the default value <b>10.247.0.0/16</b> will be used. This parameter is deprecated. Use <b>serviceNetwork</b> instead. The new field contains the IPv4 CIDR blocks.
clusterTags	No	Array of <b>ResourceTag</b> objects	Cluster resource tags.



Parameter	Mandatory	Type	Description
kubeProxyMode	No	String	Service forwarding mode. Options: <ul style="list-style-type: none"> <li>• <b>iptables</b>: Traditional kube-proxy uses iptables rules to implement Service load balancing. In this mode, too many iptables rules will be generated when many Services are deployed. In addition, non-incremental updates will cause latency and even tangible performance issues in the case of service traffic spikes.</li> <li>• <b>ipvs</b>: Optimized kube-proxy mode with higher throughput and faster speed. This mode supports incremental updates and can keep connections uninterrupted during Service updates. It is suitable for large-sized clusters.</li> </ul> <p><b>NOTE</b> <b>iptables</b> is used by default.</p>
az	No	String	AZ. This field is returned only for a query.
extendParam	No	<b>ClusterExtendedParam</b> object	Extended field to decide whether the cluster will span across AZs or belong to a specified enterprise project, or whether a dedicated CCE cluster is to be created.
supportIstio	No	Boolean	Whether Istio is supported.
configurationsOverride	No	Array of <b>PackageConfiguration</b> objects	Cluster default component configuration override. If you specify a component or parameter that is not supported, the configuration item will be ignored.

**Table 4-7** HostNetwork

Parameter	Mandatory	Type	Description
vpc	Yes	String	<p>ID of the VPC used to create a master node.</p> <p>You can obtain it in either of the following ways:</p> <ul style="list-style-type: none"> <li>• Method 1: Log in to the VPC console and view the VPC ID in the VPC details page.</li> <li>• Method 2: Use the VPC API to obtain VPC IDs.</li> </ul>
subnet	Yes	String	<p>Network ID of the subnet used to create a master node. You can obtain it in either of the following ways:</p> <ul style="list-style-type: none"> <li>• Method 1: Log in to the VPC console and click the target subnet on the <b>Subnets</b> page. You can view the network ID on the displayed page.</li> <li>• Method 2: Use the VPC API to obtain subnets.</li> </ul>
SecurityGroup	No	String	<p>Default worker node security group ID of the cluster. If specified, the cluster will be bound to the target security group. Otherwise, the system will automatically create a default worker node security group for you. The default worker node security group needs to allow access from certain ports to ensure normal communications.</p>
controlPlaneSecurityGroup	No	String	<p>Security group ID of the cluster control plane node, which is automatically generated after the cluster is created. A user-defined value will not take effect.</p>

**Table 4-8** ContainerNetwork

Parameter	Mandatory	Type	Description
mode	Yes	String	<p>Container network model. Select one of the following possible values:</p> <ul style="list-style-type: none"> <li>• <b>overlay_l2</b>: an overlay_l2 network (container tunnel network) built for containers by using OpenVSwitch (OVS).</li> <li>• <b>vpc-router</b>: an underlay_l2 network built for containers by using IPvlan and custom VPC routes.</li> </ul>
cidr	No	String	<p>Container CIDR block. Recommended: 10.0.0.0/12-19, 172.16.0.0/16-19, or 192.168.0.0/16-19. If the selected CIDR block conflicts with existing ones, an error will be reported.</p> <p>Not editable after the cluster is created. (This parameter has been discarded. A specified <b>cidrs</b> will make <b>cidr</b> invalid.) Once you have created a cluster that uses a VPC network, you can add new CIDR blocks, but you are unable to modify the existing ones. If you need to adjust the CIDR blocks, you will have to create a new cluster.</p> <p>If this parameter is left blank, a unique CIDR block is randomly allocated from 172.(16-31).0.0/16 and 10.(0   16   32   48   64   80   96   112).0.0/12.</p>

Parameter	Mandatory	Type	Description
cidrs	No	Array of <b>ContainerCIDR</b> objects	<p>List of container CIDR blocks</p> <p>In clusters of v1.21 or later, the <b>cidrs</b> field is used. When the cluster network type is <b>vpc-router</b>, you can configure a maximum of 20 container CIDR blocks. In cluster versions earlier than v1.21, if the <b>cidrs</b> field is used, the first CIDR element in the array is used as the container CIDR block.</p> <p>Once you have created a cluster that uses a container tunnel network, you cannot modify the CIDR blocks. Once you have created a cluster that uses a VPC network, you can add new CIDR blocks, but you are unable to modify the existing ones. If you need to adjust the CIDR blocks, you will have to create a new cluster.</p> <p>If this parameter is left blank, a unique CIDR block is randomly allocated from 172.(16-31).0.0/16 and 10.(0   16   32   48   64   80   96   112).0.0/12.</p>

**Table 4-9** ContainerCIDR

Parameter	Mandatory	Type	Description
cidr	Yes	String	<p>Container CIDR block.</p> <p>Recommended: 10.0.0.0/12-19, 172.16.0.0/16-19, or 192.168.0.0/16-19. If the selected CIDR block conflicts with existing ones, an error will be reported.</p>

**Table 4-10** EniNetwork

Parameter	Mandatory	Type	Description
eniSubnetId	Yes	String	IPv4 subnet ID of ENI subnet. (IPv6 is not supported and is being discarded.) You can obtain it in either of the following ways: <ul style="list-style-type: none"> <li>Method 1: Log in to the VPC console and click the target subnet on the <b>Subnets</b> page. You can view the IPv4 subnet ID on the displayed page.</li> <li>Method 2: Use the VPC API to obtain subnet IDs.</li> </ul>
eniSubnetCIDR	No	String	ENI subnet CIDR (being discarded)
subnets	Yes	Array of <a href="#">NetworkSubnet</a> objects	List of IPv4 subnet IDs

**Table 4-11** NetworkSubnet

Parameter	Mandatory	Type	Description
subnetID	Yes	String	IPv4 subnet ID of the subnet for creating master nodes. IPv6 is not supported. You can obtain it in either of the following ways: <ul style="list-style-type: none"> <li>Method 1: Log in to the VPC console and click the target subnet on the <b>Subnets</b> page. You can view the IPv4 subnet ID on the displayed page.</li> <li>Method 2: Use the VPC API to obtain subnets.</li> </ul>

**Table 4-12** ServiceNetwork

Parameter	Mandatory	Type	Description
IPv4CIDR	No	String	Value range of the Kubernetes clusterIP IPv4 CIDR blocks. If this parameter is not specified during cluster creation, the default value <b>10.247.0.0/16</b> will be used.

**Table 4-13** PublicAccess

Parameter	Mandatory	Type	Description
cidrs	No	Array of strings	Trustlist of network CIDRs that are allowed to access cluster APIs. You are advised to allow the traffic from VPC and container network CIDRs. By default, no trustlist is configured, and the value is ["0.0.0.0/0"].  <b>NOTE</b> This parameter is valid only when a cluster is created.

**Table 4-14** Authentication

Parameter	Mandatory	Type	Description
mode	No	String	Cluster authentication mode. <ul style="list-style-type: none"> <li>Clusters of Kubernetes v1.11 or earlier support <b>x509</b>, <b>rbac</b>, and <b>authenticating_proxy</b>. Defaults to <b>x509</b>.</li> <li>Clusters of Kubernetes v1.13 or later support <b>rbac</b> and <b>authenticating_proxy</b>. Defaults to <b>rbac</b>.</li> </ul>
authenticatingProxy	No	<b>AuthenticatingProxy</b> object	Configuration related to the <code>authenticating_proxy</code> mode. This field is mandatory when the authentication mode is <b>authenticating_proxy</b> .

**Table 4-15** AuthenticatingProxy

Parameter	Mandatory	Type	Description
ca	No	String	X509 CA certificate (Base64-encoded) configured in authenticating_proxy mode. This field is mandatory when the cluster authentication mode is <b>authenticating_proxy</b> . Maximum size: 1 MB
cert	No	String	Client certificate issued by the X509 CA certificate configured in authenticating_proxy mode, which is used for authentication from kube-apiserver to the extended API server. (The value must be Base64-encoded.) This field is mandatory when the cluster authentication mode is <b>authenticating_proxy</b> . Maximum size: 1 MB
privateKey	No	String	Private key of the client certificate issued by the X509 CA certificate configured in authenticating_proxy mode, which is used for authentication from kube-apiserver to the extended API server. The private key used by the Kubernetes cluster does not support password encryption. Use an unencrypted private key. (The value must be Base64-encoded.) This field is mandatory when the cluster authentication mode is <b>authenticating_proxy</b> . Maximum size: 1 MB

**Table 4-16** MasterSpec

Parameter	Mandatory	Type	Description
availabilityZone	No	String	AZ

**Table 4-17** ResourceTag

Parameter	Mandatory	Type	Description
key	No	String	Key. <ul style="list-style-type: none"> <li>• Cannot be null. Max characters: 128.</li> <li>• Use letters, digits, and spaces in UTF-8 format.</li> <li>• Can contain the following special characters: <code>_.:=+@</code></li> <li>• Cannot start with <code>_sys_</code>.</li> </ul>
value	No	String	Value. <ul style="list-style-type: none"> <li>• Can be null but not the default. Max characters: 255.</li> <li>• Use letters, digits, and spaces in UTF-8 format.</li> <li>• Can contain the following special characters: <code>._:/=+@</code>.</li> </ul>

**Table 4-18** ClusterExtendParam

Parameter	Mandatory	Type	Description
clusterAZ	No	String	AZ of master nodes in a cluster. <ul style="list-style-type: none"> <li>• <b>multi_az:</b> (Optional) The cluster will span across AZs. Multiple AZs can be configured only when a cluster with multiple master nodes is used.</li> <li>• <i>AZ of the dedicated cloud computing pool:</i> The cluster will be deployed in the DeC AZ. It is mandatory for dedicated CCE clusters.</li> </ul>



Parameter	Mandatory	Type	Description
dssMasterVolumes	No	String	<p>Whether the system and data disks of a master node use dedicated distributed storage. If this parameter is omitted or left unspecified, EVS disks are used by default.</p> <p>This parameter is mandatory for dedicated CCE clusters. It is in the following format:  <code>&lt;rootVol.dssPoolID&gt;.&lt;rootVol.volType&gt;;&lt;dataVol.dssPoolID&gt;.&lt;dataVol.volType&gt;</code></p> <p>Field description:</p> <ul style="list-style-type: none"> <li>• <b>rootVol</b> is the system disk. <b>dataVol</b> is the data disk.</li> <li>• <b>dssPoolID</b> indicates the ID of the DSS storage pool.</li> <li>• <b>volType</b> indicates the storage volume type of the DSS storage pool, such as SAS, SSD, SATA, ESSD, GPSSD, ESSD2, and GPSSD2.</li> </ul> <p>Example:  c950ee97-587c-4f24-8a74-3367e3da570f.sas;6edbc2f4-1507-44f8-ac0d-eed1d2608d38.ssd</p> <p><b>NOTE</b>  This field cannot be configured for non-dedicated CCE clusters.</p>
enterpriseProjectId	No	String	<p>ID of the enterprise project that a cluster belongs to</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• An enterprise project can be configured only after the enterprise project function is enabled.</li> </ul>

Parameter	Mandatory	Type	Description
kubeProxyMode	No	String	<p>Service forwarding mode. Two modes are available:</p> <ul style="list-style-type: none"> <li>• <b>iptables</b>: Traditional kube-proxy uses iptables rules to implement Service load balancing. In this mode, too many iptables rules will be generated when many Services are deployed. In addition, non-incremental updates will cause latency and even tangible performance issues in the case of service traffic spikes.</li> <li>• <b>ipvs</b>: Optimized kube-proxy mode with higher throughput and faster speed. This mode supports incremental updates and can keep connections uninterrupted during Service updates. It is suitable for large-sized clusters.</li> </ul> <p><b>NOTE</b> This parameter has been deprecated. If this parameter and <b>kubeProxyMode</b> in <b>ClusterSpec</b> are specified at the same time, the latter is used.</p>
clusterExternalIP	No	String	EIP of the master node

Parameter	Mandatory	Type	Description
alpha.cce/ fixPoolMask	No	String	Number of mask bits of the fixed IP address pool of the container networks. It is supported only by the VPC network model (vpc-router). The number of mask bits and <b>maxPods</b> configured during node creation determine how many IP addresses can be allocated to containers from a node and how many pods can be created in a node. For details, see <a href="#">Maximum Number of Pods That Can Be Created on a Node</a> . For integer characters, the value ranges from 24 to 28.
decMasterFlavor	No	String	Specifications of the master node in the dedicated hybrid cluster.
dockerUmaskMode	No	String	Default UmaskMode configuration of Docker in a cluster. The value can be <b>secure</b> or <b>normal</b> . If this parameter is not specified, <b>normal</b> is used by default.
kubernetes.io/ cpuManagerPolicy	No	String	Cluster CPU management policy. The value can be none (or null) or static. The default value is none (or null). <ul style="list-style-type: none"> <li>• <b>none</b> or <b>null</b>: disables pods from exclusively occupying CPUs. Select this option if you want a large pool of shareable CPU cores.</li> <li>• <b>static</b>: enables pods to exclusively occupy CPUs. Select this option if your workload is sensitive to CPU cache and scheduling latency.</li> </ul>
upgradefrom	No	String	Records of how the cluster is upgraded to the current version

**Table 4-19** PackageConfiguration

Parameter	Mandatory	Type	Description
name	No	String	Component name.
configurations	No	Array of <b>ConfigurationItem</b> objects	Component configuration item.

**Table 4-20** ConfigurationItem

Parameter	Mandatory	Type	Description
name	No	String	Component configuration item name.
value	No	Object	Component configuration item value.

**Table 4-21** ClusterStatus

Parameter	Mandatory	Type	Description
phase	No	String	<p>Cluster status. Options:</p> <ul style="list-style-type: none"> <li>● <b>Available:</b> The cluster is running properly.</li> <li>● <b>Unavailable:</b> The cluster is exhibiting unexpected behavior and needs to be manually deleted.</li> <li>● <b>ScalingUp:</b> Nodes are being added to the cluster.</li> <li>● <b>ScalingDown:</b> The cluster is being downsized to fewer nodes.</li> <li>● <b>Creating:</b> The cluster is being created.</li> <li>● <b>Deleting:</b> The cluster is being deleted.</li> <li>● <b>Upgrading:</b> The cluster is being upgraded.</li> <li>● <b>Resizing:</b> Cluster specifications are being changed.</li> <li>● <b>ResizeFailed:</b> Cluster specification changing is abnormal.</li> <li>● <b>RollingBack:</b> The cluster is being rolled back.</li> <li>● <b>RollbackFailed:</b> The cluster rollback is abnormal.</li> <li>● <b>Hibernating:</b> The cluster is being hibernated.</li> <li>● <b>Hibernation:</b> The cluster is in hibernation.</li> <li>● <b>Freezing:</b> The cluster is being frozen.</li> <li>● <b>Frozen:</b> The cluster has been frozen.</li> <li>● <b>UnFreezing:</b> The cluster is being unfrozen.</li> <li>● <b>Awaking:</b> The cluster is being woken up from hibernation.</li> <li>● <b>Empty:</b> The cluster does not have any resources. This field is discarded.</li> </ul>

Parameter	Mandatory	Type	Description
			<ul style="list-style-type: none"> <li><b>Error:</b> Resources in the cluster are abnormal. You can manually delete the cluster.</li> </ul>
jobID	No	String	<p>ID of the task associated with the cluster in the current state. Options:</p> <ul style="list-style-type: none"> <li>ID of the associated task returned when creating a cluster. You can use it to obtain the auxiliary tasks for creating a cluster.</li> <li>ID of the associated task returned when a cluster fails to be deleted or is deleted. If this parameter is not empty, you can use the task ID to obtain the auxiliary tasks for deleting a cluster.</li> </ul> <p><b>NOTE</b> Tasks are short-lived. Do not use task information in scenarios such as cluster status determination.</p>
reason	No	String	Reason of cluster state change. This parameter is returned if the cluster is not in the Available state.
message	No	String	Detailed information about why the cluster changes to the current state. This parameter is returned if the cluster is not in the Available state.
endpoints	No	Array of <a href="#">ClusterEndpoints</a> objects	Access address of kube-apiserver in the cluster.
deleteOption	No	Object	Whether to delete configurations. This parameter is contained only in the response to the deletion request.
deleteStatus	No	Object	Whether to delete the status information. This parameter is contained only in the response to the deletion request.

**Table 4-22** ClusterEndpoints

Parameter	Mandatory	Type	Description
url	No	String	Access address of kube-apiserver in the cluster.
type	No	String	Type of the cluster access address. <ul style="list-style-type: none"> <li>• <b>Internal</b>: address for internal network access</li> <li>• <b>External</b>: address for external network access</li> </ul>

## Response Parameters

Status code: 201

**Table 4-23** Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at <b>Cluster</b> or <b>cluster</b> and cannot be changed.
apiVersion	String	API version. The value is fixed at <b>v3</b> and cannot be changed.
metadata	<b>ClusterMetadata</b> object	Basic information about a cluster. Metadata is a collection of attributes.
spec	<b>ClusterSpec</b> object	Detailed description of the cluster. CCE creates or updates objects by defining or updating <b>spec</b> .
status	<b>ClusterStatus</b> object	Cluster status and job ID of the cluster creation job.

**Table 4-24** ClusterMetadata

Parameter	Type	Description
name	String	Cluster name. Enter 4 to 128 characters, starting with a lowercase letter and not ending with a hyphen (-). Only lowercase letters, digits, and hyphens (-) are allowed.

Parameter	Type	Description
uid	String	Cluster ID, which uniquely identifies a cluster. This ID is automatically generated after a cluster is created. Only the automatically generated ID will take effect.
alias	String	Alias of a cluster name displayed on the CCE console, and the name can be changed. Enter 4 to 128 characters starting with a letter and not ending with a hyphen (-). Only digits, letters, and hyphens (-) are allowed. A cluster alias must be unique. In the request body for creating or updating a cluster, if the cluster alias is not specified or is left blank, the alias of the cluster is the same as the cluster name. In the response body for obtaining a cluster, if the cluster alias is not configured, the cluster alias will not be returned.
annotations	Map<String,String>	Cluster annotations, in the format of key-value pairs. <pre>"annotations": {   "key1" : "value1",   "key2" : "value2" }</pre> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• <b>annotations:</b> Does not label or select objects. The metadata in <b>annotations</b> may be small or large, structured or unstructured, and may include characters that are not allowed in labels.</li> <li>• This field is not stored in the database and is used only to specify the add-ons to be installed in the cluster.</li> <li>• Install ICAGENT during cluster creation by adding the key-value pair "cluster.install.addons.external/install":{"addonTemplateName":"icagent"}".</li> </ul>
labels	Map<String,String>	Cluster labels, in the format of key-value pairs. <b>NOTE</b> The value of this field is automatically generated by the system and is used by the frontend to identify the features supported by the cluster during the upgrade. Customized values are invalid.
creationTimestamp	String	Time when the cluster was created.
updateTimestamp	String	Time when the cluster was updated.



**Table 4-25** ClusterSpec

Parameter	Type	Description
category	String	Cluster type. Options: <ul style="list-style-type: none"> <li>• <b>CCE</b>: CCE cluster CCE cluster supports hybrid deployment of VMs and BMSs, and heterogeneous nodes such as GPU and NPU nodes, allowing you to run your containers in a secure and stable container runtime environment based on a high-performance network model.</li> </ul>
type	String	Master node architecture: <ul style="list-style-type: none"> <li>• <b>VirtualMachine</b>: x86</li> </ul>

Parameter	Type	Description
flavor	String	<p>Cluster specifications. Specifications of clusters 1.15 or later versions can be changed after they are created. For details, see <a href="#">Modifying Cluster Specifications</a>. Options:</p> <ul style="list-style-type: none"> <li>• <b>cce.s1.small</b>: a small-scale CCE cluster with one master node and a maximum of 50 worker nodes</li> <li>• <b>cce.s1.medium</b>: a medium-scale CCE cluster with one master node and a maximum of 200 worker nodes</li> <li>• <b>cce.s2.small</b>: a small-scale CCE cluster with three master nodes and a maximum of 50 worker nodes</li> <li>• <b>cce.s2.medium</b>: a medium-scale CCE cluster with three master nodes and a maximum of 200 worker nodes</li> <li>• <b>cce.s2.large</b>: a large-scale CCE cluster with three master nodes and a maximum of 1,000 worker nodes</li> <li>• <b>cce.s2.xlarge</b>: an ultra-large-scale CCE cluster with three master nodes and a maximum of 2,000 worker nodes</li> </ul> <p><b>NOTE</b> The fields in the parameters are described as follows:</p> <ul style="list-style-type: none"> <li>• <b>s1</b>: specifies a cluster with one master node. If the master node is faulty, the cluster will become unavailable, but running workloads in the cluster are not affected.</li> <li>• <b>s2</b>: specifies an HA cluster with three master nodes. If one of the master nodes is faulty, the cluster is still available.</li> <li>• <b>small</b>: specifies that a cluster can manage a maximum of 50 worker nodes.</li> <li>• <b>medium</b>: specifies that a cluster can manage a maximum of 200 worker nodes.</li> <li>• <b>large</b>: specifies that a cluster can manage a maximum of 1,000 worker nodes.</li> <li>• <b>xlarge</b>: specifies that a cluster can manage a maximum of 2,000 worker nodes.</li> </ul>

Parameter	Type	Description
version	String	<p>Version of a cluster, which mirrors the baseline version of the Kubernetes community. The latest version is recommended.</p> <p>You can create clusters of the latest three versions on the CCE console. To learn which cluster versions are available, log in to the CCE console, create a cluster, and check supported cluster versions. You can call APIs to create clusters of other versions. However, these cluster versions will be gradually terminated. For details about the support policy, see the CCE announcement.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• If not specified, a cluster of the latest version will be created.</li> <li>• If a baseline cluster version is specified but the R version is not specified, a cluster of the latest R version will be created by default. It is a good practice not to specify the R version.</li> </ul>
platformVersion	String	<p>CCE cluster platform version, indicating the internal version under the cluster version (<b>version</b>). Platform versions are used to trace iterations in a major cluster version. They are unique within a major cluster version and recounted when the major cluster version changes. This parameter cannot be customized. When you create a cluster, the latest corresponding platform version is automatically selected.</p> <p>The format of <b>platformVersion</b> is <b>cce.X.Y</b>.</p> <ul style="list-style-type: none"> <li>• <b>X</b>: internal feature version, indicating changes in features, patches, or OS support in the cluster version. The value starts from <b>1</b> and increases monotonically.</li> <li>• <b>Y</b>: patch version of an internal feature version. It is used only for software package update after the feature version is brought online. No other modification is involved. The value starts from <b>0</b> and increases monotonically.</li> </ul>
legacyVersion	String	<p>(Discarded) Outdated version of CCE clusters, which serves no practical purpose and is only used to showcase the cluster version and platform version combination. The version number is globally unique. For example, if the cluster version is va.b and the platformVersion is cce.X.Y, the value of legacyVersion is va.b.X-rY.</p>

Parameter	Type	Description
description	String	Cluster description, for example, which purpose the cluster is intended to serve. By default, this field is left unspecified. To modify cluster description after the cluster is created, call the <a href="#">API for updating cluster information</a> or go to the cluster details page on the CCE console. Only UTF-8 encoding is supported.
customSan	Array of strings	Custom SAN field in the server certificate of the cluster API server, which must comply with the SSL and X509 format specifications. <ol style="list-style-type: none"> <li>1. Duplicate names are not allowed.</li> <li>2. Must comply with the IP address and domain name formats.</li> </ol> <p>Example:</p> <pre>SAN 1: DNS Name=example.com SAN 2: DNS Name=www.example.com SAN 3: DNS Name=example.net SAN 4: IP Address=93.184.216.34</pre>
ipv6enable	Boolean	Whether a cluster supports IPv6 addresses. It is supported by clusters 1.15 and later. <p><b>NOTE</b></p> <p>After IPv6 is enabled, iptables is not supported. The VPC network model does not support IPv4/IPv6 dual-stack.</p>
hostNetwork	<a href="#">HostNetwork</a> object	Node networking parameters, including VPC and subnet ID. This field is mandatory because nodes in a cluster communicate with each other by using a VPC.
containerNetwork	<a href="#">ContainerNetwork</a> object	Container networking parameters, including the container network model and container CIDR block.
serviceNetwork	<a href="#">ServiceNetwork</a> object	Service CIDR block, including IPv4 CIDR blocks.
publicAccess	<a href="#">PublicAccess</a> object	Cluster API access control.
authentication	<a href="#">Authentication</a> object	Configurations of the cluster authentication mode.
billingMode	Integer	Billing mode of a cluster. <ul style="list-style-type: none"> <li>• <b>0</b>: pay-per-use</li> </ul> <p>Defaults to pay-per-use.</p>
masters	Array of <a href="#">MasterSpec</a> objects	Advanced configurations of the master nodes. If this parameter is not configured, no value is returned.

Parameter	Type	Description
kubernetesSvcIpRange	String	Service CIDR blocks for Kubernetes clusterIPs. This field is available only for clusters of v1.11.7 and later. If this parameter is not specified during cluster creation, the default value <b>10.247.0.0/16</b> will be used. This parameter is deprecated. Use <b>serviceNetwork</b> instead. The new field contains the IPv4 CIDR blocks.
clusterTags	Array of <b>ResourceTag</b> objects	Cluster resource tags.
kubeProxyMode	String	Service forwarding mode. Options: <ul style="list-style-type: none"> <li>• <b>iptables</b>: Traditional kube-proxy uses iptables rules to implement Service load balancing. In this mode, too many iptables rules will be generated when many Services are deployed. In addition, non-incremental updates will cause latency and even tangible performance issues in the case of service traffic spikes.</li> <li>• <b>ipvs</b>: Optimized kube-proxy mode with higher throughput and faster speed. This mode supports incremental updates and can keep connections uninterrupted during Service updates. It is suitable for large-sized clusters.</li> </ul> <p><b>NOTE</b> <b>iptables</b> is used by default.</p>
az	String	AZ. This field is returned only for a query.
extendParam	<b>ClusterExtendedParam</b> object	Extended field to decide whether the cluster will span across AZs or belong to a specified enterprise project, or whether a dedicated CCE cluster is to be created.
supportIstio	Boolean	Whether Istio is supported.
configurationsOverride	Array of <b>PackageConfiguration</b> objects	Cluster default component configuration override. If you specify a component or parameter that is not supported, the configuration item will be ignored.

**Table 4-26** HostNetwork

Parameter	Type	Description
vpc	String	ID of the VPC used to create a master node. You can obtain it in either of the following ways: <ul style="list-style-type: none"> <li>Method 1: Log in to the VPC console and view the VPC ID in the VPC details page.</li> <li>Method 2: Use the VPC API to obtain VPC IDs.</li> </ul>
subnet	String	Network ID of the subnet used to create a master node. You can obtain it in either of the following ways: <ul style="list-style-type: none"> <li>Method 1: Log in to the VPC console and click the target subnet on the <b>Subnets</b> page. You can view the network ID on the displayed page.</li> <li>Method 2: Use the VPC API to obtain subnets.</li> </ul>
SecurityGroup	String	Default worker node security group ID of the cluster. If specified, the cluster will be bound to the target security group. Otherwise, the system will automatically create a default worker node security group for you. The default worker node security group needs to allow access from certain ports to ensure normal communications.
controlPlaneSecurityGroup	String	Security group ID of the cluster control plane node, which is automatically generated after the cluster is created. A user-defined value will not take effect.

**Table 4-27** ContainerNetwork

Parameter	Type	Description
mode	String	Container network model. Select one of the following possible values: <ul style="list-style-type: none"> <li><b>overlay_l2</b>: an overlay_l2 network (container tunnel network) built for containers by using OpenVSwitch (OVS).</li> <li><b>vpc-router</b>: an underlay_l2 network built for containers by using IPvlan and custom VPC routes.</li> </ul>

Parameter	Type	Description
cidr	String	<p>Container CIDR block. Recommended: 10.0.0.0/12-19, 172.16.0.0/16-19, or 192.168.0.0/16-19. If the selected CIDR block conflicts with existing ones, an error will be reported.</p> <p>Not editable after the cluster is created. (This parameter has been discarded. A specified <b>cidrs</b> will make <b>cidr</b> invalid.) Once you have created a cluster that uses a VPC network, you can add new CIDR blocks, but you are unable to modify the existing ones. If you need to adjust the CIDR blocks, you will have to create a new cluster.</p> <p>If this parameter is left blank, a unique CIDR block is randomly allocated from 172.(16-31).0.0/16 and 10.(0   16   32   48   64   80   96   112).0.0/12.</p>
cidrs	Array of <a href="#">ContainerCIDR</a> objects	<p>List of container CIDR blocks In clusters of v1.21 or later, the <b>cidrs</b> field is used. When the cluster network type is <b>vpc-router</b>, you can configure a maximum of 20 container CIDR blocks. In cluster versions earlier than v1.21, if the <b>cidrs</b> field is used, the first CIDR element in the array is used as the container CIDR block.</p> <p>Once you have created a cluster that uses a container tunnel network, you cannot modify the CIDR blocks. Once you have created a cluster that uses a VPC network, you can add new CIDR blocks, but you are unable to modify the existing ones. If you need to adjust the CIDR blocks, you will have to create a new cluster.</p> <p>If this parameter is left blank, a unique CIDR block is randomly allocated from 172.(16-31).0.0/16 and 10.(0   16   32   48   64   80   96   112).0.0/12.</p>

**Table 4-28** ContainerCIDR

Parameter	Type	Description
cidr	String	<p>Container CIDR block. Recommended: 10.0.0.0/12-19, 172.16.0.0/16-19, or 192.168.0.0/16-19. If the selected CIDR block conflicts with existing ones, an error will be reported.</p>

**Table 4-29** EniNetwork

Parameter	Type	Description
eniSubnetId	String	IPv4 subnet ID of ENI subnet. (IPv6 is not supported and is being discarded.) You can obtain it in either of the following ways: <ul style="list-style-type: none"> <li>Method 1: Log in to the VPC console and click the target subnet on the <b>Subnets</b> page. You can view the IPv4 subnet ID on the displayed page.</li> <li>Method 2: Use the VPC API to obtain subnet IDs.</li> </ul>
eniSubnetCIDR	String	ENI subnet CIDR (being discarded)
subnets	Array of <a href="#">NetworkSubnet</a> objects	List of IPv4 subnet IDs

**Table 4-30** NetworkSubnet

Parameter	Type	Description
subnetID	String	IPv4 subnet ID of the subnet for creating master nodes. IPv6 is not supported. You can obtain it in either of the following ways: <ul style="list-style-type: none"> <li>Method 1: Log in to the VPC console and click the target subnet on the <b>Subnets</b> page. You can view the IPv4 subnet ID on the displayed page.</li> <li>Method 2: Use the VPC API to obtain subnets.</li> </ul>

**Table 4-31** ServiceNetwork

Parameter	Type	Description
IPv4CIDR	String	Value range of the Kubernetes clusterIP IPv4 CIDR blocks. If this parameter is not specified during cluster creation, the default value <b>10.247.0.0/16</b> will be used.



**Table 4-32 PublicAccess**

Parameter	Type	Description
cidrs	Array of strings	Trustlist of network CIDRs that are allowed to access cluster APIs. You are advised to allow the traffic from VPC and container network CIDRs. By default, no trustlist is configured, and the value is ["0.0.0.0/0"].  <b>NOTE</b> This parameter is valid only when a cluster is created.

**Table 4-33 Authentication**

Parameter	Type	Description
mode	String	Cluster authentication mode. <ul style="list-style-type: none"> <li>Clusters of Kubernetes v1.11 or earlier support <b>x509</b>, <b>rbac</b>, and <b>authenticating_proxy</b>. Defaults to <b>x509</b>.</li> <li>Clusters of Kubernetes v1.13 or later support <b>rbac</b> and <b>authenticating_proxy</b>. Defaults to <b>rbac</b>.</li> </ul>
authenticatingProxy	<b>AuthenticatingProxy</b> object	Configuration related to the authenticating_proxy mode. This field is mandatory when the authentication mode is <b>authenticating_proxy</b> .

**Table 4-34 AuthenticatingProxy**

Parameter	Type	Description
ca	String	X509 CA certificate (Base64-encoded) configured in authenticating_proxy mode. This field is mandatory when the cluster authentication mode is <b>authenticating_proxy</b> . Maximum size: 1 MB
cert	String	Client certificate issued by the X509 CA certificate configured in authenticating_proxy mode, which is used for authentication from kube-apiserver to the extended API server. (The value must be Base64-encoded.) This field is mandatory when the cluster authentication mode is <b>authenticating_proxy</b> . Maximum size: 1 MB

Parameter	Type	Description
privateKey	String	Private key of the client certificate issued by the X509 CA certificate configured in <code>authenticating_proxy</code> mode, which is used for authentication from kube-apiserver to the extended API server. The private key used by the Kubernetes cluster does not support password encryption. Use an unencrypted private key. (The value must be Base64-encoded.) This field is mandatory when the cluster authentication mode is <b>authenticating_proxy</b> . Maximum size: 1 MB

**Table 4-35** MasterSpec

Parameter	Type	Description
availabilityZone	String	AZ

**Table 4-36** ResourceTag

Parameter	Type	Description
key	String	Key. <ul style="list-style-type: none"> <li>• Cannot be null. Max characters: 128.</li> <li>• Use letters, digits, and spaces in UTF-8 format.</li> <li>• Can contain the following special characters: <code>._:=-@</code></li> <li>• Cannot start with <code>_sys_</code>.</li> </ul>
value	String	Value. <ul style="list-style-type: none"> <li>• Can be null but not the default. Max characters: 255.</li> <li>• Use letters, digits, and spaces in UTF-8 format.</li> <li>• Can contain the following special characters: <code>._:=-@</code>.</li> </ul>

**Table 4-37** ClusterExtendParam

Parameter	Type	Description
clusterAZ	String	<p>AZ of master nodes in a cluster.</p> <ul style="list-style-type: none"> <li>• <b>multi_az</b>: (Optional) The cluster will span across AZs. Multiple AZs can be configured only when a cluster with multiple master nodes is used.</li> <li>• <i>AZ of the dedicated cloud computing pool</i>: The cluster will be deployed in the DeC AZ. It is mandatory for dedicated CCE clusters.</li> </ul>
dssMasterVolumes	String	<p>Whether the system and data disks of a master node use dedicated distributed storage. If this parameter is omitted or left unspecified, EVS disks are used by default.</p> <p>This parameter is mandatory for dedicated CCE clusters. It is in the following format:  <code>&lt;rootVol.dssPoolID&gt;.&lt;rootVol.volType&gt;;&lt;dataVol.dssPoolID&gt;.&lt;dataVol.volType&gt;</code></p> <p>Field description:</p> <ul style="list-style-type: none"> <li>• <b>rootVol</b> is the system disk. <b>dataVol</b> is the data disk.</li> <li>• <b>dssPoolID</b> indicates the ID of the DSS storage pool.</li> <li>• <b>volType</b> indicates the storage volume type of the DSS storage pool, such as SAS, SSD, SATA, ESSD, GPSSD, ESSD2, and GPSSD2.</li> </ul> <p>Example:  c950ee97-587c-4f24-8a74-3367e3da570f.sas;6dbc2f4-1507-44f8-ac0d-eed1d2608d38.ssd</p> <p><b>NOTE</b>  This field cannot be configured for non-dedicated CCE clusters.</p>
enterpriseProjectId	String	<p>ID of the enterprise project that a cluster belongs to</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• An enterprise project can be configured only after the enterprise project function is enabled.</li> </ul>

Parameter	Type	Description
kubeProxyMode	String	<p>Service forwarding mode. Two modes are available:</p> <ul style="list-style-type: none"> <li>• <b>iptables</b>: Traditional kube-proxy uses iptables rules to implement Service load balancing. In this mode, too many iptables rules will be generated when many Services are deployed. In addition, non-incremental updates will cause latency and even tangible performance issues in the case of service traffic spikes.</li> <li>• <b>ipvs</b>: Optimized kube-proxy mode with higher throughput and faster speed. This mode supports incremental updates and can keep connections uninterrupted during Service updates. It is suitable for large-sized clusters.</li> </ul> <p><b>NOTE</b> This parameter has been deprecated. If this parameter and <b>kubeProxyMode</b> in <b>ClusterSpec</b> are specified at the same time, the latter is used.</p>
clusterExternalIP	String	EIP of the master node
alpha.cce/fixPoolMask	String	<p>Number of mask bits of the fixed IP address pool of the container networks. It is supported only by the VPC network model (vpc-router).</p> <p>The number of mask bits and <b>maxPods</b> configured during node creation determine how many IP addresses can be allocated to containers from a node and how many pods can be created in a node. For details, see <a href="#">Maximum Number of Pods That Can Be Created on a Node</a>.</p> <p>For integer characters, the value ranges from 24 to 28.</p>
decMasterFlavor	String	Specifications of the master node in the dedicated hybrid cluster.
dockerUmaskMode	String	Default UmaskMode configuration of Docker in a cluster. The value can be <b>secure</b> or <b>normal</b> . If this parameter is not specified, <b>normal</b> is used by default.

Parameter	Type	Description
kubernetes.io/cpuManagerPolicy	String	Cluster CPU management policy. The value can be none (or null) or static. The default value is none (or null). <ul style="list-style-type: none"> <li>• <b>none</b> or <b>null</b>: disables pods from exclusively occupying CPUs. Select this option if you want a large pool of shareable CPU cores.</li> <li>• <b>static</b>: enables pods to exclusively occupy CPUs. Select this option if your workload is sensitive to CPU cache and scheduling latency.</li> </ul>
upgradefrom	String	Records of how the cluster is upgraded to the current version

**Table 4-38** PackageConfiguration

Parameter	Type	Description
name	String	Component name.
configurations	Array of <a href="#">Configuratio nItem</a> objects	Component configuration item.

**Table 4-39** ConfigurationItem

Parameter	Type	Description
name	String	Component configuration item name.
value	Object	Component configuration item value.

**Table 4-40** ClusterStatus

Parameter	Type	Description
phase	String	<p>Cluster status. Options:</p> <ul style="list-style-type: none"> <li>● <b>Available:</b> The cluster is running properly.</li> <li>● <b>Unavailable:</b> The cluster is exhibiting unexpected behavior and needs to be manually deleted.</li> <li>● <b>ScalingUp:</b> Nodes are being added to the cluster.</li> <li>● <b>ScalingDown:</b> The cluster is being downsized to fewer nodes.</li> <li>● <b>Creating:</b> The cluster is being created.</li> <li>● <b>Deleting:</b> The cluster is being deleted.</li> <li>● <b>Upgrading:</b> The cluster is being upgraded.</li> <li>● <b>Resizing:</b> Cluster specifications are being changed.</li> <li>● <b>ResizeFailed:</b> Cluster specification changing is abnormal.</li> <li>● <b>RollingBack:</b> The cluster is being rolled back.</li> <li>● <b>RollbackFailed:</b> The cluster rollback is abnormal.</li> <li>● <b>Hibernating:</b> The cluster is being hibernated.</li> <li>● <b>Hibernation:</b> The cluster is in hibernation.</li> <li>● <b>Freezing:</b> The cluster is being frozen.</li> <li>● <b>Frozen:</b> The cluster has been frozen.</li> <li>● <b>UnFreezing:</b> The cluster is being unfrozen.</li> <li>● <b>Awaking:</b> The cluster is being woken up from hibernation.</li> <li>● <b>Empty:</b> The cluster does not have any resources. This field is discarded.</li> <li>● <b>Error:</b> Resources in the cluster are abnormal. You can manually delete the cluster.</li> </ul>

Parameter	Type	Description
jobID	String	<p>ID of the task associated with the cluster in the current state. Options:</p> <ul style="list-style-type: none"> <li>• ID of the associated task returned when creating a cluster. You can use it to obtain the auxiliary tasks for creating a cluster.</li> <li>• ID of the associated task returned when a cluster fails to be deleted or is deleted. If this parameter is not empty, you can use the task ID to obtain the auxiliary tasks for deleting a cluster.</li> </ul> <p><b>NOTE</b> Tasks are short-lived. Do not use task information in scenarios such as cluster status determination.</p>
reason	String	Reason of cluster state change. This parameter is returned if the cluster is not in the Available state.
message	String	Detailed information about why the cluster changes to the current state. This parameter is returned if the cluster is not in the Available state.
endpoints	Array of <a href="#">ClusterEndpoints</a> objects	Access address of kube-apiserver in the cluster.
deleteOption	Object	Whether to delete configurations. This parameter is contained only in the response to the deletion request.
deleteStatus	Object	Whether to delete the status information. This parameter is contained only in the response to the deletion request.

**Table 4-41** ClusterEndpoints

Parameter	Type	Description
url	String	Access address of kube-apiserver in the cluster.
type	String	<p>Type of the cluster access address.</p> <ul style="list-style-type: none"> <li>• <b>Internal:</b> address for internal network access</li> <li>• <b>External:</b> address for external network access</li> </ul>

## Example Requests

- Create an HA CCE cluster of v1.29 that has three master nodes and a maximum of 50 worker nodes.

```
/api/v3/projects/{project_id}/clusters

{
  "kind": "Cluster",
  "apiVersion": "v3",
  "metadata": {
    "name": "cluster"
  },
  "spec": {
    "category": "CCE",
    "flavor": "cce.s2.small",
    "version": "v1.29",
    "hostNetwork": {
      "vpc": "030bfb19-5fa7-42ad-8a0d-c0721d268867",
      "subnet": "ca964acf-8468-4735-8229-97940ef6c881"
    },
    "containerNetwork": {
      "mode": "vpc-router",
      "cidr": "10.0.0.0/16"
    },
    "kubernetesSvclpRange": "10.247.0.0/16",
    "description": "",
    "billingMode": 0,
    "extendParam": {
      "kubeProxyMode": "iptables",
      "alpha.cce/fixPoolMask": "25",
      "enterpriseProjectId": "0"
    },
    "authentication": {
      "mode": "rbac"
    },
    "ipv6enable": false
  }
}
```

- Create an HA CCE cluster of v1.29 that has three master nodes and a maximum of 50 worker nodes and install ICAgent in the cluster.

```
/api/v3/projects/{project_id}/clusters

{
  "kind": "Cluster",
  "apiVersion": "v3",
  "metadata": {
    "name": "cluster",
    "annotations": {
      "cluster.install.addons.external/install": "[{\"addonTemplateName\":\"icagent\"}]"
    }
  },
  "spec": {
    "category": "CCE",
    "flavor": "cce.s2.small",
    "version": "v1.29",
    "hostNetwork": {
      "vpc": "030bfb19-5fa7-42ad-8a0d-c0721d268867",
      "subnet": "ca964acf-8468-4735-8229-97940ef6c881"
    },
    "containerNetwork": {
      "mode": "vpc-router",
      "cidr": "10.0.0.0/16"
    },
    "kubernetesSvclpRange": "10.247.0.0/16",
    "description": "",
    "billingMode": 0,
    "extendParam": {
      "kubeProxyMode": "iptables",

```



```

    "alpha.cce/fixPoolMask" : "25",
    "enterpriseProjectId" : "0"
  },
  "authentication" : {
    "mode" : "rbac"
  },
  "ipv6enable" : false
}
}

```

- Create an HA CCE cluster of v1.29 that has three master nodes and a maximum of 50 worker nodes and configure the custom default security group for nodes in the cluster.

```

/api/v3/projects/{project_id}/clusters

{
  "kind" : "Cluster",
  "apiVersion" : "v3",
  "metadata" : {
    "name" : "cluster"
  },
  "spec" : {
    "category" : "CCE",
    "flavor" : "cce.s2.small",
    "version" : "v1.29",
    "hostNetwork" : {
      "vpc" : "030bfb19-5fa7-42ad-8a0d-c0721d268867",
      "subnet" : "ca964acf-8468-4735-8229-97940ef6c881",
      "SecurityGroup" : "a4ef108c-2ec6-492f-a6c4-7b64e25ae490"
    },
    "containerNetwork" : {
      "mode" : "vpc-router",
      "cidr" : "10.0.0.0/16"
    },
    "kubernetesSvcIpRange" : "10.247.0.0/16",
    "description" : "",
    "billingMode" : 0,
    "extendParam" : {
      "kubeProxyMode" : "iptables",
      "alpha.cce/fixPoolMask" : "25",
      "enterpriseProjectId" : "0"
    },
    "authentication" : {
      "mode" : "rbac"
    },
    "ipv6enable" : false
  }
}

```

- Create an HA CCE cluster of v1.25 with three master nodes and a maximum of 50 worker nodes.

```

/api/v3/projects/{project_id}/clusters

{
  "kind" : "Cluster",
  "apiVersion" : "v3",
  "metadata" : {
    "name" : "cluster"
  },
  "spec" : {
    "category" : "Turbo",
    "flavor" : "cce.s2.small",
    "version" : "v1.25",
    "type" : "VirtualMachine",
    "hostNetwork" : {
      "vpc" : "030bfb19-5fa7-42ad-8a0d-c0721d268867",
      "subnet" : "ca964acf-8468-4735-8229-97940ef6c881"
    },
    "containerNetwork" : {

```

```

    "mode" : "eni"
  },
  "eniNetwork" : {
    "eniSubnetId" : "861fb11d-2f0e-4c10-a98a-166dc26e4ff7",
    "eniSubnetCIDR" : "192.168.0.0/24",
    "subnets" : [ {
      "subnetID" : "861fb11d-2f0e-4c10-a98a-166dc26e4ff7"
    } ]
  },
  "serviceNetwork" : {
    "IPv4CIDR" : "10.247.0.0/16"
  },
  "description" : "",
  "billingMode" : 0,
  "extendParam" : {
    "kubeProxyMode" : "iptables",
    "enterpriseProjectId" : "0"
  },
  "authentication" : {
    "mode" : "rbac"
  },
  "ipv6enable" : false
}

```

- Configure the trustlist network CIDR for accessing the cluster API when creating a cluster.

```

/api/v3/projects/{project_id}/clusters

{
  "kind" : "Cluster",
  "apiVersion" : "v3",
  "metadata" : {
    "name" : "cluster"
  },
  "spec" : {
    "category" : "CCE",
    "flavor" : "cce.s2.small",
    "version" : "v1.19",
    "hostNetWork" : {
      "vpc" : "030bfb19-5fa7-42ad-8a0d-c0721d268867",
      "subnet" : "ca964acf-8468-4735-8229-97940ef6c881"
    },
    "publicAccess" : {
      "cidrs" : [ "192.168.0.0/16" ]
    },
    "containerNetWork" : {
      "mode" : "vpc-router",
      "cidr" : "10.0.0.0/16"
    },
    "serviceNetwork" : {
      "IPv4CIDR" : "10.247.0.0/16",
      "IPv6CIDR" : "fc00::/112"
    },
    "kubernetesSvcIcpRange" : "10.247.0.0/16",
    "description" : "",
    "billingMode" : 0,
    "extendParam" : {
      "kubeProxyMode" : "iptables",
      "alpha.cce/fixPoolMask" : "25",
      "enterpriseProjectId" : "0"
    },
    "authentication" : {
      "mode" : "rbac"
    },
    "ipv6enable" : false
  }
}

```

## Example Responses

### Status code: 201

The cluster creation job is successfully delivered.

```
{
  "kind": "Cluster",
  "apiVersion": "v3",
  "metadata": {
    "name": "cluster",
    "uid": "1df09f9a-5b9e-11ef-8f52-0255ac10003e",
    "creationTimestamp": "2024-08-16 07:06:53.704389459 +0000 UTC",
    "updateTimestamp": "2024-08-16 07:06:53.704389529 +0000 UTC",
    "annotations": {
      "jobid": "1e50bfbe-5b9e-11ef-8f52-0255ac10003e",
      "resourceJobId": "1df0ec6b-5b9e-11ef-8f52-0255ac10003e"
    },
    "timezone": "*****",
    "ownerReferences": { }
  },
  "spec": {
    "publicAccess": { },
    "category": "CCE",
    "type": "VirtualMachine",
    "enableAutopilot": false,
    "maintenanceWindow": null,
    "flavor": "cce.s1.small",
    "version": "v1.29",
    "platformVersion": "cce.4.0",
    "configurationsOverride": [ {
      "name": "kube-apiserver",
      "configurations": [ {
        "name": "support-overload",
        "value": true
      } ]
    } ],
    "hostNetwork": {
      "vpc": "0538a5d0-9a65-4c1d-a8bf-e9acee237980",
      "subnet": "bc81be88-6e34-4b02-83bd-df0a1f7672c5"
    },
    "containerNetwork": {
      "mode": "vpc-router",
      "cidr": "172.17.0.0/16",
      "cidrs": [ {
        "cidr": "172.17.0.0/16"
      } ]
    },
    "eniNetwork": { },
    "serviceNetwork": {
      "IPv4CIDR": "10.247.0.0/16"
    },
    "clusterOps": { },
    "authentication": {
      "mode": "rbac",
      "authenticatingProxy": { }
    },
    "billingMode": 0,
    "kubernetesSvclpRange": "10.247.0.0/16",
    "kubeProxyMode": "iptables",
    "extendParam": {
      "alpha.cce/fixPoolMask": "25",
      "enterpriseProjectId": "0",
      "orderId": ""
    },
    "encryptionConfig": {
      "mode": "Default"
    },
    "deletionProtection": false
  }
}
```

```

},
"status" : {
  "phase" : "Creating",
  "jobID" : "1e50bfbe-5b9e-11ef-8f52-0255ac10003e"
}
}

```

## Status Codes

Status Code	Description
201	The cluster creation job is successfully delivered.

## Error Codes

See [Error Codes](#).

## 4.2.2 Reading a Specified Cluster

### Function

This API is used to obtain details about a specified cluster.

#### NOTE

The URL for cluster management is in the format of `https://Endpoint/uri`. In the URL, **uri** indicates the resource path, that is, the path for API access.

### URI

GET `/api/v3/projects/{project_id}/clusters/{cluster_id}`

**Table 4-42** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

**Table 4-43** Query Parameters

Parameter	Mandatory	Type	Description
detail	No	String	Whether the details about a cluster are queried. If this parameter is set to <b>true</b> , the total number of nodes (totalNodesNumber), number of normal nodes (activeNodesNumber), total CPUs (totalNodesCPU), total memory size (totalNodesMemory), and installed add-ons (installedAddonInstances) will be added to the annotation. The details of installed add-ons include the add-on name (addonTemplateName), version (version), and status (status).

## Request Parameters

**Table 4-44** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

## Response Parameters

**Status code: 200**

**Table 4-45** Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at <b>Cluster</b> or <b>cluster</b> and cannot be changed.
apiVersion	String	API version. The value is fixed at <b>v3</b> and cannot be changed.
metadata	<b>ClusterMetadata</b> object	Basic information about a cluster. Metadata is a collection of attributes.
spec	<b>ClusterSpec</b> object	Detailed description of the cluster. CCE creates or updates objects by defining or updating <b>spec</b> .
status	<b>ClusterStatus</b> object	Cluster status.

**Table 4-46** ClusterMetadata

Parameter	Type	Description
name	String	Cluster name. Enter 4 to 128 characters, starting with a lowercase letter and not ending with a hyphen (-). Only lowercase letters, digits, and hyphens (-) are allowed.
uid	String	Cluster ID, which uniquely identifies a cluster. This ID is automatically generated after a cluster is created. Only the automatically generated ID will take effect.
alias	String	Alias of a cluster name displayed on the CCE console, and the name can be changed. Enter 4 to 128 characters starting with a letter and not ending with a hyphen (-). Only digits, letters, and hyphens (-) are allowed. A cluster alias must be unique. In the request body for creating or updating a cluster, if the cluster alias is not specified or is left blank, the alias of the cluster is the same as the cluster name. In the response body for obtaining a cluster, if the cluster alias is not configured, the cluster alias will not be returned.

Parameter	Type	Description
annotations	Map<String,String>	<p>Cluster annotations, in the format of key-value pairs.</p> <pre>"annotations": {   "key1" : "value1",   "key2" : "value2" }</pre> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• <b>annotations:</b> Does not label or select objects. The metadata in <b>annotations</b> may be small or large, structured or unstructured, and may include characters that are not allowed in labels.</li> <li>• This field is not stored in the database and is used only to specify the add-ons to be installed in the cluster.</li> <li>• Install ICAgent during cluster creation by adding the key-value pair "cluster.install.addons.external/install":{"addonTemplateName":"icagent"}".</li> </ul>
labels	Map<String,String>	<p>Cluster labels, in the format of key-value pairs.</p> <p><b>NOTE</b></p> <p>The value of this field is automatically generated by the system and is used by the frontend to identify the features supported by the cluster during the upgrade. Customized values are invalid.</p>
creationTimestamp	String	Time when the cluster was created.
updateTimestamp	String	Time when the cluster was updated.

**Table 4-47** ClusterSpec

Parameter	Type	Description
category	String	<p>Cluster type. Options:</p> <ul style="list-style-type: none"> <li>• <b>CCE:</b> CCE cluster CCE cluster supports hybrid deployment of VMs and BMSs, and heterogeneous nodes such as GPU and NPU nodes, allowing you to run your containers in a secure and stable container runtime environment based on a high-performance network model.</li> </ul>
type	String	<p>Master node architecture:</p> <ul style="list-style-type: none"> <li>• <b>VirtualMachine:</b> x86</li> </ul>

Parameter	Type	Description
flavor	String	<p>Cluster specifications. Specifications of clusters 1.15 or later versions can be changed after they are created. For details, see <a href="#">Modifying Cluster Specifications</a>. Options:</p> <ul style="list-style-type: none"> <li>● <b>cce.s1.small</b>: a small-scale CCE cluster with one master node and a maximum of 50 worker nodes</li> <li>● <b>cce.s1.medium</b>: a medium-scale CCE cluster with one master node and a maximum of 200 worker nodes</li> <li>● <b>cce.s2.small</b>: a small-scale CCE cluster with three master nodes and a maximum of 50 worker nodes</li> <li>● <b>cce.s2.medium</b>: a medium-scale CCE cluster with three master nodes and a maximum of 200 worker nodes</li> <li>● <b>cce.s2.large</b>: a large-scale CCE cluster with three master nodes and a maximum of 1,000 worker nodes</li> <li>● <b>cce.s2.xlarge</b>: an ultra-large-scale CCE cluster with three master nodes and a maximum of 2,000 worker nodes</li> </ul> <p><b>NOTE</b> The fields in the parameters are described as follows:</p> <ul style="list-style-type: none"> <li>● <b>s1</b>: specifies a cluster with one master node. If the master node is faulty, the cluster will become unavailable, but running workloads in the cluster are not affected.</li> <li>● <b>s2</b>: specifies an HA cluster with three master nodes. If one of the master nodes is faulty, the cluster is still available.</li> <li>● <b>small</b>: specifies that a cluster can manage a maximum of 50 worker nodes.</li> <li>● <b>medium</b>: specifies that a cluster can manage a maximum of 200 worker nodes.</li> <li>● <b>large</b>: specifies that a cluster can manage a maximum of 1,000 worker nodes.</li> <li>● <b>xlarge</b>: specifies that a cluster can manage a maximum of 2,000 worker nodes.</li> </ul>



Parameter	Type	Description
version	String	<p>Version of a cluster, which mirrors the baseline version of the Kubernetes community. The latest version is recommended.</p> <p>You can create clusters of the latest three versions on the CCE console. To learn which cluster versions are available, log in to the CCE console, create a cluster, and check supported cluster versions. You can call APIs to create clusters of other versions. However, these cluster versions will be gradually terminated. For details about the support policy, see the CCE announcement.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• If not specified, a cluster of the latest version will be created.</li> <li>• If a baseline cluster version is specified but the R version is not specified, a cluster of the latest R version will be created by default. It is a good practice not to specify the R version.</li> </ul>
platformVersion	String	<p>CCE cluster platform version, indicating the internal version under the cluster version (<b>version</b>). Platform versions are used to trace iterations in a major cluster version. They are unique within a major cluster version and recounted when the major cluster version changes. This parameter cannot be customized. When you create a cluster, the latest corresponding platform version is automatically selected.</p> <p>The format of <b>platformVersion</b> is <b>cce.X.Y</b>.</p> <ul style="list-style-type: none"> <li>• <b>X</b>: internal feature version, indicating changes in features, patches, or OS support in the cluster version. The value starts from <b>1</b> and increases monotonically.</li> <li>• <b>Y</b>: patch version of an internal feature version. It is used only for software package update after the feature version is brought online. No other modification is involved. The value starts from <b>0</b> and increases monotonically.</li> </ul>
legacyVersion	String	<p>(Discarded) Outdated version of CCE clusters, which serves no practical purpose and is only used to showcase the cluster version and platform version combination. The version number is globally unique. For example, if the cluster version is va.b and the platformVersion is cce.X.Y, the value of legacyVersion is va.b.X-rY.</p>

Parameter	Type	Description
description	String	Cluster description, for example, which purpose the cluster is intended to serve. By default, this field is left unspecified. To modify cluster description after the cluster is created, call the <a href="#">API for updating cluster information</a> or go to the cluster details page on the CCE console. Only UTF-8 encoding is supported.
customSan	Array of strings	Custom SAN field in the server certificate of the cluster API server, which must comply with the SSL and X509 format specifications. <ol style="list-style-type: none"> <li>1. Duplicate names are not allowed.</li> <li>2. Must comply with the IP address and domain name formats.</li> </ol> <p>Example:</p> <pre>SAN 1: DNS Name=example.com SAN 2: DNS Name=www.example.com SAN 3: DNS Name=example.net SAN 4: IP Address=93.184.216.34</pre>
ipv6enable	Boolean	Whether a cluster supports IPv6 addresses. It is supported by clusters 1.15 and later. <p><b>NOTE</b></p> <p>After IPv6 is enabled, iptables is not supported. The VPC network model does not support IPv4/IPv6 dual-stack.</p>
hostNetwork	<a href="#">HostNetwork</a> object	Node networking parameters, including VPC and subnet ID. This field is mandatory because nodes in a cluster communicate with each other by using a VPC.
containerNetwork	<a href="#">ContainerNetwork</a> object	Container networking parameters, including the container network model and container CIDR block.
serviceNetwork	<a href="#">ServiceNetwork</a> object	Service CIDR block, including IPv4 CIDR blocks.
publicAccess	<a href="#">PublicAccess</a> object	Cluster API access control.
authentication	<a href="#">Authentication</a> object	Configurations of the cluster authentication mode.
billingMode	Integer	Billing mode of a cluster. <ul style="list-style-type: none"> <li>• <b>0</b>: pay-per-use</li> </ul> <p>Defaults to pay-per-use.</p>
masters	Array of <a href="#">MasterSpec</a> objects	Advanced configurations of the master nodes. If this parameter is not configured, no value is returned.

Parameter	Type	Description
kubernetesSvcIpRange	String	Service CIDR blocks for Kubernetes clusterIPs. This field is available only for clusters of v1.11.7 and later. If this parameter is not specified during cluster creation, the default value <b>10.247.0.0/16</b> will be used. This parameter is deprecated. Use <b>serviceNetwork</b> instead. The new field contains the IPv4 CIDR blocks.
clusterTags	Array of <b>ResourceTag</b> objects	Cluster resource tags.
kubeProxyMode	String	Service forwarding mode. Options: <ul style="list-style-type: none"> <li>• <b>iptables</b>: Traditional kube-proxy uses iptables rules to implement Service load balancing. In this mode, too many iptables rules will be generated when many Services are deployed. In addition, non-incremental updates will cause latency and even tangible performance issues in the case of service traffic spikes.</li> <li>• <b>ipvs</b>: Optimized kube-proxy mode with higher throughput and faster speed. This mode supports incremental updates and can keep connections uninterrupted during Service updates. It is suitable for large-sized clusters.</li> </ul> <p><b>NOTE</b> <b>iptables</b> is used by default.</p>
az	String	AZ. This field is returned only for a query.
extendParam	<b>ClusterExtendedParam</b> object	Extended field to decide whether the cluster will span across AZs or belong to a specified enterprise project, or whether a dedicated CCE cluster is to be created.
supportIstio	Boolean	Whether Istio is supported.
configurationsOverride	Array of <b>PackageConfiguration</b> objects	Cluster default component configuration override. If you specify a component or parameter that is not supported, the configuration item will be ignored.

**Table 4-48** HostNetwork

Parameter	Type	Description
vpc	String	ID of the VPC used to create a master node. You can obtain it in either of the following ways: <ul style="list-style-type: none"> <li>Method 1: Log in to the VPC console and view the VPC ID in the VPC details page.</li> <li>Method 2: Use the VPC API to obtain VPC IDs.</li> </ul>
subnet	String	Network ID of the subnet used to create a master node. You can obtain it in either of the following ways: <ul style="list-style-type: none"> <li>Method 1: Log in to the VPC console and click the target subnet on the <b>Subnets</b> page. You can view the network ID on the displayed page.</li> <li>Method 2: Use the VPC API to obtain subnets.</li> </ul>
SecurityGroup	String	Default worker node security group ID of the cluster. If specified, the cluster will be bound to the target security group. Otherwise, the system will automatically create a default worker node security group for you. The default worker node security group needs to allow access from certain ports to ensure normal communications.
controlPlaneSecurityGroup	String	Security group ID of the cluster control plane node, which is automatically generated after the cluster is created. A user-defined value will not take effect.

**Table 4-49** ContainerNetwork

Parameter	Type	Description
mode	String	Container network model. Select one of the following possible values: <ul style="list-style-type: none"> <li><b>overlay_l2</b>: an overlay_l2 network (container tunnel network) built for containers by using OpenVSwitch (OVS).</li> <li><b>vpc-router</b>: an underlay_l2 network built for containers by using IPvlan and custom VPC routes.</li> </ul>

Parameter	Type	Description
cidr	String	<p>Container CIDR block. Recommended: 10.0.0.0/12-19, 172.16.0.0/16-19, or 192.168.0.0/16-19. If the selected CIDR block conflicts with existing ones, an error will be reported.</p> <p>Not editable after the cluster is created. (This parameter has been discarded. A specified <b>cidrs</b> will make <b>cidr</b> invalid.) Once you have created a cluster that uses a VPC network, you can add new CIDR blocks, but you are unable to modify the existing ones. If you need to adjust the CIDR blocks, you will have to create a new cluster.</p> <p>If this parameter is left blank, a unique CIDR block is randomly allocated from 172.(16-31).0.0/16 and 10.(0   16   32   48   64   80   96   112).0.0/12.</p>
cidrs	Array of <a href="#">ContainerCIDR</a> objects	<p>List of container CIDR blocks In clusters of v1.21 or later, the <b>cidrs</b> field is used. When the cluster network type is <b>vpc-router</b>, you can configure a maximum of 20 container CIDR blocks. In cluster versions earlier than v1.21, if the <b>cidrs</b> field is used, the first CIDR element in the array is used as the container CIDR block.</p> <p>Once you have created a cluster that uses a container tunnel network, you cannot modify the CIDR blocks. Once you have created a cluster that uses a VPC network, you can add new CIDR blocks, but you are unable to modify the existing ones. If you need to adjust the CIDR blocks, you will have to create a new cluster.</p> <p>If this parameter is left blank, a unique CIDR block is randomly allocated from 172.(16-31).0.0/16 and 10.(0   16   32   48   64   80   96   112).0.0/12.</p>

**Table 4-50** ContainerCIDR

Parameter	Type	Description
cidr	String	<p>Container CIDR block. Recommended: 10.0.0.0/12-19, 172.16.0.0/16-19, or 192.168.0.0/16-19. If the selected CIDR block conflicts with existing ones, an error will be reported.</p>

**Table 4-51** EniNetwork

Parameter	Type	Description
eniSubnetId	String	IPv4 subnet ID of ENI subnet. (IPv6 is not supported and is being discarded.) You can obtain it in either of the following ways: <ul style="list-style-type: none"> <li>Method 1: Log in to the VPC console and click the target subnet on the <b>Subnets</b> page. You can view the IPv4 subnet ID on the displayed page.</li> <li>Method 2: Use the VPC API to obtain subnet IDs.</li> </ul>
eniSubnetCIDR	String	ENI subnet CIDR (being discarded)
subnets	Array of <a href="#">NetworkSubnet</a> objects	List of IPv4 subnet IDs

**Table 4-52** NetworkSubnet

Parameter	Type	Description
subnetID	String	IPv4 subnet ID of the subnet for creating master nodes. IPv6 is not supported. You can obtain it in either of the following ways: <ul style="list-style-type: none"> <li>Method 1: Log in to the VPC console and click the target subnet on the <b>Subnets</b> page. You can view the IPv4 subnet ID on the displayed page.</li> <li>Method 2: Use the VPC API to obtain subnets.</li> </ul>

**Table 4-53** ServiceNetwork

Parameter	Type	Description
IPv4CIDR	String	Value range of the Kubernetes clusterIP IPv4 CIDR blocks. If this parameter is not specified during cluster creation, the default value <b>10.247.0.0/16</b> will be used.

**Table 4-54 PublicAccess**

Parameter	Type	Description
cidrs	Array of strings	Trustlist of network CIDRs that are allowed to access cluster APIs. You are advised to allow the traffic from VPC and container network CIDRs. By default, no trustlist is configured, and the value is ["0.0.0.0/0"].  <b>NOTE</b> This parameter is valid only when a cluster is created.

**Table 4-55 Authentication**

Parameter	Type	Description
mode	String	Cluster authentication mode. <ul style="list-style-type: none"> <li>Clusters of Kubernetes v1.11 or earlier support <b>x509</b>, <b>rbac</b>, and <b>authenticating_proxy</b>. Defaults to <b>x509</b>.</li> <li>Clusters of Kubernetes v1.13 or later support <b>rbac</b> and <b>authenticating_proxy</b>. Defaults to <b>rbac</b>.</li> </ul>
authenticatingProxy	<b>AuthenticatingProxy</b> object	Configuration related to the authenticating_proxy mode. This field is mandatory when the authentication mode is <b>authenticating_proxy</b> .

**Table 4-56 AuthenticatingProxy**

Parameter	Type	Description
ca	String	X509 CA certificate (Base64-encoded) configured in authenticating_proxy mode. This field is mandatory when the cluster authentication mode is <b>authenticating_proxy</b> . Maximum size: 1 MB
cert	String	Client certificate issued by the X509 CA certificate configured in authenticating_proxy mode, which is used for authentication from kube-apiserver to the extended API server. (The value must be Base64-encoded.) This field is mandatory when the cluster authentication mode is <b>authenticating_proxy</b> . Maximum size: 1 MB

Parameter	Type	Description
privateKey	String	Private key of the client certificate issued by the X509 CA certificate configured in <code>authenticating_proxy</code> mode, which is used for authentication from kube-apiserver to the extended API server. The private key used by the Kubernetes cluster does not support password encryption. Use an unencrypted private key. (The value must be Base64-encoded.) This field is mandatory when the cluster authentication mode is <b>authenticating_proxy</b> . Maximum size: 1 MB

**Table 4-57** MasterSpec

Parameter	Type	Description
availabilityZone	String	AZ

**Table 4-58** ResourceTag

Parameter	Type	Description
key	String	Key. <ul style="list-style-type: none"> <li>• Cannot be null. Max characters: 128.</li> <li>• Use letters, digits, and spaces in UTF-8 format.</li> <li>• Can contain the following special characters: <code>._:=-@</code></li> <li>• Cannot start with <code>_sys_</code>.</li> </ul>
value	String	Value. <ul style="list-style-type: none"> <li>• Can be null but not the default. Max characters: 255.</li> <li>• Use letters, digits, and spaces in UTF-8 format.</li> <li>• Can contain the following special characters: <code>._:/=-@</code>.</li> </ul>



**Table 4-59** ClusterExtendParam

Parameter	Type	Description
clusterAZ	String	<p>AZ of master nodes in a cluster.</p> <ul style="list-style-type: none"> <li>• <b>multi_az</b>: (Optional) The cluster will span across AZs. Multiple AZs can be configured only when a cluster with multiple master nodes is used.</li> <li>• <i>AZ of the dedicated cloud computing pool</i>: The cluster will be deployed in the DeC AZ. It is mandatory for dedicated CCE clusters.</li> </ul>
dssMasterVolumes	String	<p>Whether the system and data disks of a master node use dedicated distributed storage. If this parameter is omitted or left unspecified, EVS disks are used by default.</p> <p>This parameter is mandatory for dedicated CCE clusters. It is in the following format:  <code>&lt;rootVol.dssPoolID&gt;.&lt;rootVol.volType&gt;;&lt;dataVol.dssPoolID&gt;.&lt;dataVol.volType&gt;</code></p> <p>Field description:</p> <ul style="list-style-type: none"> <li>• <b>rootVol</b> is the system disk. <b>dataVol</b> is the data disk.</li> <li>• <b>dssPoolID</b> indicates the ID of the DSS storage pool.</li> <li>• <b>volType</b> indicates the storage volume type of the DSS storage pool, such as SAS, SSD, SATA, ESSD, GPSSD, ESSD2, and GPSSD2.</li> </ul> <p>Example:  <code>c950ee97-587c-4f24-8a74-3367e3da570f.sas;6dbc2f4-1507-44f8-ac0d-eed1d2608d38.ssd</code></p> <p><b>NOTE</b>  This field cannot be configured for non-dedicated CCE clusters.</p>
enterpriseProjectId	String	<p>ID of the enterprise project that a cluster belongs to</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• An enterprise project can be configured only after the enterprise project function is enabled.</li> </ul>

Parameter	Type	Description
kubeProxyMode	String	<p>Service forwarding mode. Two modes are available:</p> <ul style="list-style-type: none"> <li>• <b>iptables</b>: Traditional kube-proxy uses iptables rules to implement Service load balancing. In this mode, too many iptables rules will be generated when many Services are deployed. In addition, non-incremental updates will cause latency and even tangible performance issues in the case of service traffic spikes.</li> <li>• <b>ipvs</b>: Optimized kube-proxy mode with higher throughput and faster speed. This mode supports incremental updates and can keep connections uninterrupted during Service updates. It is suitable for large-sized clusters.</li> </ul> <p><b>NOTE</b> This parameter has been deprecated. If this parameter and <b>kubeProxyMode</b> in <b>ClusterSpec</b> are specified at the same time, the latter is used.</p>
clusterExternalIP	String	EIP of the master node
alpha.cce/fixPoolMask	String	<p>Number of mask bits of the fixed IP address pool of the container networks. It is supported only by the VPC network model (vpc-router).</p> <p>The number of mask bits and <b>maxPods</b> configured during node creation determine how many IP addresses can be allocated to containers from a node and how many pods can be created in a node. For details, see <a href="#">Maximum Number of Pods That Can Be Created on a Node</a>.</p> <p>For integer characters, the value ranges from 24 to 28.</p>
decMasterFlavor	String	Specifications of the master node in the dedicated hybrid cluster.
dockerUmaskMode	String	Default UmaskMode configuration of Docker in a cluster. The value can be <b>secure</b> or <b>normal</b> . If this parameter is not specified, <b>normal</b> is used by default.

Parameter	Type	Description
kubernetes.io/cpuManagerPolicy	String	Cluster CPU management policy. The value can be none (or null) or static. The default value is none (or null). <ul style="list-style-type: none"> <li>• <b>none</b> or <b>null</b>: disables pods from exclusively occupying CPUs. Select this option if you want a large pool of shareable CPU cores.</li> <li>• <b>static</b>: enables pods to exclusively occupy CPUs. Select this option if your workload is sensitive to CPU cache and scheduling latency.</li> </ul>
upgradefrom	String	Records of how the cluster is upgraded to the current version

**Table 4-60** PackageConfiguration

Parameter	Type	Description
name	String	Component name.
configurations	Array of <a href="#">Configuratio nItem</a> objects	Component configuration item.

**Table 4-61** ConfigurationItem

Parameter	Type	Description
name	String	Component configuration item name.
value	Object	Component configuration item value.

**Table 4-62** ClusterStatus

Parameter	Type	Description
phase	String	<p>Cluster status. Options:</p> <ul style="list-style-type: none"> <li>● <b>Available:</b> The cluster is running properly.</li> <li>● <b>Unavailable:</b> The cluster is exhibiting unexpected behavior and needs to be manually deleted.</li> <li>● <b>ScalingUp:</b> Nodes are being added to the cluster.</li> <li>● <b>ScalingDown:</b> The cluster is being downsized to fewer nodes.</li> <li>● <b>Creating:</b> The cluster is being created.</li> <li>● <b>Deleting:</b> The cluster is being deleted.</li> <li>● <b>Upgrading:</b> The cluster is being upgraded.</li> <li>● <b>Resizing:</b> Cluster specifications are being changed.</li> <li>● <b>ResizeFailed:</b> Cluster specification changing is abnormal.</li> <li>● <b>RollingBack:</b> The cluster is being rolled back.</li> <li>● <b>RollbackFailed:</b> The cluster rollback is abnormal.</li> <li>● <b>Hibernating:</b> The cluster is being hibernated.</li> <li>● <b>Hibernation:</b> The cluster is in hibernation.</li> <li>● <b>Freezing:</b> The cluster is being frozen.</li> <li>● <b>Frozen:</b> The cluster has been frozen.</li> <li>● <b>UnFreezing:</b> The cluster is being unfrozen.</li> <li>● <b>Awaking:</b> The cluster is being woken up from hibernation.</li> <li>● <b>Empty:</b> The cluster does not have any resources. This field is discarded.</li> <li>● <b>Error:</b> Resources in the cluster are abnormal. You can manually delete the cluster.</li> </ul>

Parameter	Type	Description
jobID	String	<p>ID of the task associated with the cluster in the current state. Options:</p> <ul style="list-style-type: none"> <li>• ID of the associated task returned when creating a cluster. You can use it to obtain the auxiliary tasks for creating a cluster.</li> <li>• ID of the associated task returned when a cluster fails to be deleted or is deleted. If this parameter is not empty, you can use the task ID to obtain the auxiliary tasks for deleting a cluster.</li> </ul> <p><b>NOTE</b> Tasks are short-lived. Do not use task information in scenarios such as cluster status determination.</p>
reason	String	Reason of cluster state change. This parameter is returned if the cluster is not in the Available state.
message	String	Detailed information about why the cluster changes to the current state. This parameter is returned if the cluster is not in the Available state.
endpoints	Array of <a href="#">ClusterEndpoints</a> objects	Access address of kube-apiserver in the cluster.
deleteOption	Object	Whether to delete configurations. This parameter is contained only in the response to the deletion request.
deleteStatus	Object	Whether to delete the status information. This parameter is contained only in the response to the deletion request.

**Table 4-63** ClusterEndpoints

Parameter	Type	Description
url	String	Access address of kube-apiserver in the cluster.
type	String	<p>Type of the cluster access address.</p> <ul style="list-style-type: none"> <li>• <b>Internal:</b> address for internal network access</li> <li>• <b>External:</b> address for external network access</li> </ul>

## Example Requests

None

## Example Responses

**Status code: 200**

Information about the specified cluster is successfully obtained.

```
{
  "kind": "Cluster",
  "apiVersion": "v3",
  "metadata": {
    "alias": "mycluster",
    "name": "mycluster",
    "uid": "4d1ecb2c-229a-11e8-9c75-0255ac100ceb",
    "creationTimestamp": "2018-08-02 03:48:58.968214406 +0000 UTC",
    "updateTimestamp": "2018-08-02 04:05:29.386391813 +0000 UTC"
  },
  "spec": {
    "type": "VirtualMachine",
    "flavor": "cce.s1.small",
    "version": "v1.25",
    "platformVersion": "cce.6.0",
    "legacyVersion": "v1.25.6-r0",
    "description": "this is a demo cluster",
    "customSan": [ "192.168.1.0", "example.com" ],
    "hostNetwork": {
      "vpc": "4d1ecb2c-229a-11e8-9c75-0255ac100ceb",
      "subnet": "4d1ecb2c-229a-11e8-9c75-0255ac100ceb",
      "controlPlaneSecurityGroup": "14834251-ac69-460a-bfbd-7ac84274c52b"
    },
    "containerNetwork": {
      "mode": "overlay_l2",
      "cidr": "172.16.0.0/16"
    },
    "authentication": {
      "mode": "x509",
      "authenticatingProxy": { }
    },
    "billingMode": 0
  },
  "status": {
    "phase": "Available",
    "endpoints": [ {
      "url": "https://192.168.0.11:5443",
      "type": "Internal"
    } ]
  }
}
```

## Status Codes

Status Code	Description
200	Information about the specified cluster is successfully obtained.

## Error Codes

See [Error Codes](#).

## 4.2.3 Listing Clusters in a Specified Project

### Function

This API is used to obtain details about all clusters in a specified project.

### URI

GET /api/v3/projects/{project\_id}/clusters

**Table 4-64** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

**Table 4-65** Query Parameters

Parameter	Mandatory	Type	Description
detail	No	String	Whether the details about a cluster are obtained. If this parameter is set to <b>true</b> , the total number of nodes ( <b>totalNodesNumber</b> ), number of normal nodes ( <b>activeNodesNumber</b> ), total CPUs ( <b>totalNodesCPU</b> ), total memory size ( <b>totalNodesMemory</b> ), and installed add-ons ( <b>installedAddonInstances</b> ) will be added to the annotation. The details of installed add-ons include the add-on name ( <b>addonTemplateName</b> ), version ( <b>version</b> ), status ( <b>status</b> ), and node pool scaling supported or not ( <b>supportNodePoolScaleGroup</b> ).

Parameter	Mandatory	Type	Description
status	No	String	<p>Cluster status. Options:</p> <ul style="list-style-type: none"> <li>● <b>Available:</b> The cluster is running properly.</li> <li>● <b>Unavailable:</b> The cluster is exhibiting unexpected behavior and needs to be manually deleted.</li> <li>● <b>ScalingUp:</b> Nodes are being added to the cluster.</li> <li>● <b>ScalingDown:</b> The cluster is being downsized to fewer nodes.</li> <li>● <b>Creating:</b> The cluster is being created.</li> <li>● <b>Deleting:</b> The cluster is being deleted.</li> <li>● <b>Upgrading:</b> The cluster is being upgraded.</li> <li>● <b>Resizing:</b> Cluster specifications are being changed.</li> <li>● <b>ResizeFailed:</b> Cluster specification changing is abnormal.</li> <li>● <b>RollingBack:</b> The cluster is being rolled back.</li> <li>● <b>RollbackFailed:</b> The cluster rollback is abnormal.</li> <li>● <b>Hibernating:</b> The cluster is being hibernated.</li> <li>● <b>Hibernation:</b> The cluster is in hibernation.</li> <li>● <b>Freezing:</b> The cluster is being frozen.</li> <li>● <b>Frozen:</b> The cluster has been frozen.</li> <li>● <b>UnFreezing:</b> The cluster is being unfrozen.</li> <li>● <b>Awaking:</b> The cluster is being woken up from hibernation.</li> <li>● <b>Empty:</b> The cluster does not have any resources. This field is discarded.</li> </ul>



Parameter	Mandatory	Type	Description
			<ul style="list-style-type: none"> <li>• <b>Error:</b> Resources in the cluster are abnormal. You can manually delete the cluster.</li> </ul>
type	No	String	Cluster type. Options: <ul style="list-style-type: none"> <li>• <b>VirtualMachine:</b> CCE cluster</li> <li>• <b>ARM64:</b> Kunpeng cluster</li> </ul>
version	No	String	Cluster version filtering.

## Request Parameters

**Table 4-66** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

## Response Parameters

Status code: 200

**Table 4-67** Response body parameters

Parameter	Type	Description
kind	String	Api type
apiVersion	String	API version
items	Array of <a href="#">Cluster</a> objects	A list of details for all clusters in the current project. You can filter clusters by <b>items.metadata.name</b> .

**Table 4-68** Cluster

Parameter	Type	Description
kind	String	API type. The value is fixed at <b>Cluster</b> or <b>cluster</b> and cannot be changed.
apiVersion	String	API version. The value is fixed at <b>v3</b> and cannot be changed.
metadata	<b>ClusterMetadata</b> object	Basic information about a cluster. Metadata is a collection of attributes.
spec	<b>ClusterSpec</b> object	Detailed description of the cluster. CCE creates or updates objects by defining or updating <b>spec</b> .
status	<b>ClusterStatus</b> object	Cluster status and job ID of the cluster creation job.

**Table 4-69** ClusterMetadata

Parameter	Type	Description
name	String	Cluster name. Enter 4 to 128 characters, starting with a lowercase letter and not ending with a hyphen (-). Only lowercase letters, digits, and hyphens (-) are allowed.
uid	String	Cluster ID, which uniquely identifies a cluster. This ID is automatically generated after a cluster is created. Only the automatically generated ID will take effect.
alias	String	Alias of a cluster name displayed on the CCE console, and the name can be changed. Enter 4 to 128 characters starting with a letter and not ending with a hyphen (-). Only digits, letters, and hyphens (-) are allowed. A cluster alias must be unique. In the request body for creating or updating a cluster, if the cluster alias is not specified or is left blank, the alias of the cluster is the same as the cluster name. In the response body for obtaining a cluster, if the cluster alias is not configured, the cluster alias will not be returned.

Parameter	Type	Description
annotations	Map<String,String>	<p>Cluster annotations, in the format of key-value pairs.</p> <pre>"annotations": {   "key1" : "value1",   "key2" : "value2" }</pre> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• <b>annotations:</b> Does not label or select objects. The metadata in <b>annotations</b> may be small or large, structured or unstructured, and may include characters that are not allowed in labels.</li> <li>• This field is not stored in the database and is used only to specify the add-ons to be installed in the cluster.</li> <li>• Install ICAgent during cluster creation by adding the key-value pair "cluster.install.addons.external/install":{"addonTemplateName":"icagent"}".</li> </ul>
labels	Map<String,String>	<p>Cluster labels, in the format of key-value pairs.</p> <p><b>NOTE</b></p> <p>The value of this field is automatically generated by the system and is used by the frontend to identify the features supported by the cluster during the upgrade. Customized values are invalid.</p>
creationTimestamp	String	Time when the cluster was created.
updateTimestamp	String	Time when the cluster was updated.

**Table 4-70** ClusterSpec

Parameter	Type	Description
category	String	<p>Cluster type. Options:</p> <ul style="list-style-type: none"> <li>• <b>CCE:</b> CCE cluster CCE cluster supports hybrid deployment of VMs and BMSs, and heterogeneous nodes such as GPU and NPU nodes, allowing you to run your containers in a secure and stable container runtime environment based on a high-performance network model.</li> </ul>
type	String	<p>Master node architecture:</p> <ul style="list-style-type: none"> <li>• <b>VirtualMachine:</b> x86</li> </ul>

Parameter	Type	Description
flavor	String	<p>Cluster specifications. Specifications of clusters 1.15 or later versions can be changed after they are created. For details, see <a href="#">Modifying Cluster Specifications</a>. Options:</p> <ul style="list-style-type: none"> <li>● <b>cce.s1.small</b>: a small-scale CCE cluster with one master node and a maximum of 50 worker nodes</li> <li>● <b>cce.s1.medium</b>: a medium-scale CCE cluster with one master node and a maximum of 200 worker nodes</li> <li>● <b>cce.s2.small</b>: a small-scale CCE cluster with three master nodes and a maximum of 50 worker nodes</li> <li>● <b>cce.s2.medium</b>: a medium-scale CCE cluster with three master nodes and a maximum of 200 worker nodes</li> <li>● <b>cce.s2.large</b>: a large-scale CCE cluster with three master nodes and a maximum of 1,000 worker nodes</li> <li>● <b>cce.s2.xlarge</b>: an ultra-large-scale CCE cluster with three master nodes and a maximum of 2,000 worker nodes</li> </ul> <p><b>NOTE</b> The fields in the parameters are described as follows:</p> <ul style="list-style-type: none"> <li>● <b>s1</b>: specifies a cluster with one master node. If the master node is faulty, the cluster will become unavailable, but running workloads in the cluster are not affected.</li> <li>● <b>s2</b>: specifies an HA cluster with three master nodes. If one of the master nodes is faulty, the cluster is still available.</li> <li>● <b>small</b>: specifies that a cluster can manage a maximum of 50 worker nodes.</li> <li>● <b>medium</b>: specifies that a cluster can manage a maximum of 200 worker nodes.</li> <li>● <b>large</b>: specifies that a cluster can manage a maximum of 1,000 worker nodes.</li> <li>● <b>xlarge</b>: specifies that a cluster can manage a maximum of 2,000 worker nodes.</li> </ul>

Parameter	Type	Description
version	String	<p>Version of a cluster, which mirrors the baseline version of the Kubernetes community. The latest version is recommended.</p> <p>You can create clusters of the latest three versions on the CCE console. To learn which cluster versions are available, log in to the CCE console, create a cluster, and check supported cluster versions. You can call APIs to create clusters of other versions. However, these cluster versions will be gradually terminated. For details about the support policy, see the CCE announcement.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• If not specified, a cluster of the latest version will be created.</li> <li>• If a baseline cluster version is specified but the R version is not specified, a cluster of the latest R version will be created by default. It is a good practice not to specify the R version.</li> </ul>
platformVersion	String	<p>CCE cluster platform version, indicating the internal version under the cluster version (<b>version</b>). Platform versions are used to trace iterations in a major cluster version. They are unique within a major cluster version and recounted when the major cluster version changes. This parameter cannot be customized. When you create a cluster, the latest corresponding platform version is automatically selected.</p> <p>The format of <b>platformVersion</b> is <b>cce.X.Y</b>.</p> <ul style="list-style-type: none"> <li>• <b>X</b>: internal feature version, indicating changes in features, patches, or OS support in the cluster version. The value starts from <b>1</b> and increases monotonically.</li> <li>• <b>Y</b>: patch version of an internal feature version. It is used only for software package update after the feature version is brought online. No other modification is involved. The value starts from <b>0</b> and increases monotonically.</li> </ul>
legacyVersion	String	<p>(Discarded) Outdated version of CCE clusters, which serves no practical purpose and is only used to showcase the cluster version and platform version combination. The version number is globally unique. For example, if the cluster version is va.b and the platformVersion is cce.X.Y, the value of legacyVersion is va.b.X-rY.</p>

Parameter	Type	Description
description	String	Cluster description, for example, which purpose the cluster is intended to serve. By default, this field is left unspecified. To modify cluster description after the cluster is created, call the <a href="#">API for updating cluster information</a> or go to the cluster details page on the CCE console. Only UTF-8 encoding is supported.
customSan	Array of strings	Custom SAN field in the server certificate of the cluster API server, which must comply with the SSL and X509 format specifications. <ol style="list-style-type: none"> <li>1. Duplicate names are not allowed.</li> <li>2. Must comply with the IP address and domain name formats.</li> </ol> <p>Example:</p> <pre>SAN 1: DNS Name=example.com SAN 2: DNS Name=www.example.com SAN 3: DNS Name=example.net SAN 4: IP Address=93.184.216.34</pre>
ipv6enable	Boolean	Whether a cluster supports IPv6 addresses. It is supported by clusters 1.15 and later. <p><b>NOTE</b></p> <p>After IPv6 is enabled, iptables is not supported. The VPC network model does not support IPv4/IPv6 dual-stack.</p>
hostNetwork	<a href="#">HostNetwork</a> object	Node networking parameters, including VPC and subnet ID. This field is mandatory because nodes in a cluster communicate with each other by using a VPC.
containerNetwork	<a href="#">ContainerNetwork</a> object	Container networking parameters, including the container network model and container CIDR block.
serviceNetwork	<a href="#">ServiceNetwork</a> object	Service CIDR block, including IPv4 CIDR blocks.
publicAccess	<a href="#">PublicAccess</a> object	Cluster API access control
authentication	<a href="#">Authentication</a> object	Configurations of the cluster authentication mode.
billingMode	Integer	Billing mode of a cluster. <ul style="list-style-type: none"> <li>• <b>0</b>: pay-per-use</li> </ul> <p>Defaults to pay-per-use.</p>
masters	Array of <a href="#">MasterSpec</a> objects	Advanced configurations of the master nodes. If this parameter is not configured, no value is returned.

Parameter	Type	Description
kubernetesSvcIpRange	String	Service CIDR blocks for Kubernetes cluster IPs. This field is available only for clusters of v1.11.7 and later. If this parameter is not specified during cluster creation, the default value <b>10.247.0.0/16</b> will be used. This parameter is deprecated. Use <b>serviceNetwork</b> instead. The new field contains the IPv4 CIDR blocks.
clusterTags	Array of <b>ResourceTag</b> objects	Cluster resource tags.
kubeProxyMode	String	Service forwarding mode. Options: <ul style="list-style-type: none"> <li>• <b>iptables</b>: Traditional kube-proxy uses iptables rules to implement Service load balancing. In this mode, too many iptables rules will be generated when many Services are deployed. In addition, non-incremental updates will cause latency and even tangible performance issues in the case of service traffic spikes.</li> <li>• <b>ipvs</b>: Optimized kube-proxy mode with higher throughput and faster speed. This mode supports incremental updates and can keep connections uninterrupted during Service updates. It is suitable for large-sized clusters.</li> </ul> <p><b>NOTE</b> <b>iptables</b> is used by default.</p>
az	String	AZ. This field is returned only for a query.
extendParam	<b>ClusterExtendedParam</b> object	Extended field to decide whether the cluster will span across AZs or belong to a specified enterprise project, or whether a dedicated CCE cluster is to be created.
supportIstio	Boolean	Whether Istio is supported.
configurationsOverride	Array of <b>PackageConfiguration</b> objects	Cluster default component configuration override. If you specify a component or parameter that is not supported, the configuration item will be ignored.

**Table 4-71** HostNetwork

Parameter	Type	Description
vpc	String	ID of the VPC used to create a master node. You can obtain it in either of the following ways: <ul style="list-style-type: none"> <li>• Method 1: Log in to the VPC console and view the VPC ID in the VPC details page.</li> <li>• Method 2: Use the VPC API to obtain VPC IDs.</li> </ul>
subnet	String	Network ID of the subnet used to create a master node. You can obtain it in either of the following ways: <ul style="list-style-type: none"> <li>• Method 1: Log in to the VPC console and click the target subnet on the <b>Subnets</b> page. You can view the network ID on the displayed page.</li> <li>• Method 2: Use the VPC API to obtain subnets.</li> </ul>
SecurityGroup	String	Default worker node security group ID of the cluster. If specified, the cluster will be bound to the target security group. Otherwise, the system will automatically create a default worker node security group for you. The default worker node security group needs to allow access from certain ports to ensure normal communications.
controlPlaneSecurityGroup	String	Security group ID of the cluster control plane node, which is automatically generated after the cluster is created. A user-defined value will not take effect.

**Table 4-72** ContainerNetwork

Parameter	Type	Description
mode	String	Container network model. Select one of the following possible values: <ul style="list-style-type: none"> <li>• <b>overlay_l2</b>: an overlay_l2 network (container tunnel network) built for containers by using OpenVSwitch (OVS).</li> <li>• <b>vpc-router</b>: an underlay_l2 network built for containers by using IPvlan and custom VPC routes.</li> </ul>



Parameter	Type	Description
cidr	String	<p>Container CIDR block. Recommended: 10.0.0.0/12-19, 172.16.0.0/16-19, or 192.168.0.0/16-19. If the selected CIDR block conflicts with existing ones, an error will be reported.</p> <p>Not editable after the cluster is created. (This parameter has been discarded. A specified <b>cidrs</b> will make <b>cidr</b> invalid.) Once you have created a cluster that uses a VPC network, you can add new CIDR blocks, but you are unable to modify the existing ones. If you need to adjust the CIDR blocks, you will have to create a new cluster.</p> <p>If this parameter is left blank, a unique CIDR block is randomly allocated from 172.(16-31).0.0/16 and 10.(0   16   32   48   64   80   96   112).0.0/12.</p>
cidrs	Array of <a href="#">ContainerCIDR</a> objects	<p>List of container CIDR blocks In clusters of v1.21 or later, the <b>cidrs</b> field is used. When the cluster network type is <b>vpc-router</b>, you can configure a maximum of 20 container CIDR blocks. In cluster versions earlier than v1.21, if the <b>cidrs</b> field is used, the first CIDR element in the array is used as the container CIDR block.</p> <p>Once you have created a cluster that uses a container tunnel network, you cannot modify the CIDR blocks. Once you have created a cluster that uses a VPC network, you can add new CIDR blocks, but you are unable to modify the existing ones. If you need to adjust the CIDR blocks, you will have to create a new cluster.</p> <p>If this parameter is left blank, a unique CIDR block is randomly allocated from 172.(16-31).0.0/16 and 10.(0   16   32   48   64   80   96   112).0.0/12.</p>

**Table 4-73** ContainerCIDR

Parameter	Type	Description
cidr	String	<p>Container CIDR block. Recommended: 10.0.0.0/12-19, 172.16.0.0/16-19, or 192.168.0.0/16-19. If the selected CIDR block conflicts with existing ones, an error will be reported.</p>

**Table 4-74** EniNetwork

Parameter	Type	Description
eniSubnetId	String	IPv4 subnet ID of ENI subnet. (IPv6 is not supported and is being discarded.) You can obtain it in either of the following ways: <ul style="list-style-type: none"> <li>Method 1: Log in to the VPC console and click the target subnet on the <b>Subnets</b> page. You can view the IPv4 subnet ID on the displayed page.</li> <li>Method 2: Use the VPC API to obtain subnet IDs.</li> </ul>
eniSubnetCIDR	String	ENI subnet CIDR (being discarded)
subnets	Array of <a href="#">NetworkSubnet</a> objects	List of IPv4 subnet IDs

**Table 4-75** NetworkSubnet

Parameter	Type	Description
subnetID	String	IPv4 subnet ID of the subnet for creating master nodes. IPv6 is not supported. You can obtain it in either of the following ways: <ul style="list-style-type: none"> <li>Method 1: Log in to the VPC console and click the target subnet on the <b>Subnets</b> page. You can view the IPv4 subnet ID on the displayed page.</li> <li>Method 2: Use the VPC API to obtain subnets.</li> </ul>

**Table 4-76** ServiceNetwork

Parameter	Type	Description
IPv4CIDR	String	Value range of the Kubernetes clusterIP IPv4 CIDR blocks. If this parameter is not specified during cluster creation, the default value <b>10.247.0.0/16</b> will be used.

**Table 4-77 PublicAccess**

Parameter	Type	Description
cidrs	Array of strings	Trustlist of network CIDRs that are allowed to access cluster APIs. You are advised to allow the traffic from VPC and container network CIDRs. By default, no trustlist is configured, and the value is ["0.0.0.0/0"].  <b>NOTE</b> This parameter is valid only when a cluster is created.

**Table 4-78 Authentication**

Parameter	Type	Description
mode	String	Cluster authentication mode. <ul style="list-style-type: none"> <li>Clusters of Kubernetes v1.11 or earlier support <b>x509</b>, <b>rbac</b>, and <b>authenticating_proxy</b>. Defaults to <b>x509</b>.</li> <li>Clusters of Kubernetes v1.13 or later support <b>rbac</b> and <b>authenticating_proxy</b>. Defaults to <b>rbac</b>.</li> </ul>
authenticatingProxy	<b>AuthenticatingProxy</b> object	Configuration related to the authenticating_proxy mode. This field is mandatory when the authentication mode is <b>authenticating_proxy</b> .

**Table 4-79 AuthenticatingProxy**

Parameter	Type	Description
ca	String	X509 CA certificate (Base64-encoded) configured in authenticating_proxy mode. This field is mandatory when the cluster authentication mode is <b>authenticating_proxy</b> . Maximum size: 1 MB
cert	String	Client certificate issued by the X509 CA certificate configured in authenticating_proxy mode, which is used for authentication from kube-apiserver to the extended API server. (The value must be Base64-encoded.) This field is mandatory when the cluster authentication mode is <b>authenticating_proxy</b> . Maximum size: 1 MB

Parameter	Type	Description
privateKey	String	Private key of the client certificate issued by the X509 CA certificate configured in <code>authenticating_proxy</code> mode, which is used for authentication from kube-apiserver to the extended API server. The private key used by the Kubernetes cluster does not support password encryption. Use an unencrypted private key. (The value must be Base64-encoded.) This field is mandatory when the cluster authentication mode is <b>authenticating_proxy</b> . Maximum size: 1 MB

**Table 4-80** MasterSpec

Parameter	Type	Description
availabilityZone	String	AZ

**Table 4-81** ResourceTag

Parameter	Type	Description
key	String	Key. <ul style="list-style-type: none"> <li>• Cannot be null. Max characters: 128.</li> <li>• Use letters, digits, and spaces in UTF-8 format.</li> <li>• Can contain the following special characters: <code>._:=-@</code></li> <li>• Cannot start with <code>_sys_</code>.</li> </ul>
value	String	Value. <ul style="list-style-type: none"> <li>• Can be null but not the default. Max characters: 255.</li> <li>• Use letters, digits, and spaces in UTF-8 format.</li> <li>• Can contain the following special characters: <code>._:=-@</code>.</li> </ul>

**Table 4-82** ClusterExtendParam

Parameter	Type	Description
clusterAZ	String	<p>AZ of master nodes in a cluster.</p> <ul style="list-style-type: none"> <li>• <b>multi_az</b>: (Optional) The cluster will span across AZs. Multiple AZs can be configured only when a cluster with multiple master nodes is used.</li> <li>• <i>AZ of the dedicated cloud computing pool</i>: The cluster will be deployed in the DeC AZ. It is mandatory for dedicated CCE clusters.</li> </ul>
dssMasterVolumes	String	<p>Whether the system and data disks of a master node use dedicated distributed storage. If this parameter is omitted or left unspecified, EVS disks are used by default.</p> <p>This parameter is mandatory for dedicated CCE clusters. It is in the following format:  <code>&lt;rootVol.dssPoolID&gt;.&lt;rootVol.volType&gt;;&lt;dataVol.dssPoolID&gt;.&lt;dataVol.volType&gt;</code></p> <p>Field description:</p> <ul style="list-style-type: none"> <li>• <b>rootVol</b> is the system disk. <b>dataVol</b> is the data disk.</li> <li>• <b>dssPoolID</b> indicates the ID of the DSS storage pool.</li> <li>• <b>volType</b> indicates the storage volume type of the DSS storage pool, such as SAS, SSD, SATA, ESSD, GPSSD, ESSD2, and GPSSD2.</li> </ul> <p>Example:  <code>c950ee97-587c-4f24-8a74-3367e3da570f.sas;6dbc2f4-1507-44f8-ac0d-eed1d2608d38.ssd</code></p> <p><b>NOTE</b>  This field cannot be configured for non-dedicated CCE clusters.</p>
enterpriseProjectId	String	<p>ID of the enterprise project that a cluster belongs to</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• An enterprise project can be configured only after the enterprise project function is enabled.</li> </ul>

Parameter	Type	Description
kubeProxyMode	String	<p>Service forwarding mode. Two modes are available:</p> <ul style="list-style-type: none"> <li>• <b>iptables</b>: Traditional kube-proxy uses iptables rules to implement Service load balancing. In this mode, too many iptables rules will be generated when many Services are deployed. In addition, non-incremental updates will cause latency and even tangible performance issues in the case of service traffic spikes.</li> <li>• <b>ipvs</b>: Optimized kube-proxy mode with higher throughput and faster speed. This mode supports incremental updates and can keep connections uninterrupted during Service updates. It is suitable for large-sized clusters.</li> </ul> <p><b>NOTE</b> This parameter has been deprecated. If this parameter and <b>kubeProxyMode</b> in <b>ClusterSpec</b> are specified at the same time, the latter is used.</p>
clusterExternalIP	String	EIP of the master node
alpha.cce/fixPoolMask	String	<p>Number of mask bits of the fixed IP address pool of the container networks. It is supported only by the VPC network model (vpc-router).</p> <p>The number of mask bits and <b>maxPods</b> configured during node creation determine how many IP addresses can be allocated to containers from a node and how many pods can be created in a node. For details, see <a href="#">Maximum Number of Pods That Can Be Created on a Node</a>.</p> <p>For integer characters, the value ranges from 24 to 28.</p>
decMasterFlavor	String	Specifications of the master node in the dedicated hybrid cluster.
dockerUmaskMode	String	Default UmaskMode configuration of Docker in a cluster. The value can be <b>secure</b> or <b>normal</b> . If this parameter is not specified, <b>normal</b> is used by default.

Parameter	Type	Description
kubernetes.io/cpuManagerPolicy	String	Cluster CPU management policy. The value can be none (or null) or static. The default value is none (or null). <ul style="list-style-type: none"> <li>• <b>none</b> or <b>null</b>: disables pods from exclusively occupying CPUs. Select this option if you want a large pool of shareable CPU cores.</li> <li>• <b>static</b>: enables pods to exclusively occupy CPUs. Select this option if your workload is sensitive to CPU cache and scheduling latency.</li> </ul>
upgradefrom	String	Records of how the cluster is upgraded to the current version

**Table 4-83** PackageConfiguration

Parameter	Type	Description
name	String	Component name.
configurations	Array of <a href="#">Configuratio nItem</a> objects	Component configuration item.

**Table 4-84** ConfigurationItem

Parameter	Type	Description
name	String	Component configuration item name.
value	Object	Component configuration item value.

**Table 4-85** ClusterStatus

Parameter	Type	Description
phase	String	<p>Cluster status. Options:</p> <ul style="list-style-type: none"> <li>● <b>Available:</b> The cluster is running properly.</li> <li>● <b>Unavailable:</b> The cluster is exhibiting unexpected behavior and needs to be manually deleted.</li> <li>● <b>ScalingUp:</b> Nodes are being added to the cluster.</li> <li>● <b>ScalingDown:</b> The cluster is being downsized to fewer nodes.</li> <li>● <b>Creating:</b> The cluster is being created.</li> <li>● <b>Deleting:</b> The cluster is being deleted.</li> <li>● <b>Upgrading:</b> The cluster is being upgraded.</li> <li>● <b>Resizing:</b> Cluster specifications are being changed.</li> <li>● <b>ResizeFailed:</b> Cluster specification changing is abnormal.</li> <li>● <b>RollingBack:</b> The cluster is being rolled back.</li> <li>● <b>RollbackFailed:</b> The cluster rollback is abnormal.</li> <li>● <b>Hibernating:</b> The cluster is being hibernated.</li> <li>● <b>Hibernation:</b> The cluster is in hibernation.</li> <li>● <b>Freezing:</b> The cluster is being frozen.</li> <li>● <b>Frozen:</b> The cluster has been frozen.</li> <li>● <b>UnFreezing:</b> The cluster is being unfrozen.</li> <li>● <b>Awaking:</b> The cluster is being woken up from hibernation.</li> <li>● <b>Empty:</b> The cluster does not have any resources. This field is discarded.</li> <li>● <b>Error:</b> Resources in the cluster are abnormal. You can manually delete the cluster.</li> </ul>



Parameter	Type	Description
jobID	String	<p>ID of the task associated with the cluster in the current state. Options:</p> <ul style="list-style-type: none"> <li>• ID of the associated task returned when creating a cluster. You can use it to obtain the auxiliary tasks for creating a cluster.</li> <li>• ID of the associated task returned when a cluster fails to be deleted or is deleted. If this parameter is not empty, you can use the task ID to obtain the auxiliary tasks for deleting a cluster.</li> </ul> <p><b>NOTE</b> Tasks are short-lived. Do not use task information in scenarios such as cluster status determination.</p>
reason	String	Reason of cluster state change. This parameter is returned if the cluster is not in the Available state.
message	String	Detailed information about why the cluster changes to the current state. This parameter is returned if the cluster is not in the Available state.
endpoints	Array of <a href="#">ClusterEndpoints</a> objects	Access address of kube-apiserver in the cluster.
deleteOption	Object	Whether to delete configurations. This parameter is contained only in the response to the deletion request.
deleteStatus	Object	Whether to delete the status information. This parameter is contained only in the response to the deletion request.

**Table 4-86** ClusterEndpoints

Parameter	Type	Description
url	String	Access address of kube-apiserver in the cluster.
type	String	<p>Type of the cluster access address.</p> <ul style="list-style-type: none"> <li>• <b>Internal:</b> address for internal network access</li> <li>• <b>External:</b> address for external network access</li> </ul>

## Example Requests

None

## Example Responses

**Status code: 200**

The cluster list is successfully obtained.

```
{
  "kind": "Cluster",
  "apiVersion": "v3",
  "items": [ {
    "kind": "Cluster",
    "apiVersion": "v3",
    "metadata": {
      "alias": "mycluster",
      "name": "mycluster",
      "uid": "4d1ecb2c-229a-11e8-9c75-0255ac100ceb",
      "creationTimestamp": "2018-08-02 03:48:58.968214406 +0000 UTC",
      "updateTimestamp": "2018-08-02 04:05:29.386391813 +0000 UTC"
    },
    "spec": {
      "type": "VirtualMachine",
      "flavor": "cce.s1.small",
      "version": "v1.25",
      "platformVersion": "cce.6.0",
      "legacyVersion": "v1.25.6-r0",
      "description": "awesome cluster",
      "customSan": [ "192.168.1.0", "example.com" ],
      "hostNetwork": {
        "vpc": "f0c12911-4fdb-4284-9230-7ffb0860826a",
        "subnet": "ac274229-fd2e-4695-9f01-a0c1372b8006",
        "controlPlaneSecurityGroup": "14834251-ac69-460a-bfbd-7ac84274c52b"
      },
      "containerNetwork": {
        "mode": "overlay_l2",
        "cidr": "172.16.0.0/16"
      },
      "authentication": {
        "mode": "x509",
        "authenticatingProxy": { }
      },
      "billingMode": 0
    },
    "status": {
      "phase": "Available",
      "endpoints": [ {
        "url": "https://192.168.0.11:5443",
        "type": "Internal"
      } ]
    }
  } ]
}
```

## Status Codes

Status Code	Description
200	The cluster list is successfully obtained.

## Error Codes

See [Error Codes](#).

## 4.2.4 Updating a Specified Cluster

### Function

This API is used to update information about a specified cluster.

#### NOTE

The URL for cluster management is in the format of https://Endpoint/uri. In the URL, **uri** indicates the resource path, that is, the path for API access.

### URI

PUT /api/v3/projects/{project\_id}/clusters/{cluster\_id}

**Table 4-87** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

### Request Parameters

**Table 4-88** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

**Table 4-89** Request body parameters

Parameter	Mandatory	Type	Description
spec	Yes	<b>ClusterInformationSpec</b> object	Detailed cluster parameters
metadata	No	<b>ClusterMetadataForUpdate</b> object	Basic cluster information, including name-related fields

**Table 4-90** ClusterInformationSpec

Parameter	Mandatory	Type	Description
description	No	String	Cluster description. <ol style="list-style-type: none"> <li>1. A maximum of 200 characters are allowed. The value cannot contain the following special characters: ~\$%^&amp;*&lt;&gt;[]{}()'"#\</li> <li>2. Only clusters in the Available, ScalingUp, and ScalingDown states can be modified.</li> </ol>
customSan	No	Array of strings	Custom SAN field in the server certificate of the cluster API server, which must comply with the SSL and X509 format specifications. <ol style="list-style-type: none"> <li>1. Duplicate names are not allowed.</li> <li>2. Must comply with the IP address and domain name formats.</li> </ol> <p>Example:  SAN 1: DNS Name=example.com  SAN 2: DNS Name=www.example.com  SAN 3: DNS Name=example.net  SAN 4: IP Address=93.184.216.34</p>
containerNetwork	No	<b>ContainerNetworkUpdate</b> object	Container networking parameters, including information about the container CIDR block.

Parameter	Mandatory	Type	Description
eniNetwork	No	<b>EniNetworkUpdate</b> object	Cloud Native Network 2.0 network configuration, including the container subnet information of the CCE Turbo cluster.
hostNetwork	No	<b>hostNetwork</b> object	Node network parameters, including the default security group settings.

**Table 4-91** ContainerNetworkUpdate

Parameter	Mandatory	Type	Description
cidrs	No	Array of <b>ContainerCIDR</b> objects	List of container CIDR blocks. For a clusters of v1.21 or later, if it uses a VPC network, a maximum of 20 container CIDR blocks can be added incrementally.  The configuration cannot be modified after the cluster is updated.

**Table 4-92** ContainerCIDR

Parameter	Mandatory	Type	Description
cidr	Yes	String	Container CIDR block. Recommended: 10.0.0.0/12-19, 172.16.0.0/16-19, or 192.168.0.0/16-19. If the selected CIDR block conflicts with existing ones, an error will be reported.

**Table 4-93** EniNetworkUpdate

Parameter	Mandatory	Type	Description
subnets	No	Array of <a href="#">NetworkSubnet</a> objects	List of IPv4 subnet IDs. CCE Turbo clusters of version 1.19.10 or later support multiple container subnets and allow for incremental updates to the subnet list. The request body must contain all existing subnets.

**Table 4-94** NetworkSubnet

Parameter	Mandatory	Type	Description
subnetID	Yes	String	IPv4 subnet ID of the subnet for creating master nodes. IPv6 is not supported. You can obtain it in either of the following ways: <ul style="list-style-type: none"> <li>Method 1: Log in to the VPC console and click the target subnet on the <b>Subnets</b> page. You can view the IPv4 subnet ID on the displayed page.</li> <li>Method 2: Use the VPC API to obtain subnets.</li> </ul>

**Table 4-95** hostNetwork

Parameter	Mandatory	Type	Description
SecurityGroup	No	String	The default security group of the worker node in the cluster needs to allow traffic through some ports to ensure normal communication. The modified security group applies only to nodes newly created or accepted. For existing nodes, you need to manually modify the security group rules for them.

**Table 4-96** ClusterMetadataForUpdate

Parameter	Mandatory	Type	Description
alias	No	String	Alias of a cluster name. Enter 4 to 128 characters not starting or ending with a hyphen (-). Only digits, letters, and hyphens (-) are allowed. A cluster alias must be unique. If this parameter is left blank, the alias is not changed.

## Response Parameters

Status code: 200

**Table 4-97** Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at <b>Cluster</b> or <b>cluster</b> and cannot be changed.
apiVersion	String	API version. The value is fixed at <b>v3</b> and cannot be changed.
metadata	<b>ClusterMetadata</b> object	Basic information about a cluster. Metadata is a collection of attributes.
spec	<b>ClusterSpec</b> object	Detailed description of the cluster. CCE creates or updates objects by defining or updating <b>spec</b> .
status	<b>ClusterStatus</b> object	Cluster status and job ID of the cluster creation job.

**Table 4-98** ClusterMetadata

Parameter	Type	Description
name	String	Cluster name. Enter 4 to 128 characters, starting with a lowercase letter and not ending with a hyphen (-). Only lowercase letters, digits, and hyphens (-) are allowed.

Parameter	Type	Description
uid	String	Cluster ID, which uniquely identifies a cluster. This ID is automatically generated after a cluster is created. Only the automatically generated ID will take effect.
alias	String	Alias of a cluster name displayed on the CCE console, and the name can be changed. Enter 4 to 128 characters starting with a letter and not ending with a hyphen (-). Only digits, letters, and hyphens (-) are allowed. A cluster alias must be unique. In the request body for creating or updating a cluster, if the cluster alias is not specified or is left blank, the alias of the cluster is the same as the cluster name. In the response body for obtaining a cluster, if the cluster alias is not configured, the cluster alias will not be returned.
annotations	Map<String,String>	Cluster annotations, in the format of key-value pairs. <pre>"annotations": {   "key1" : "value1",   "key2" : "value2" }</pre> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• <b>annotations:</b> Does not label or select objects. The metadata in <b>annotations</b> may be small or large, structured or unstructured, and may include characters that are not allowed in labels.</li> <li>• This field is not stored in the database and is used only to specify the add-ons to be installed in the cluster.</li> <li>• Install ICAgent during cluster creation by adding the key-value pair "cluster.install.addons.external/install":{"addonTemplateName":"icagent"}".</li> </ul>
labels	Map<String,String>	Cluster labels, in the format of key-value pairs. <b>NOTE</b> The value of this field is automatically generated by the system and is used by the frontend to identify the features supported by the cluster during the upgrade. Customized values are invalid.
creationTimestamp	String	Time when the cluster was created.
updateTimestamp	String	Time when the cluster was updated.



**Table 4-99** ClusterSpec

Parameter	Type	Description
category	String	Cluster type. Options: <ul style="list-style-type: none"> <li>• <b>CCE</b>: CCE cluster CCE cluster supports hybrid deployment of VMs and BMSs, and heterogeneous nodes such as GPU and NPU nodes, allowing you to run your containers in a secure and stable container runtime environment based on a high-performance network model.</li> </ul>
type	String	Master node architecture: <ul style="list-style-type: none"> <li>• <b>VirtualMachine</b>: x86</li> </ul>

Parameter	Type	Description
flavor	String	<p>Cluster specifications. Specifications of clusters 1.15 or later versions can be changed after they are created. For details, see <a href="#">Modifying Cluster Specifications</a>. Options:</p> <ul style="list-style-type: none"> <li>• <b>cce.s1.small</b>: a small-scale CCE cluster with one master node and a maximum of 50 worker nodes</li> <li>• <b>cce.s1.medium</b>: a medium-scale CCE cluster with one master node and a maximum of 200 worker nodes</li> <li>• <b>cce.s2.small</b>: a small-scale CCE cluster with three master nodes and a maximum of 50 worker nodes</li> <li>• <b>cce.s2.medium</b>: a medium-scale CCE cluster with three master nodes and a maximum of 200 worker nodes</li> <li>• <b>cce.s2.large</b>: a large-scale CCE cluster with three master nodes and a maximum of 1,000 worker nodes</li> <li>• <b>cce.s2.xlarge</b>: an ultra-large-scale CCE cluster with three master nodes and a maximum of 2,000 worker nodes</li> </ul> <p><b>NOTE</b> The fields in the parameters are described as follows:</p> <ul style="list-style-type: none"> <li>• <b>s1</b>: specifies a cluster with one master node. If the master node is faulty, the cluster will become unavailable, but running workloads in the cluster are not affected.</li> <li>• <b>s2</b>: specifies an HA cluster with three master nodes. If one of the master nodes is faulty, the cluster is still available.</li> <li>• <b>small</b>: specifies that a cluster can manage a maximum of 50 worker nodes.</li> <li>• <b>medium</b>: specifies that a cluster can manage a maximum of 200 worker nodes.</li> <li>• <b>large</b>: specifies that a cluster can manage a maximum of 1,000 worker nodes.</li> <li>• <b>xlarge</b>: specifies that a cluster can manage a maximum of 2,000 worker nodes.</li> </ul>

Parameter	Type	Description
version	String	<p>Version of a cluster, which mirrors the baseline version of the Kubernetes community. The latest version is recommended.</p> <p>You can create clusters of the latest three versions on the CCE console. To learn which cluster versions are available, log in to the CCE console, create a cluster, and check supported cluster versions. You can call APIs to create clusters of other versions. However, these cluster versions will be gradually terminated. For details about the support policy, see the CCE announcement.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• If not specified, a cluster of the latest version will be created.</li> <li>• If a baseline cluster version is specified but the R version is not specified, a cluster of the latest R version will be created by default. It is a good practice not to specify the R version.</li> </ul>
platformVersion	String	<p>CCE cluster platform version, indicating the internal version under the cluster version (<b>version</b>). Platform versions are used to trace iterations in a major cluster version. They are unique within a major cluster version and recounted when the major cluster version changes. This parameter cannot be customized. When you create a cluster, the latest corresponding platform version is automatically selected.</p> <p>The format of <b>platformVersion</b> is <b>cce.X.Y</b>.</p> <ul style="list-style-type: none"> <li>• <b>X</b>: internal feature version, indicating changes in features, patches, or OS support in the cluster version. The value starts from <b>1</b> and increases monotonically.</li> <li>• <b>Y</b>: patch version of an internal feature version. It is used only for software package update after the feature version is brought online. No other modification is involved. The value starts from <b>0</b> and increases monotonically.</li> </ul>
legacyVersion	String	<p>(Discarded) Outdated version of CCE clusters, which serves no practical purpose and is only used to showcase the cluster version and platform version combination. The version number is globally unique. For example, if the cluster version is va.b and the platformVersion is cce.X.Y, the value of legacyVersion is va.b.X-rY.</p>

Parameter	Type	Description
description	String	Cluster description, for example, which purpose the cluster is intended to serve. By default, this field is left unspecified. To modify cluster description after the cluster is created, call the <a href="#">API for updating cluster information</a> or go to the cluster details page on the CCE console. Only UTF-8 encoding is supported.
customSan	Array of strings	Custom SAN field in the server certificate of the cluster API server, which must comply with the SSL and X509 format specifications. <ol style="list-style-type: none"> <li>1. Duplicate names are not allowed.</li> <li>2. Must comply with the IP address and domain name formats.</li> </ol> <p>Example:</p> <pre>SAN 1: DNS Name=example.com SAN 2: DNS Name=www.example.com SAN 3: DNS Name=example.net SAN 4: IP Address=93.184.216.34</pre>
ipv6enable	Boolean	Whether a cluster supports IPv6 addresses. It is supported by clusters 1.15 and later. <p><b>NOTE</b></p> <p>After IPv6 is enabled, iptables is not supported. The VPC network model does not support IPv4/IPv6 dual-stack.</p>
hostNetwork	<a href="#">HostNetwork</a> object	Node networking parameters, including VPC and subnet ID. This field is mandatory because nodes in a cluster communicate with each other by using a VPC.
containerNetwork	<a href="#">ContainerNetwork</a> object	Container networking parameters, including the container network model and container CIDR block.
serviceNetwork	<a href="#">ServiceNetwork</a> object	Service CIDR block, including IPv4 CIDR blocks.
publicAccess	<a href="#">PublicAccess</a> object	Cluster API access control
authentication	<a href="#">Authentication</a> object	Configurations of the cluster authentication mode.
billingMode	Integer	Billing mode of a cluster. <ul style="list-style-type: none"> <li>• <b>0</b>: pay-per-use</li> </ul> <p>Defaults to pay-per-use.</p>
masters	Array of <a href="#">MasterSpec</a> objects	Advanced configurations of the master nodes. If this parameter is not configured, no value is returned.

Parameter	Type	Description
kubernetesSvcIpRange	String	Service CIDR blocks for Kubernetes clusterIPs. This field is available only for clusters of v1.11.7 and later. If this parameter is not specified during cluster creation, the default value <b>10.247.0.0/16</b> will be used. This parameter is deprecated. Use <b>serviceNetwork</b> instead. The new field contains the IPv4 CIDR blocks.
clusterTags	Array of <b>ResourceTag</b> objects	Cluster resource tags.
kubeProxyMode	String	Service forwarding mode. Options: <ul style="list-style-type: none"> <li>• <b>iptables</b>: Traditional kube-proxy uses iptables rules to implement Service load balancing. In this mode, too many iptables rules will be generated when many Services are deployed. In addition, non-incremental updates will cause latency and even tangible performance issues in the case of service traffic spikes.</li> <li>• <b>ipvs</b>: Optimized kube-proxy mode with higher throughput and faster speed. This mode supports incremental updates and can keep connections uninterrupted during Service updates. It is suitable for large-sized clusters.</li> </ul> <p><b>NOTE</b> <b>iptables</b> is used by default.</p>
az	String	AZ. This field is returned only for a query.
extendParam	<b>ClusterExtendedParam</b> object	Extended field to decide whether the cluster will span across AZs or belong to a specified enterprise project, or whether a dedicated CCE cluster is to be created.
supportIstio	Boolean	Whether Istio is supported.
configurationsOverride	Array of <b>PackageConfiguration</b> objects	Cluster default component configuration override. If you specify a component or parameter that is not supported, the configuration item will be ignored.

**Table 4-100** HostNetwork

Parameter	Type	Description
vpc	String	ID of the VPC used to create a master node. You can obtain it in either of the following ways: <ul style="list-style-type: none"> <li>Method 1: Log in to the VPC console and view the VPC ID in the VPC details page.</li> <li>Method 2: Use the VPC API to obtain VPC IDs.</li> </ul>
subnet	String	Network ID of the subnet used to create a master node. You can obtain it in either of the following ways: <ul style="list-style-type: none"> <li>Method 1: Log in to the VPC console and click the target subnet on the <b>Subnets</b> page. You can view the network ID on the displayed page.</li> <li>Method 2: Use the VPC API to obtain subnets.</li> </ul>
SecurityGroup	String	Default worker node security group ID of the cluster. If specified, the cluster will be bound to the target security group. Otherwise, the system will automatically create a default worker node security group for you. The default worker node security group needs to allow access from certain ports to ensure normal communications.
controlPlaneSecurityGroup	String	Security group ID of the cluster control plane node, which is automatically generated after the cluster is created. A user-defined value will not take effect.

**Table 4-101** ContainerNetwork

Parameter	Type	Description
mode	String	Container network model. Select one of the following possible values: <ul style="list-style-type: none"> <li><b>overlay_l2</b>: an overlay_l2 network (container tunnel network) built for containers by using OpenVSwitch (OVS).</li> <li><b>vpc-router</b>: an underlay_l2 network built for containers by using IPvlan and custom VPC routes.</li> </ul>

Parameter	Type	Description
cidr	String	<p>Container CIDR block. Recommended: 10.0.0.0/12-19, 172.16.0.0/16-19, or 192.168.0.0/16-19. If the selected CIDR block conflicts with existing ones, an error will be reported.</p> <p>Not editable after the cluster is created. (This parameter has been discarded. A specified <b>cidrs</b> will make <b>cidr</b> invalid.) Once you have created a cluster that uses a VPC network, you can add new CIDR blocks, but you are unable to modify the existing ones. If you need to adjust the CIDR blocks, you will have to create a new cluster.</p> <p>If this parameter is left blank, a unique CIDR block is randomly allocated from 172.(16-31).0.0/16 and 10.(0   16   32   48   64   80   96   112).0.0/12.</p>
cidrs	Array of <a href="#">ContainerCIDR</a> objects	<p>List of container CIDR blocks In clusters of v1.21 or later, the <b>cidrs</b> field is used. When the cluster network type is <b>vpc-router</b>, you can configure a maximum of 20 container CIDR blocks. In cluster versions earlier than v1.21, if the <b>cidrs</b> field is used, the first CIDR element in the array is used as the container CIDR block.</p> <p>Once you have created a cluster that uses a container tunnel network, you cannot modify the CIDR blocks. Once you have created a cluster that uses a VPC network, you can add new CIDR blocks, but you are unable to modify the existing ones. If you need to adjust the CIDR blocks, you will have to create a new cluster.</p> <p>If this parameter is left blank, a unique CIDR block is randomly allocated from 172.(16-31).0.0/16 and 10.(0   16   32   48   64   80   96   112).0.0/12.</p>

**Table 4-102** ContainerCIDR

Parameter	Type	Description
cidr	String	<p>Container CIDR block. Recommended: 10.0.0.0/12-19, 172.16.0.0/16-19, or 192.168.0.0/16-19. If the selected CIDR block conflicts with existing ones, an error will be reported.</p>

**Table 4-103** EniNetwork

Parameter	Type	Description
eniSubnetId	String	IPv4 subnet ID of ENI subnet. (IPv6 is not supported and is being discarded.) You can obtain it in either of the following ways: <ul style="list-style-type: none"> <li>Method 1: Log in to the VPC console and click the target subnet on the <b>Subnets</b> page. You can view the IPv4 subnet ID on the displayed page.</li> <li>Method 2: Use the VPC API to obtain subnet IDs.</li> </ul>
eniSubnetCIDR	String	ENI subnet CIDR (being discarded)
subnets	Array of <a href="#">NetworkSubnet</a> objects	List of IPv4 subnet IDs

**Table 4-104** NetworkSubnet

Parameter	Type	Description
subnetID	String	IPv4 subnet ID of the subnet for creating master nodes. IPv6 is not supported. You can obtain it in either of the following ways: <ul style="list-style-type: none"> <li>Method 1: Log in to the VPC console and click the target subnet on the <b>Subnets</b> page. You can view the IPv4 subnet ID on the displayed page.</li> <li>Method 2: Use the VPC API to obtain subnets.</li> </ul>

**Table 4-105** ServiceNetwork

Parameter	Type	Description
IPv4CIDR	String	Value range of the Kubernetes clusterIP IPv4 CIDR blocks. If this parameter is not specified during cluster creation, the default value <b>10.247.0.0/16</b> will be used.



**Table 4-106** PublicAccess

Parameter	Type	Description
cidrs	Array of strings	Trustlist of network CIDRs that are allowed to access cluster APIs. You are advised to allow the traffic from VPC and container network CIDRs. By default, no trustlist is configured, and the value is ["0.0.0.0/0"].  <b>NOTE</b> This parameter is valid only when a cluster is created.

**Table 4-107** Authentication

Parameter	Type	Description
mode	String	Cluster authentication mode. <ul style="list-style-type: none"> <li>Clusters of Kubernetes v1.11 or earlier support <b>x509</b>, <b>rbac</b>, and <b>authenticating_proxy</b>. Defaults to <b>x509</b>.</li> <li>Clusters of Kubernetes v1.13 or later support <b>rbac</b> and <b>authenticating_proxy</b>. Defaults to <b>rbac</b>.</li> </ul>
authenticatingProxy	<b>AuthenticatingProxy</b> object	Configuration related to the authenticating_proxy mode. This field is mandatory when the authentication mode is <b>authenticating_proxy</b> .

**Table 4-108** AuthenticatingProxy

Parameter	Type	Description
ca	String	X509 CA certificate (Base64-encoded) configured in authenticating_proxy mode. This field is mandatory when the cluster authentication mode is <b>authenticating_proxy</b> . Maximum size: 1 MB
cert	String	Client certificate issued by the X509 CA certificate configured in authenticating_proxy mode, which is used for authentication from kube-apiserver to the extended API server. (The value must be Base64-encoded.) This field is mandatory when the cluster authentication mode is <b>authenticating_proxy</b> . Maximum size: 1 MB

Parameter	Type	Description
privateKey	String	Private key of the client certificate issued by the X509 CA certificate configured in <code>authenticating_proxy</code> mode, which is used for authentication from kube-apiserver to the extended API server. The private key used by the Kubernetes cluster does not support password encryption. Use an unencrypted private key. (The value must be Base64-encoded.) This field is mandatory when the cluster authentication mode is <b>authenticating_proxy</b> . Maximum size: 1 MB

**Table 4-109** MasterSpec

Parameter	Type	Description
availabilityZone	String	AZ

**Table 4-110** ResourceTag

Parameter	Type	Description
key	String	Key. <ul style="list-style-type: none"> <li>• Cannot be null. Max characters: 128.</li> <li>• Use letters, digits, and spaces in UTF-8 format.</li> <li>• Can contain the following special characters: <code>._:=-@</code></li> <li>• Cannot start with <code>_sys_</code>.</li> </ul>
value	String	Value. <ul style="list-style-type: none"> <li>• Can be null but not the default. Max characters: 255.</li> <li>• Use letters, digits, and spaces in UTF-8 format.</li> <li>• Can contain the following special characters: <code>._:=-@</code>.</li> </ul>

**Table 4-111** ClusterExtendParam

Parameter	Type	Description
clusterAZ	String	<p>AZ of master nodes in a cluster.</p> <ul style="list-style-type: none"> <li>• <b>multi_az</b>: (Optional) The cluster will span across AZs. Multiple AZs can be configured only when a cluster with multiple master nodes is used.</li> <li>• <i>AZ of the dedicated cloud computing pool</i>: The cluster will be deployed in the DeC AZ. It is mandatory for dedicated CCE clusters.</li> </ul>
dssMasterVolumes	String	<p>Whether the system and data disks of a master node use dedicated distributed storage. If this parameter is omitted or left unspecified, EVS disks are used by default.</p> <p>This parameter is mandatory for dedicated CCE clusters. It is in the following format:  <code>&lt;rootVol.dssPoolID&gt;.&lt;rootVol.volType&gt;;&lt;dataVol.dssPoolID&gt;.&lt;dataVol.volType&gt;</code></p> <p>Field description:</p> <ul style="list-style-type: none"> <li>• <b>rootVol</b> is the system disk. <b>dataVol</b> is the data disk.</li> <li>• <b>dssPoolID</b> indicates the ID of the DSS storage pool.</li> <li>• <b>volType</b> indicates the storage volume type of the DSS storage pool, such as SAS, SSD, SATA, ESSD, GPSSD, ESSD2, and GPSSD2.</li> </ul> <p>Example:  <code>c950ee97-587c-4f24-8a74-3367e3da570f.sas;6dbc2f4-1507-44f8-ac0d-eed1d2608d38.ssd</code></p> <p><b>NOTE</b>  This field cannot be configured for non-dedicated CCE clusters.</p>
enterpriseProjectId	String	<p>ID of the enterprise project that a cluster belongs to</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• An enterprise project can be configured only after the enterprise project function is enabled.</li> </ul>

Parameter	Type	Description
kubeProxyMode	String	<p>Service forwarding mode. Two modes are available:</p> <ul style="list-style-type: none"> <li>• <b>iptables</b>: Traditional kube-proxy uses iptables rules to implement Service load balancing. In this mode, too many iptables rules will be generated when many Services are deployed. In addition, non-incremental updates will cause latency and even tangible performance issues in the case of service traffic spikes.</li> <li>• <b>ipvs</b>: Optimized kube-proxy mode with higher throughput and faster speed. This mode supports incremental updates and can keep connections uninterrupted during Service updates. It is suitable for large-sized clusters.</li> </ul> <p><b>NOTE</b> This parameter has been deprecated. If this parameter and <b>kubeProxyMode</b> in <b>ClusterSpec</b> are specified at the same time, the latter is used.</p>
clusterExternalIP	String	EIP of the master node
alpha.cce/fixPoolMask	String	<p>Number of mask bits of the fixed IP address pool of the container networks. It is supported only by the VPC network model (vpc-router).</p> <p>The number of mask bits and <b>maxPods</b> configured during node creation determine how many IP addresses can be allocated to containers from a node and how many pods can be created in a node. For details, see <a href="#">Maximum Number of Pods That Can Be Created on a Node</a>.</p> <p>For integer characters, the value ranges from 24 to 28.</p>
decMasterFlavor	String	Specifications of the master node in the dedicated hybrid cluster.
dockerUmaskMode	String	Default UmaskMode configuration of Docker in a cluster. The value can be <b>secure</b> or <b>normal</b> . If this parameter is not specified, <b>normal</b> is used by default.

Parameter	Type	Description
kubernetes.io/ cpuManagerPolicy	String	Cluster CPU management policy. The value can be none (or null) or static. The default value is none (or null). <ul style="list-style-type: none"> <li>• <b>none</b> or <b>null</b>: disables pods from exclusively occupying CPUs. Select this option if you want a large pool of shareable CPU cores.</li> <li>• <b>static</b>: enables pods to exclusively occupy CPUs. Select this option if your workload is sensitive to CPU cache and scheduling latency.</li> </ul>
upgradefrom	String	Records of how the cluster is upgraded to the current version

**Table 4-112** PackageConfiguration

Parameter	Type	Description
name	String	Component name.
configurations	Array of <b>ConfigurationItem</b> objects	Component configuration item.

**Table 4-113** ConfigurationItem

Parameter	Type	Description
name	String	Component configuration item name.
value	Object	Component configuration item value.

**Table 4-114** ClusterStatus

Parameter	Type	Description
phase	String	<p>Cluster status. Options:</p> <ul style="list-style-type: none"> <li>● <b>Available:</b> The cluster is running properly.</li> <li>● <b>Unavailable:</b> The cluster is exhibiting unexpected behavior and needs to be manually deleted.</li> <li>● <b>ScalingUp:</b> Nodes are being added to the cluster.</li> <li>● <b>ScalingDown:</b> The cluster is being downsized to fewer nodes.</li> <li>● <b>Creating:</b> The cluster is being created.</li> <li>● <b>Deleting:</b> The cluster is being deleted.</li> <li>● <b>Upgrading:</b> The cluster is being upgraded.</li> <li>● <b>Resizing:</b> Cluster specifications are being changed.</li> <li>● <b>ResizeFailed:</b> Cluster specification changing is abnormal.</li> <li>● <b>RollingBack:</b> The cluster is being rolled back.</li> <li>● <b>RollbackFailed:</b> The cluster rollback is abnormal.</li> <li>● <b>Hibernating:</b> The cluster is being hibernated.</li> <li>● <b>Hibernation:</b> The cluster is in hibernation.</li> <li>● <b>Freezing:</b> The cluster is being frozen.</li> <li>● <b>Frozen:</b> The cluster has been frozen.</li> <li>● <b>UnFreezing:</b> The cluster is being unfrozen.</li> <li>● <b>Awaking:</b> The cluster is being woken up from hibernation.</li> <li>● <b>Empty:</b> The cluster does not have any resources. This field is discarded.</li> <li>● <b>Error:</b> Resources in the cluster are abnormal. You can manually delete the cluster.</li> </ul>

Parameter	Type	Description
jobID	String	<p>ID of the task associated with the cluster in the current state. Options:</p> <ul style="list-style-type: none"> <li>• ID of the associated task returned when creating a cluster. You can use it to obtain the auxiliary tasks for creating a cluster.</li> <li>• ID of the associated task returned when a cluster fails to be deleted or is deleted. If this parameter is not empty, you can use the task ID to obtain the auxiliary tasks for deleting a cluster.</li> </ul> <p><b>NOTE</b> Tasks are short-lived. Do not use task information in scenarios such as cluster status determination.</p>
reason	String	Reason of cluster state change. This parameter is returned if the cluster is not in the Available state.
message	String	Detailed information about why the cluster changes to the current state. This parameter is returned if the cluster is not in the Available state.
endpoints	Array of <a href="#">ClusterEndpoints</a> objects	Access address of kube-apiserver in the cluster.
deleteOption	Object	Whether to delete configurations. This parameter is contained only in the response to the deletion request.
deleteStatus	Object	Whether to delete the status information. This parameter is contained only in the response to the deletion request.

**Table 4-115** ClusterEndpoints

Parameter	Type	Description
url	String	Access address of kube-apiserver in the cluster.
type	String	<p>Type of the cluster access address.</p> <ul style="list-style-type: none"> <li>• <b>Internal:</b> address for internal network access</li> <li>• <b>External:</b> address for external network access</li> </ul>

## Example Requests

- Updating only the description of a cluster

```
{
  "spec" : {
    "description" : "new description"
  }
}
```

- Updating only the custom SAN of the cluster certificate

```
{
  "spec" : {
    "customSan" : [ "192.168.1.0", "example.com" ]
  }
}
```

- Updating both the cluster description and custom certificate SAN

```
{
  "spec" : {
    "description" : "new description",
    "customSan" : [ "192.168.1.0", "example.com" ]
  }
}
```

- Configure a CIDR block for a cluster that uses VPCs and is of a version later than v1.21.

```
{
  "spec" : {
    "containerNetwork" : {
      "cidrs" : [ {
        "cidr" : "10.10.0.0/16"
      }, {
        "cidr" : "10.11.0.0/16"
      } ]
    }
  }
}
```

- Modify the default security group of a node in a cluster.

```
/api/v3/projects/{project_id}/clusters/{cluster_id}

{
  "spec" : {
    "hostNetwork" : {
      "SecurityGroup" : "6ee29825-8f49-4796-b33a-fc76f84a59ae"
    }
  }
}
```

## Example Responses

**Status code: 200**

Information about the specified cluster is successfully updated.

```
{
  "kind" : "Cluster",
  "apiVersion" : "v3",
  "metadata" : {
    "alias" : "mycluster",
    "name" : "mycluster",
    "uid" : "4d1ecb2c-229a-11e8-9c75-0255ac100ceb",
    "creationTimestamp" : "2018-08-02 03:48:58.968214406 +0000 UTC",
    "updateTimestamp" : "2018-08-02 06:39:36.844676088 +0000 UTC"
  },
  "spec" : {
    "type" : "VirtualMachine",
    "flavor" : "cce.s1.small",
    "version" : "v1.7.3-r13",
  }
}
```



```

"description" : "new description",
"customSan" : [ "192.168.1.0", "example.com" ],
"hostNetwork" : {
  "vpc" : "4d1ecb2c-229a-11e8-9c75-0255ac100ceb",
  "subnet" : "4d1ecb2c-229a-11e8-9c75-0255ac100ceb",
  "SecurityGroup" : "6ee29825-8f49-4796-b33a-fc76f84a59ae"
},
"containerNetwork" : {
  "mode" : "overlay_l2",
  "cidr" : "172.17.0.0/16"
},
"authentication" : {
  "mode" : "x509",
  "authenticatingProxy" : { }
},
"billingMode" : 0
},
"status" : {
  "phase" : "Available",
  "endpoints" : [ {
    "url" : "https://192.168.0.11:5443",
    "type" : "Internal"
  } ]
}
}

```

## Status Codes

Status Code	Description
200	Information about the specified cluster is successfully updated.

## Error Codes

See [Error Codes](#).

## 4.2.5 Deleting a Cluster

### Function

This API is used to delete a specified cluster.

#### NOTE

The URL for cluster management is in the format of https://Endpoint/uri. In the URL, **uri** indicates the resource path, that is, the path for API access.

### URI

DELETE /api/v3/projects/{project\_id}/clusters/{cluster\_id}

**Table 4-116** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

**Table 4-117** Query Parameters

Parameter	Mandatory	Type	Description
delete_efs	No	String	Whether to delete an SFS Turbo volume. Options: <ul style="list-style-type: none"> <li>• <b>true</b> or <b>block</b>: The system starts to delete the object. If the deletion fails, subsequent processes are blocked.</li> <li>• <b>try</b>: The system starts to delete the object. If the deletion fails, no deletion retry is performed, and subsequent processes will proceed.</li> <li>• <b>false</b> or <b>skip</b>: The deletion is skipped. This is the default option.</li> </ul>

Parameter	Mandatory	Type	Description
delete_evs	No	String	<p>Whether to delete an EVS disk. Options:</p> <ul style="list-style-type: none"> <li>• <b>true</b> or <b>block</b>: The system starts to delete the object. If the deletion fails, subsequent processes are blocked.</li> <li>• <b>try</b>: The system starts to delete the object. If the deletion fails, no deletion retry is performed, and subsequent processes will proceed.</li> <li>• <b>false</b> or <b>skip</b>: The deletion is skipped. This is the default option.</li> </ul>
delete_net	No	String	<p>Whether to delete cluster Service or ingress resources, such as a load balancer. Options:</p> <ul style="list-style-type: none"> <li>• <b>true</b> or <b>block</b>: The system starts to delete the object. If the deletion fails, subsequent processes are blocked. This is the default option.</li> <li>• <b>try</b>: The system starts to delete the object. If the deletion fails, no deletion retry is performed, and subsequent processes will proceed.</li> <li>• <b>false</b> or <b>skip</b>: The deletion is skipped.</li> </ul>

Parameter	Mandatory	Type	Description
delete_obs	No	String	<p>Whether to delete an OBS volume. Options:</p> <ul style="list-style-type: none"> <li>• <b>true</b> or <b>block</b>: The system starts to delete the object. If the deletion fails, subsequent processes are blocked.</li> <li>• <b>try</b>: The system starts to delete the object. If the deletion fails, no deletion retry is performed, and subsequent processes will proceed.</li> <li>• <b>false</b> or <b>skip</b>: The deletion is skipped. This is the default option.</li> </ul>
delete_sfs	No	String	<p>Whether to delete an SFS volume. Options:</p> <ul style="list-style-type: none"> <li>• <b>true</b> or <b>block</b>: The system starts to delete the object. If the deletion fails, subsequent processes are blocked.</li> <li>• <b>try</b>: The system starts to delete the object. If the deletion fails, no deletion retry is performed, and subsequent processes will proceed.</li> <li>• <b>false</b> or <b>skip</b>: The deletion is skipped. This is the default option.</li> </ul>

Parameter	Mandatory	Type	Description
delete_sfs30	No	String	Whether to delete an SFS 3.0 volume. Options: <ul style="list-style-type: none"> <li>• <b>true</b> or <b>block</b>: The system starts to delete the object. If the deletion fails, subsequent processes are blocked.</li> <li>• <b>try</b>: The system starts to delete the object. If the deletion fails, no deletion retry is performed, and subsequent processes will proceed.</li> <li>• <b>false</b> or <b>skip</b>: The deletion is skipped. This is the default option.</li> </ul>
lts_reclaim_policy	No	String	Whether to delete an LTS resource, such as a log group or a log stream. Options: <ul style="list-style-type: none"> <li>• <b>Delete_Log_Group</b>: The system starts to delete a log group. If the deletion fails, no deletion retry is performed, and subsequent processes will proceed.</li> <li>• <b>Delete_Master_Log_Stream</b>: The system starts to delete a master log stream. If the deletion fails, no deletion retry is performed, and subsequent processes will proceed. This is the default option.</li> <li>• <b>Retain</b>: The deletion is skipped.</li> </ul>

## Request Parameters

**Table 4-118** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

## Response Parameters

Status code: 200

**Table 4-119** Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at <b>Cluster</b> or <b>cluster</b> and cannot be changed.
apiVersion	String	API version. The value is fixed at <b>v3</b> and cannot be changed.
metadata	<a href="#">ClusterMetadata</a> object	Basic information about a cluster. Metadata is a collection of attributes.
spec	<a href="#">ClusterSpec</a> object	Detailed description of the cluster. CCE creates or updates objects by defining or updating <b>spec</b> .
status	<a href="#">ClusterStatus</a> object	Cluster status and job ID of the cluster creation job.

**Table 4-120** ClusterMetadata

Parameter	Type	Description
name	String	Cluster name. Enter 4 to 128 characters, starting with a lowercase letter and not ending with a hyphen (-). Only lowercase letters, digits, and hyphens (-) are allowed.
uid	String	Cluster ID, which uniquely identifies a cluster. This ID is automatically generated after a cluster is created. Only the automatically generated ID will take effect.

Parameter	Type	Description
alias	String	<p>Alias of a cluster name displayed on the CCE console, and the name can be changed.</p> <p>Enter 4 to 128 characters starting with a letter and not ending with a hyphen (-). Only digits, letters, and hyphens (-) are allowed.</p> <p>A cluster alias must be unique.</p> <p>In the request body for creating or updating a cluster, if the cluster alias is not specified or is left blank, the alias of the cluster is the same as the cluster name. In the response body for obtaining a cluster, if the cluster alias is not configured, the cluster alias will not be returned.</p>
annotations	Map<String,String>	<p>Cluster annotations, in the format of key-value pairs.</p> <pre>"annotations": {   "key1": "value1",   "key2": "value2" }</pre> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• <b>annotations:</b> Does not label or select objects. The metadata in <b>annotations</b> may be small or large, structured or unstructured, and may include characters that are not allowed in labels.</li> <li>• This field is not stored in the database and is used only to specify the add-ons to be installed in the cluster.</li> <li>• Install ICAgent during cluster creation by adding the key-value pair "cluster.install.addons.external/install":{"addonTemplateName":"icagent"}".</li> </ul>
labels	Map<String,String>	<p>Cluster labels, in the format of key-value pairs.</p> <p><b>NOTE</b></p> <p>The value of this field is automatically generated by the system and is used by the frontend to identify the features supported by the cluster during the upgrade. Customized values are invalid.</p>
creationTimestamp	String	Time when the cluster was created.
updateTimestamp	String	Time when the cluster was updated.

**Table 4-121** ClusterSpec

Parameter	Type	Description
category	String	Cluster type. Options: <ul style="list-style-type: none"> <li>• <b>CCE</b>: CCE cluster CCE cluster supports hybrid deployment of VMs and BMSs, and heterogeneous nodes such as GPU and NPU nodes, allowing you to run your containers in a secure and stable container runtime environment based on a high-performance network model.</li> </ul>
type	String	Master node architecture: <ul style="list-style-type: none"> <li>• <b>VirtualMachine</b>: x86</li> </ul>



Parameter	Type	Description
flavor	String	<p>Cluster specifications. Specifications of clusters 1.15 or later versions can be changed after they are created. For details, see <a href="#">Modifying Cluster Specifications</a>. Options:</p> <ul style="list-style-type: none"> <li>● <b>cce.s1.small</b>: a small-scale CCE cluster with one master node and a maximum of 50 worker nodes</li> <li>● <b>cce.s1.medium</b>: a medium-scale CCE cluster with one master node and a maximum of 200 worker nodes</li> <li>● <b>cce.s2.small</b>: a small-scale CCE cluster with three master nodes and a maximum of 50 worker nodes</li> <li>● <b>cce.s2.medium</b>: a medium-scale CCE cluster with three master nodes and a maximum of 200 worker nodes</li> <li>● <b>cce.s2.large</b>: a large-scale CCE cluster with three master nodes and a maximum of 1,000 worker nodes</li> <li>● <b>cce.s2.xlarge</b>: an ultra-large-scale CCE cluster with three master nodes and a maximum of 2,000 worker nodes</li> </ul> <p><b>NOTE</b> The fields in the parameters are described as follows:</p> <ul style="list-style-type: none"> <li>● <b>s1</b>: specifies a cluster with one master node. If the master node is faulty, the cluster will become unavailable, but running workloads in the cluster are not affected.</li> <li>● <b>s2</b>: specifies an HA cluster with three master nodes. If one of the master nodes is faulty, the cluster is still available.</li> <li>● <b>small</b>: specifies that a cluster can manage a maximum of 50 worker nodes.</li> <li>● <b>medium</b>: specifies that a cluster can manage a maximum of 200 worker nodes.</li> <li>● <b>large</b>: specifies that a cluster can manage a maximum of 1,000 worker nodes.</li> <li>● <b>xlarge</b>: specifies that a cluster can manage a maximum of 2,000 worker nodes.</li> </ul>

Parameter	Type	Description
version	String	<p>Version of a cluster, which mirrors the baseline version of the Kubernetes community. The latest version is recommended.</p> <p>You can create clusters of the latest three versions on the CCE console. To learn which cluster versions are available, log in to the CCE console, create a cluster, and check supported cluster versions. You can call APIs to create clusters of other versions. However, these cluster versions will be gradually terminated. For details about the support policy, see the CCE announcement.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• If not specified, a cluster of the latest version will be created.</li> <li>• If a baseline cluster version is specified but the R version is not specified, a cluster of the latest R version will be created by default. It is a good practice not to specify the R version.</li> </ul>
platformVersion	String	<p>CCE cluster platform version, indicating the internal version under the cluster version (<b>version</b>). Platform versions are used to trace iterations in a major cluster version. They are unique within a major cluster version and recounted when the major cluster version changes. This parameter cannot be customized. When you create a cluster, the latest corresponding platform version is automatically selected.</p> <p>The format of <b>platformVersion</b> is <b>cce.X.Y</b>.</p> <ul style="list-style-type: none"> <li>• <b>X</b>: internal feature version, indicating changes in features, patches, or OS support in the cluster version. The value starts from <b>1</b> and increases monotonically.</li> <li>• <b>Y</b>: patch version of an internal feature version. It is used only for software package update after the feature version is brought online. No other modification is involved. The value starts from <b>0</b> and increases monotonically.</li> </ul>
legacyVersion	String	<p>(Discarded) Outdated version of CCE clusters, which serves no practical purpose and is only used to showcase the cluster version and platform version combination. The version number is globally unique. For example, if the cluster version is va.b and the platformVersion is cce.X.Y, the value of legacyVersion is va.b.X-rY.</p>

Parameter	Type	Description
description	String	Cluster description, for example, which purpose the cluster is intended to serve. By default, this field is left unspecified. To modify cluster description after the cluster is created, call the <a href="#">API for updating cluster information</a> or go to the cluster details page on the CCE console. Only UTF-8 encoding is supported.
customSan	Array of strings	Custom SAN field in the server certificate of the cluster API server, which must comply with the SSL and X509 format specifications. <ol style="list-style-type: none"> <li>1. Duplicate names are not allowed.</li> <li>2. Must comply with the IP address and domain name formats.</li> </ol> <p>Example:</p> <pre>SAN 1: DNS Name=example.com SAN 2: DNS Name=www.example.com SAN 3: DNS Name=example.net SAN 4: IP Address=93.184.216.34</pre>
ipv6enable	Boolean	Whether a cluster supports IPv6 addresses. It is supported by clusters 1.15 and later. <p><b>NOTE</b></p> <p>After IPv6 is enabled, iptables is not supported. The VPC network model does not support IPv4/IPv6 dual-stack.</p>
hostNetwork	<a href="#">HostNetwork</a> object	Node networking parameters, including VPC and subnet ID. This field is mandatory because nodes in a cluster communicate with each other by using a VPC.
containerNetwork	<a href="#">ContainerNetwork</a> object	Container networking parameters, including the container network model and container CIDR block.
serviceNetwork	<a href="#">ServiceNetwork</a> object	Service CIDR block, including IPv4 CIDR blocks.
publicAccess	<a href="#">PublicAccess</a> object	Cluster API access control
authentication	<a href="#">Authentication</a> object	Configurations of the cluster authentication mode.
billingMode	Integer	Billing mode of a cluster. <ul style="list-style-type: none"> <li>• <b>0</b>: pay-per-use</li> </ul> <p>Defaults to pay-per-use.</p>
masters	Array of <a href="#">MasterSpec</a> objects	Advanced configurations of the master nodes. If this parameter is not configured, no value is returned.

Parameter	Type	Description
kubernetesSvcIpRange	String	Service CIDR blocks for Kubernetes clusterIPs. This field is available only for clusters of v1.11.7 and later. If this parameter is not specified during cluster creation, the default value <b>10.247.0.0/16</b> will be used. This parameter is deprecated. Use <b>serviceNetwork</b> instead. The new field contains the IPv4 CIDR blocks.
clusterTags	Array of <b>ResourceTag</b> objects	Cluster resource tags.
kubeProxyMode	String	Service forwarding mode. Options: <ul style="list-style-type: none"> <li>• <b>iptables</b>: Traditional kube-proxy uses iptables rules to implement Service load balancing. In this mode, too many iptables rules will be generated when many Services are deployed. In addition, non-incremental updates will cause latency and even tangible performance issues in the case of service traffic spikes.</li> <li>• <b>ipvs</b>: Optimized kube-proxy mode with higher throughput and faster speed. This mode supports incremental updates and can keep connections uninterrupted during Service updates. It is suitable for large-sized clusters.</li> </ul> <p><b>NOTE</b> <b>iptables</b> is used by default.</p>
az	String	AZ. This field is returned only for a query.
extendParam	<b>ClusterExtendedParam</b> object	Extended field to decide whether the cluster will span across AZs or belong to a specified enterprise project, or whether a dedicated CCE cluster is to be created.
supportIstio	Boolean	Whether Istio is supported.
configurationsOverride	Array of <b>PackageConfiguration</b> objects	Cluster default component configuration override. If you specify a component or parameter that is not supported, the configuration item will be ignored.

**Table 4-122** HostNetwork

Parameter	Type	Description
vpc	String	ID of the VPC used to create a master node. You can obtain it in either of the following ways: <ul style="list-style-type: none"> <li>Method 1: Log in to the VPC console and view the VPC ID in the VPC details page.</li> <li>Method 2: Use the VPC API to obtain VPC IDs.</li> </ul>
subnet	String	Network ID of the subnet used to create a master node. You can obtain it in either of the following ways: <ul style="list-style-type: none"> <li>Method 1: Log in to the VPC console and click the target subnet on the <b>Subnets</b> page. You can view the network ID on the displayed page.</li> <li>Method 2: Use the VPC API to obtain subnets.</li> </ul>
SecurityGroup	String	Default worker node security group ID of the cluster. If specified, the cluster will be bound to the target security group. Otherwise, the system will automatically create a default worker node security group for you. The default worker node security group needs to allow access from certain ports to ensure normal communications.
controlPlaneSecurityGroup	String	Security group ID of the cluster control plane node, which is automatically generated after the cluster is created. A user-defined value will not take effect.

**Table 4-123** ContainerNetwork

Parameter	Type	Description
mode	String	Container network model. Select one of the following possible values: <ul style="list-style-type: none"> <li><b>overlay_l2</b>: an overlay_l2 network (container tunnel network) built for containers by using OpenVSwitch (OVS).</li> <li><b>vpc-router</b>: an underlay_l2 network built for containers by using IPvlan and custom VPC routes.</li> </ul>

Parameter	Type	Description
cidr	String	<p>Container CIDR block. Recommended: 10.0.0.0/12-19, 172.16.0.0/16-19, or 192.168.0.0/16-19. If the selected CIDR block conflicts with existing ones, an error will be reported.</p> <p>Not editable after the cluster is created. (This parameter has been discarded. A specified <b>cidrs</b> will make <b>cidr</b> invalid.) Once you have created a cluster that uses a VPC network, you can add new CIDR blocks, but you are unable to modify the existing ones. If you need to adjust the CIDR blocks, you will have to create a new cluster.</p> <p>If this parameter is left blank, a unique CIDR block is randomly allocated from 172.(16-31).0.0/16 and 10.(0   16   32   48   64   80   96   112).0.0/12.</p>
cidrs	Array of <a href="#">ContainerCIDR</a> objects	<p>List of container CIDR blocks In clusters of v1.21 or later, the <b>cidrs</b> field is used. When the cluster network type is <b>vpc-router</b>, you can configure a maximum of 20 container CIDR blocks. In cluster versions earlier than v1.21, if the <b>cidrs</b> field is used, the first CIDR element in the array is used as the container CIDR block.</p> <p>Once you have created a cluster that uses a container tunnel network, you cannot modify the CIDR blocks. Once you have created a cluster that uses a VPC network, you can add new CIDR blocks, but you are unable to modify the existing ones. If you need to adjust the CIDR blocks, you will have to create a new cluster.</p> <p>If this parameter is left blank, a unique CIDR block is randomly allocated from 172.(16-31).0.0/16 and 10.(0   16   32   48   64   80   96   112).0.0/12.</p>

**Table 4-124** ContainerCIDR

Parameter	Type	Description
cidr	String	<p>Container CIDR block. Recommended: 10.0.0.0/12-19, 172.16.0.0/16-19, or 192.168.0.0/16-19. If the selected CIDR block conflicts with existing ones, an error will be reported.</p>

**Table 4-125** EniNetwork

Parameter	Type	Description
eniSubnetId	String	IPv4 subnet ID of ENI subnet. (IPv6 is not supported and is being discarded.) You can obtain it in either of the following ways: <ul style="list-style-type: none"> <li>Method 1: Log in to the VPC console and click the target subnet on the <b>Subnets</b> page. You can view the IPv4 subnet ID on the displayed page.</li> <li>Method 2: Use the VPC API to obtain subnet IDs.</li> </ul>
eniSubnetCIDR	String	ENI subnet CIDR (being discarded)
subnets	Array of <a href="#">NetworkSubnet</a> objects	List of IPv4 subnet IDs

**Table 4-126** NetworkSubnet

Parameter	Type	Description
subnetID	String	IPv4 subnet ID of the subnet for creating master nodes. IPv6 is not supported. You can obtain it in either of the following ways: <ul style="list-style-type: none"> <li>Method 1: Log in to the VPC console and click the target subnet on the <b>Subnets</b> page. You can view the IPv4 subnet ID on the displayed page.</li> <li>Method 2: Use the VPC API to obtain subnets.</li> </ul>

**Table 4-127** ServiceNetwork

Parameter	Type	Description
IPv4CIDR	String	Value range of the Kubernetes clusterIP IPv4 CIDR blocks. If this parameter is not specified during cluster creation, the default value <b>10.247.0.0/16</b> will be used.

**Table 4-128** PublicAccess

Parameter	Type	Description
cidrs	Array of strings	Trustlist of network CIDRs that are allowed to access cluster APIs. You are advised to allow the traffic from VPC and container network CIDRs. By default, no trustlist is configured, and the value is ["0.0.0.0/0"].  <b>NOTE</b> This parameter is valid only when a cluster is created.

**Table 4-129** Authentication

Parameter	Type	Description
mode	String	Cluster authentication mode. <ul style="list-style-type: none"> <li>Clusters of Kubernetes v1.11 or earlier support <b>x509</b>, <b>rbac</b>, and <b>authenticating_proxy</b>. Defaults to <b>x509</b>.</li> <li>Clusters of Kubernetes v1.13 or later support <b>rbac</b> and <b>authenticating_proxy</b>. Defaults to <b>rbac</b>.</li> </ul>
authenticatingProxy	<b>AuthenticatingProxy</b> object	Configuration related to the authenticating_proxy mode. This field is mandatory when the authentication mode is <b>authenticating_proxy</b> .

**Table 4-130** AuthenticatingProxy

Parameter	Type	Description
ca	String	X509 CA certificate (Base64-encoded) configured in authenticating_proxy mode. This field is mandatory when the cluster authentication mode is <b>authenticating_proxy</b> . Maximum size: 1 MB
cert	String	Client certificate issued by the X509 CA certificate configured in authenticating_proxy mode, which is used for authentication from kube-apiserver to the extended API server. (The value must be Base64-encoded.) This field is mandatory when the cluster authentication mode is <b>authenticating_proxy</b> . Maximum size: 1 MB



Parameter	Type	Description
privateKey	String	Private key of the client certificate issued by the X509 CA certificate configured in <code>authenticating_proxy</code> mode, which is used for authentication from kube-apiserver to the extended API server. The private key used by the Kubernetes cluster does not support password encryption. Use an unencrypted private key. (The value must be Base64-encoded.) This field is mandatory when the cluster authentication mode is <b>authenticating_proxy</b> . Maximum size: 1 MB

**Table 4-131** MasterSpec

Parameter	Type	Description
availabilityZone	String	AZ

**Table 4-132** ResourceTag

Parameter	Type	Description
key	String	Key. <ul style="list-style-type: none"> <li>• Cannot be null. Max characters: 128.</li> <li>• Use letters, digits, and spaces in UTF-8 format.</li> <li>• Can contain the following special characters: <code>._:=-@</code></li> <li>• Cannot start with <code>_sys_</code>.</li> </ul>
value	String	Value. <ul style="list-style-type: none"> <li>• Can be null but not the default. Max characters: 255.</li> <li>• Use letters, digits, and spaces in UTF-8 format.</li> <li>• Can contain the following special characters: <code>._:=-@</code>.</li> </ul>

**Table 4-133** ClusterExtendParam

Parameter	Type	Description
clusterAZ	String	<p>AZ of master nodes in a cluster.</p> <ul style="list-style-type: none"> <li>• <b>multi_az</b>: (Optional) The cluster will span across AZs. Multiple AZs can be configured only when a cluster with multiple master nodes is used.</li> <li>• <i>AZ of the dedicated cloud computing pool</i>: The cluster will be deployed in the DeC AZ. It is mandatory for dedicated CCE clusters.</li> </ul>
dssMasterVolumes	String	<p>Whether the system and data disks of a master node use dedicated distributed storage. If this parameter is omitted or left unspecified, EVS disks are used by default.</p> <p>This parameter is mandatory for dedicated CCE clusters. It is in the following format:  <code>&lt;rootVol.dssPoolID&gt;.&lt;rootVol.volType&gt;;&lt;dataVol.dssPoolID&gt;.&lt;dataVol.volType&gt;</code></p> <p>Field description:</p> <ul style="list-style-type: none"> <li>• <b>rootVol</b> is the system disk. <b>dataVol</b> is the data disk.</li> <li>• <b>dssPoolID</b> indicates the ID of the DSS storage pool.</li> <li>• <b>volType</b> indicates the storage volume type of the DSS storage pool, such as SAS, SSD, SATA, ESSD, GPSSD, ESSD2, and GPSSD2.</li> </ul> <p>Example:  <code>c950ee97-587c-4f24-8a74-3367e3da570f.sas;6dbc2f4-1507-44f8-ac0d-eed1d2608d38.ssd</code></p> <p><b>NOTE</b>  This field cannot be configured for non-dedicated CCE clusters.</p>
enterpriseProjectId	String	<p>ID of the enterprise project that a cluster belongs to</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• An enterprise project can be configured only after the enterprise project function is enabled.</li> </ul>

Parameter	Type	Description
kubeProxyMode	String	<p>Service forwarding mode. Two modes are available:</p> <ul style="list-style-type: none"> <li>• <b>iptables</b>: Traditional kube-proxy uses iptables rules to implement Service load balancing. In this mode, too many iptables rules will be generated when many Services are deployed. In addition, non-incremental updates will cause latency and even tangible performance issues in the case of service traffic spikes.</li> <li>• <b>ipvs</b>: Optimized kube-proxy mode with higher throughput and faster speed. This mode supports incremental updates and can keep connections uninterrupted during Service updates. It is suitable for large-sized clusters.</li> </ul> <p><b>NOTE</b> This parameter has been deprecated. If this parameter and <b>kubeProxyMode</b> in <b>ClusterSpec</b> are specified at the same time, the latter is used.</p>
clusterExternalIP	String	EIP of the master node
alpha.cce/fixPoolMask	String	<p>Number of mask bits of the fixed IP address pool of the container networks. It is supported only by the VPC network model (vpc-router).</p> <p>The number of mask bits and <b>maxPods</b> configured during node creation determine how many IP addresses can be allocated to containers from a node and how many pods can be created in a node. For details, see <a href="#">Maximum Number of Pods That Can Be Created on a Node</a>.</p> <p>For integer characters, the value ranges from 24 to 28.</p>
decMasterFlavor	String	Specifications of the master node in the dedicated hybrid cluster.
dockerUmaskMode	String	Default UmaskMode configuration of Docker in a cluster. The value can be <b>secure</b> or <b>normal</b> . If this parameter is not specified, <b>normal</b> is used by default.

Parameter	Type	Description
kubernetes.io/ cpuManagerPolicy	String	Cluster CPU management policy. The value can be none (or null) or static. The default value is none (or null). <ul style="list-style-type: none"> <li>• <b>none</b> or <b>null</b>: disables pods from exclusively occupying CPUs. Select this option if you want a large pool of shareable CPU cores.</li> <li>• <b>static</b>: enables pods to exclusively occupy CPUs. Select this option if your workload is sensitive to CPU cache and scheduling latency.</li> </ul>
upgradefrom	String	Records of how the cluster is upgraded to the current version

**Table 4-134** PackageConfiguration

Parameter	Type	Description
name	String	Component name.
configurations	Array of <a href="#">Configuratio nItem</a> objects	Component configuration item.

**Table 4-135** ConfigurationItem

Parameter	Type	Description
name	String	Component configuration item name.
value	Object	Component configuration item value.

**Table 4-136** ClusterStatus

Parameter	Type	Description
phase	String	<p>Cluster status. Options:</p> <ul style="list-style-type: none"> <li>● <b>Available:</b> The cluster is running properly.</li> <li>● <b>Unavailable:</b> The cluster is exhibiting unexpected behavior and needs to be manually deleted.</li> <li>● <b>ScalingUp:</b> Nodes are being added to the cluster.</li> <li>● <b>ScalingDown:</b> The cluster is being downsized to fewer nodes.</li> <li>● <b>Creating:</b> The cluster is being created.</li> <li>● <b>Deleting:</b> The cluster is being deleted.</li> <li>● <b>Upgrading:</b> The cluster is being upgraded.</li> <li>● <b>Resizing:</b> Cluster specifications are being changed.</li> <li>● <b>ResizeFailed:</b> Cluster specification changing is abnormal.</li> <li>● <b>RollingBack:</b> The cluster is being rolled back.</li> <li>● <b>RollbackFailed:</b> The cluster rollback is abnormal.</li> <li>● <b>Hibernating:</b> The cluster is being hibernated.</li> <li>● <b>Hibernation:</b> The cluster is in hibernation.</li> <li>● <b>Freezing:</b> The cluster is being frozen.</li> <li>● <b>Frozen:</b> The cluster has been frozen.</li> <li>● <b>UnFreezing:</b> The cluster is being unfrozen.</li> <li>● <b>Awaking:</b> The cluster is being woken up from hibernation.</li> <li>● <b>Empty:</b> The cluster does not have any resources. This field is discarded.</li> <li>● <b>Error:</b> Resources in the cluster are abnormal. You can manually delete the cluster.</li> </ul>

Parameter	Type	Description
jobID	String	<p>ID of the task associated with the cluster in the current state. Options:</p> <ul style="list-style-type: none"> <li>• ID of the associated task returned when creating a cluster. You can use it to obtain the auxiliary tasks for creating a cluster.</li> <li>• ID of the associated task returned when a cluster fails to be deleted or is deleted. If this parameter is not empty, you can use the task ID to obtain the auxiliary tasks for deleting a cluster.</li> </ul> <p><b>NOTE</b> Tasks are short-lived. Do not use task information in scenarios such as cluster status determination.</p>
reason	String	Reason of cluster state change. This parameter is returned if the cluster is not in the Available state.
message	String	Detailed information about why the cluster changes to the current state. This parameter is returned if the cluster is not in the Available state.
endpoints	Array of <a href="#">ClusterEndpoints</a> objects	Access address of kube-apiserver in the cluster.
deleteOption	Object	Whether to delete configurations. This parameter is contained only in the response to the deletion request.
deleteStatus	Object	Whether to delete the status information. This parameter is contained only in the response to the deletion request.

**Table 4-137** ClusterEndpoints

Parameter	Type	Description
url	String	Access address of kube-apiserver in the cluster.
type	String	<p>Type of the cluster access address.</p> <ul style="list-style-type: none"> <li>• <b>Internal:</b> address for internal network access</li> <li>• <b>External:</b> address for external network access</li> </ul>

## Example Requests

None

## Example Responses

**Status code: 200**

The job for deleting a cluster is successfully delivered.

```
{
  "kind": "Cluster",
  "apiVersion": "v3",
  "metadata": {
    "alias": "mycluster",
    "name": "mycluster",
    "uid": "fc563b3c-9552-11e8-8beb-0255ac106311",
    "creationTimestamp": "2018-08-01 06:20:28.81667161 +0000 UTC",
    "updateTimestamp": "2018-08-01 09:23:38.944333282 +0000 UTC"
  },
  "spec": {
    "type": "VirtualMachine",
    "flavor": "cce.s1.small",
    "version": "v1.7.3-r13",
    "description": "new description",
    "hostNetwork": {
      "vpc": "cbed56e8-03e7-4304-a477-b54bef0857c3",
      "subnet": "5de50062-2be2-4a52-893e-e0906e3e9c9d"
    },
    "containerNetwork": {
      "mode": "overlay_l2",
      "cidr": "172.16.0.0/16"
    },
    "authentication": {
      "mode": "x509",
      "authenticatingProxy": { }
    },
    "billingMode": 0
  },
  "status": {
    "phase": "Available",
    "jobID": "e8ebf96c-956d-11e8-a949-0255ac10575d",
    "endpoints": [ {
      "url": "https://192.168.0.16:5443",
      "type": "Internal"
    } ]
  }
}
```

## Status Codes

Status Code	Description
200	The job for deleting a cluster is successfully delivered.

## Error Codes

See [Error Codes](#).

## 4.2.6 Hibernating a Cluster

### Function

This API is used to hibernate a running cluster. After a cluster is hibernated, master node resources stop being billed.

### Constraints

1. After a cluster is hibernated, resources such as workloads cannot be created or managed in the cluster.

### URI

POST /api/v3/projects/{project\_id}/clusters/{cluster\_id}/operation/hibernate

**Table 4-138** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

### Request Parameters

**Table 4-139** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

### Response Parameters

None



## Example Requests

None

## Example Responses

None

## Status Codes

Status Code	Description
200	The cluster hibernation job is successfully delivered. Keep querying the cluster status. When the cluster status changes to <b>Hibernation</b> , the cluster is hibernated.

## Error Codes

See [Error Codes](#).

## 4.2.7 Waking Up a Cluster

### Function

This API is used to wake up a hibernated cluster. After the cluster is woken up, master node resources continue to be billed.

### URI

POST /api/v3/projects/{project\_id}/clusters/{cluster\_id}/operation/awake

**Table 4-140** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

## Request Parameters

**Table 4-141** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

## Response Parameters

None

## Example Requests

None

## Example Responses

None

## Status Codes

Status Code	Description
200	The cluster wakeup job is successfully delivered. Keep querying the cluster status. When the cluster status changes to <b>Available</b> , the cluster is woken up successfully.

## Error Codes

See [Error Codes](#).

## 4.2.8 Obtaining a Cluster Certificate

### Function

This API is used to obtain a certificate of a specified cluster.

## Constraints

This API is applicable to clusters of v1.13 and later.

## URI

POST /api/v3/projects/{project\_id}/clusters/{cluster\_id}/clustercert

**Table 4-142** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

## Request Parameters

**Table 4-143** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

**Table 4-144** Request body parameters

Parameter	Mandatory	Type	Description
duration	Yes	Integer	Validity period of the cluster certificate. The minimum value is 1 day and the maximum value is 5 years. Therefore, the value ranges from 1 to 1827. (The unit is day. The actual limit depends on the number of leap years in the five years. For example, if there is a leap year in the five years, the upper limit is 1826 days.) If this parameter is set to -1, the maximum value is 5 years.

## Response Parameters

**Status code: 200**

**Table 4-145** Response header parameters

Parameter	Type	Description
Port-ID	String	Port ID of the cluster master node

**Table 4-146** Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at <b>Config</b> and cannot be changed.
apiVersion	String	API version. The value is fixed at <b>v1</b> .
preferences	Object	This field is not used currently and is left unspecified by default.
clusters	Array of <b>Clusters</b> objects	Cluster list
users	Array of <b>Users</b> objects	Certificate information and client key information of a specified user
contexts	Array of <b>Contexts</b> objects	Context list

Parameter	Type	Description
current-context	String	Current context. If <b>publicip</b> (VM EIP) exists, the value is <b>external</b> . If <b>publicip</b> does not exist, the value is <b>internal</b> .

**Table 4-147** Clusters

Parameter	Type	Description
name	String	Cluster name. <ul style="list-style-type: none"> <li>If <b>publicip</b> does not exist (that is, no VM EIP exists), there is only one cluster in the cluster list, and the value of this parameter is <b>internalCluster</b>.</li> <li>If <b>publicip</b> exists (that is, the EIP exists), there are at least two clusters in the cluster list, and the value of this parameter is <b>externalCluster</b>.</li> </ul>
cluster	<b>ClusterCert</b> object	Cluster information

**Table 4-148** ClusterCert

Parameter	Type	Description
server	String	Server IP address
certificate-authority-data	String	Certificate authorization data
insecure-skip-tls-verify	Boolean	Whether to skip the server certificate verification. If the cluster type is <b>externalCluster</b> , the value is <b>true</b> .

**Table 4-149** Users

Parameter	Type	Description
name	String	The value is fixed at <b>user</b> .
user	<b>User</b> object	Certificate information and client key information of a specified user

**Table 4-150** User

Parameter	Type	Description
client-certificate-data	String	Client certificate
client-key-data	String	PEM encoding data from the TLS client key file

**Table 4-151** Contexts

Parameter	Type	Description
name	String	Context name. <ul style="list-style-type: none"> <li>If <b>publicip</b> does not exist (that is, no VM EIP exists), there is only one cluster in the cluster list, and the value of this parameter is <b>internal</b>.</li> <li>If <b>publicip</b> exists (that is, the EIP exists), there are at least two clusters in the cluster list, and the value of this field for all extension contexts is <b>external</b>.</li> </ul>
context	<b>Context</b> object	Context information

**Table 4-152** Context

Parameter	Type	Description
cluster	String	Cluster context
user	String	User context

## Example Requests

Applying for a cluster access certificate valid for 30 days

```
{
  "duration" : 30
}
```

## Example Responses

**Status code: 200**

The certificate of the specified cluster is successfully obtained. For details about the certificate file format, see the Kubernetes v1.Config structure.

```
{
  "kind": "Config",
  "apiVersion": "v1",
  "preferences": { },
  "clusters": [ {
    "name": "internalCluster",
    "cluster": {
      "server": "https://192.168.1.7:5443",
      "certificate-authority-data": "Q2VydGlmaWNhdGU6*****FTkQgQ0VSVElGSUNBVEUtLS0tLQo="
    }
  } ],
  "users": [ {
    "name": "user",
    "user": {
      "client-certificate-data": "LS0tLS1CRUdJTjBDR*****QVRFLS0tLS0K",
      "client-key-data": "LS0tLS1CRUdJTj*****BLRVktLS0tLQo="
    }
  } ],
  "contexts": [ {
    "name": "internal",
    "context": {
      "cluster": "internalCluster",
      "user": "user"
    }
  } ],
  "current-context": "internal"
}
```

## Status Codes

Status Code	Description
200	The certificate of the specified cluster is successfully obtained. For details about the certificate file format, see the Kubernetes v1.Config structure.

## Error Codes

See [Error Codes](#).

## 4.2.9 Modifying Cluster Specifications

### Function

This API is used to modify the specifications of a cluster.

#### NOTE

- The URL for cluster management is in the format of **https://Endpoint/uri**, in which *uri* indicates the resource path, that is, the path for API access.

### URI

POST /api/v3/projects/{project\_id}/clusters/{cluster\_id}/operation/resize

**Table 4-153** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

## Request Parameters

**Table 4-154** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .



**Table 4-155** Request body parameters

Parameter	Mandatory	Type	Description
flavorResize	Yes	String	<p>The target flavor to be changed. You can only adjust the maximum number of worker nodes in a cluster. The number of master nodes is fixed and cannot be altered, and the cluster flavors cannot be downgraded. Suppose the original cluster flavor is <b>cce.s2.medium</b>. In that case, you can only upgrade it to <b>cce.s2.large</b> or higher, but you cannot downgrade it to <b>cce.s2.small</b> or <b>cce.s1.medium</b>.</p> <ul style="list-style-type: none"> <li>• <b>cce.s1.small</b>: a small-scale CCE cluster with one master node and a maximum of 50 worker nodes</li> <li>• <b>cce.s1.medium</b>: a medium-scale CCE cluster with one master node and a maximum of 200 worker nodes</li> <li>• <b>cce.s2.small</b>: a small-scale CCE cluster with three master nodes and a maximum of 50 worker nodes</li> <li>• <b>cce.s2.medium</b>: a medium-scale CCE cluster with three master nodes and a maximum of 200 worker nodes</li> <li>• <b>cce.s2.large</b>: a large-scale CCE cluster with three master nodes and a maximum of 1,000 worker nodes</li> <li>• <b>cce.s2.xlarge</b>: an ultra-large-scale CCE cluster with three master nodes and a maximum of 2,000 worker nodes</li> </ul>

Parameter	Mandatory	Type	Description
			<p><b>NOTE</b></p> <p>The fields in the parameters are described as follows:</p> <ul style="list-style-type: none"> <li>• <b>s1</b>: specifies a cluster with one master node. If the master node is faulty, the cluster will become unavailable, but running workloads in the cluster are not affected.</li> <li>• <b>s2</b>: specifies an HA cluster with three master nodes. If one of the master nodes is faulty, the cluster is still available.</li> <li>• <b>small</b>: specifies that a cluster can manage a maximum of 50 worker nodes.</li> <li>• <b>medium</b>: specifies that a cluster can manage a maximum of 200 worker nodes.</li> <li>• <b>large</b>: specifies that a cluster can manage a maximum of 1,000 worker nodes.</li> <li>• <b>xlarge</b>: specifies that a cluster can manage a maximum of 2,000 worker nodes.</li> </ul>
extendParam	No	<a href="#">extendParam</a> object	-

**Table 4-156** extendParam

Parameter	Mandatory	Type	Description
decMasterFlavor	No	String	Specifications of the master node in the dedicated hybrid cluster

## Response Parameters

Status code: 201

**Table 4-157** Response body parameters

Parameter	Type	Description
jobID	String	ID of the job

## Example Requests

- Modifying the specifications of the pay-per-use cluster

```
POST /api/v3/projects/{project_id}/clusters/{cluster_id}/operation/resize
```

```
{
  "flavorResize" : "cce.s1.medium"
}
```

## Example Responses

**Status code: 201**

The job for modifying the specifications of the pay-per-use cluster is delivered.

```
{
  "jobID" : "13b8d958-8fcf-11ed-aef3-0255ac1001bd"
}
```

## Status Codes

Status Code	Description
201	The job for modifying the specifications of the pay-per-use cluster is delivered.

## Error Codes

See [Error Codes](#).

### 4.2.10 Querying a Job

#### Function

This API is used to obtain information about a job via the job ID returned after a query request is delivered.

#### NOTE

- The URL for cluster management is in the format of **https://Endpoint/uri**. In the URL, **uri** indicates the resource path, that is, the path for API access.
- You can call this API when:
  - Creating or deleting a cluster
  - Creating or deleting a node

#### URI

```
GET /api/v3/projects/{project_id}/jobs/{job_id}
```

**Table 4-158** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
job_id	Yes	String	Job ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

## Request Parameters

**Table 4-159** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

## Response Parameters

**Status code: 200**

**Table 4-160** Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at <b>Job</b> and cannot be changed.
apiVersion	String	API version. The value is fixed at <b>v3</b> and cannot be changed.
metadata	<a href="#">JobMetadata</a> object	Job metadata
spec	<a href="#">JobSpec</a> object	Detailed job parameters

Parameter	Type	Description
status	<b>JobStatus</b> object	Job status

**Table 4-161** JobSpec

Parameter	Type	Description
type	String	Job type. An example value is <b>CreateCluster</b> .
clusterUID	String	ID of the cluster where the job runs.
resourceID	String	ID of the resource on which the job is executed.
resourceName	String	Name of the resource on which the job is executed.
extendParam	Map<String,String>	Extended parameters
subJobs	Array of <b>Job</b> objects	Subjob list. <ul style="list-style-type: none"> <li>The list contains details about all subjobs.</li> <li>Generally, a cluster/node creation job consists of multiple subjobs. The job is complete only after all subjobs are complete.</li> </ul>

**Table 4-162** Job

Parameter	Type	Description
kind	String	API type. The value is fixed at <b>Job</b> and cannot be changed.
apiVersion	String	API version. The value is fixed at <b>v3</b> and cannot be changed.
metadata	<b>JobMetadata</b> object	Job metadata
spec	<b>JobSpec</b> object	Detailed job parameters
status	<b>JobStatus</b> object	Job status

**Table 4-163** JobMetadata

Parameter	Type	Description
uid	String	Job ID
creationTimes tamp	String	Time when the job was created
updateTimest amp	String	Time when the job was updated

**Table 4-164** JobStatus

Parameter	Type	Description
phase	String	Job status. Possible values: <ul style="list-style-type: none"> <li>● JobPhaseInitializing JobPhase = "Initializing"</li> <li>● JobPhaseRunning JobPhase = "Running"</li> <li>● JobPhaseFailed JobPhase = "Failed"</li> <li>● JobPhaseSuccess JobPhase = "Success"</li> </ul>
reason	String	Reason why the job is in the current state.

## Example Requests

None

## Example Responses

**Status code: 200**

The progress of the specified job is successfully obtained.

```
{
  "kind": "Job",
  "apiVersion": "v3",
  "metadata": {
    "uid": "354331b2c-229a-11e8-9c75-0255ac100ceb",
    "creationTimestamp": "2018-08-02 08:12:40.672772389 +0000 UTC",
    "updateTimestamp": "2018-08-02 08:21:50.478108569 +0000 UTC"
  },
  "spec": {
    "type": "CreateCluster",
    "clusterUID": "4d1ecb2c-229a-11e8-9c75-0255ac100ceb",
    "resourceID": "6f4dcb2c-229a-11e8-9c75-0255ac100ceb",
    "resourceName": "cluster-name",
    "extendParam": {
      "serverID": "bc467e3a-2338-11e8-825b-0255ac100c13"
    }
  },
  "subJobs": [ {
    "kind": "Job",
    "apiVersion": "v3",
    "metadata": {
      "uid": "fd474fab-9606-11e8-baa9-0255ac10215d",
      "creationTimestamp": "2018-08-02 03:52:34.615819618 +0000 UTC",
```

```

    "updateTimestamp" : "2018-08-02 04:05:29.196243031 +0000 UTC"
  },
  "spec" : {
    "type" : "InstallMaster",
    "clusterUID" : "fcc72de0-9606-11e8-baa8-0255ac10215d",
    "resourceID" : "fd3b4ac0-9606-11e8-baa8-0255ac10215d",
    "extendParam" : {
      "serverID" : "fd3b4ac0-9606-11e8-baa8-0255ac10215d"
    }
  },
  "status" : {
    "phase" : "Success"
  }
}, {
  "kind" : "Job",
  "apiVersion" : "v3",
  "metadata" : {
    "uid" : "fd474f82-9606-11e8-baa8-0255ac10215d",
    "creationTimestamp" : "2018-08-02 03:52:33.859150791 +0000 UTC",
    "updateTimestamp" : "2018-08-02 03:52:34.615655429 +0000 UTC"
  },
  "spec" : {
    "type" : "CreatePSMCert",
    "clusterUID" : "fcc72de0-9606-11e8-baa8-0255ac10215d"
  },
  "status" : {
    "phase" : "Success"
  }
} ]
},
"status" : {
  "phase" : "Running",
  "reason" : ""
}
}

```

## Status Codes

Status Code	Description
200	The progress of the specified job is successfully obtained.

## Error Codes

See [Error Codes](#).

### 4.2.11 Binding/Unbinding Public API Server Address

#### Function

This API is used to bind or unbind the public API server address of the cluster by cluster ID.

#### NOTE

The URL for cluster management is in the format of `https://Endpoint/uri`. In the URL, **uri** indicates the resource path, that is, the path for API access.

## URI

PUT /api/v3/projects/{project\_id}/clusters/{cluster\_id}/mastereip

**Table 4-165** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

## Request Parameters

**Table 4-166** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

**Table 4-167** Request body parameters

Parameter	Mandatory	Type	Description
spec	Yes	<a href="#">MasterEIPRequestSpec</a> object	Parameters in the request for binding or unbinding the public APIServer address of a cluster



**Table 4-168** MasterEIPRequestSpec

Parameter	Mandatory	Type	Description
action	No	String	Binding or unbinding. Mandatory. <ul style="list-style-type: none"> <li>Binding: The value is fixed at {"action":"bind"}.</li> <li>Unbinding: The value is fixed at {"action":"unbind"}.</li> </ul>
spec	No	<b>spec</b> object	Configuration attributes of the elastic IP address to be bound
bandwidth	No	String	Bandwidth (This field has expired and is not recommended.)
elasticip	No	String	ENI IP (This field has expired and is not recommended.)

**Table 4-169** spec

Parameter	Mandatory	Type	Description
id	No	String	Specifies the ENI ID. This parameter is mandatory for binding an ENI and is invalid for unbinding an ENI.

## Response Parameters

Status code: 200

**Table 4-170** Response body parameters

Parameter	Type	Description
metadata	<b>Metadata</b> object	Basic information about the object. Metadata is a collection of attributes.
spec	<b>MasterEIPResponseSpec</b> object	Configuration of the bound public API Server address of a cluster
status	<b>status</b> object	Status information

**Table 4-171** Metadata

Parameter	Type	Description
uid	String	Unique ID.
name	String	Resource name
labels	Map<String,String>	Resource labels in key-value pairs. This is a reserved field and does not take effect.
annotations	Map<String,String>	Resource annotations in the format of key-value pairs.
updateTimestamp	String	Update time.
creationTimestamp	String	Creation time.

**Table 4-172** MasterEIPResponseSpec

Parameter	Type	Description
action	String	A binding operation
spec	<b>spec</b> object	Configuration attributes of the elastic IP address to be bound
elasticIp	String	EIP

**Table 4-173** spec

Parameter	Type	Description
id	String	ENI ID
eip	<b>EipSpec</b> object	EIP details
isDynamic	Boolean	Dynamic provisioning or not

**Table 4-174** EipSpec

Parameter	Type	Description
bandwidth	<b>bandwidth</b> object	Bandwidth information

**Table 4-175** bandwidth

Parameter	Type	Description
size	Integer	Bandwidth size
sharetype	String	Bandwidth type

**Table 4-176** status

Parameter	Type	Description
privateEndpoint	String	Private IP for accessing the cluster (VIP in the case of an HA cluster)
publicEndpoint	String	Public IP for accessing the cluster

## Example Requests

Bind the public API server address to a cluster.

```
{
  "spec" : {
    "action" : "bind",
    "spec" : {
      "id" : "a757a69e-f920-455a-b1ba-d7a22db0fd50"
    }
  }
}
```

## Example Responses

**Status code: 200**

This indicates that the public API server address of the cluster is successfully bound. No response body is returned if the unbinding is successful.

```
{
  "metadata" : { },
  "spec" : {
    "action" : "bind",
    "spec" : {
      "id" : "a757a69e-f920-455a-b1ba-d7a22db0fd50",
      "eip" : {
        "bandwidth" : {
          "size" : 5,
          "sharetype" : "PER"
        }
      }
    },
    "isDynamic" : false
  },
  "elasticIp" : "8.8.8.8"
},
"status" : {
  "privateEndpoint" : "https://192.168.3.238:5443",
  "publicEndpoint" : "https://8.8.8.8:5443"
}
}
```

## Status Codes

Status Code	Description
200	This indicates that the public API server address of the cluster is successfully bound. No response body is returned if the unbinding is successful.

## Error Codes

See [Error Codes](#).

# 4.3 Node Management

## 4.3.1 Creating a Node

### Function

This API is used to create a node in a specified cluster.

#### NOTE

- If there is no cluster, [create one](#).
- The URL for cluster management is in the format of **https://Endpoint/uri**. In the URL, **uri** indicates the resource path, that is, the path for API access.

### Constraints

Only KVM nodes can be created. Non-KVM nodes cannot be used after being created.

### URI

POST /api/v3/projects/{project\_id}/clusters/{cluster\_id}/nodes

**Table 4-177** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

**Table 4-178** Query Parameters

Parameter	Mandatory	Type	Description
nodepoolScaleUp	No	String	Whether the request is delivered by the node pool. If the value is not <b>NodepoolScaleUp</b> , the number of pods in the corresponding node pool is automatically updated.

## Request Parameters

**Table 4-179** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

**Table 4-180** Request body parameters

Parameter	Mandatory	Type	Description
kind	Yes	String	API type. The value is fixed at <b>Node</b> and cannot be changed.
apiVersion	Yes	String	API version. The value is fixed at <b>v3</b> and cannot be changed.
metadata	No	<a href="#">NodeMetadata</a> object	Node metadata, which is a collection of attributes.
spec	Yes	<a href="#">NodeSpec</a> object	Detailed description of the node. CCE creates or updates objects by defining or updating <b>spec</b> .

**Table 4-181** NodeMetadata

Parameter	Mandatory	Type	Description
name	No	String	<p>Node name.</p> <p><b>NOTE</b> Enter 1 to 56 characters starting with a lowercase letter and not ending with a hyphen (-). Only lowercase letters, digits, hyphens (-), and periods (.) are allowed. If <b>name</b> is not configured or is left blank, the node name is generated based on the default rule, which is <i>Cluster name-Random characters</i>. Only the first 36 characters are kept for a long cluster name. If the number of nodes (<b>count</b>) is greater than 1, some random characters are added to the end of the entered node name according to the default rule, which is <i>User-defined name-Random characters</i>. If the length of the entered node name exceeds 50 characters, the system uses the first 50 characters and adds a random character string to the end of the name.</p>
uid	No	String	<p>Node ID, which is unique and automatically generated after the resource is created. A user-defined ID will not take effect.</p>
labels	No	Map<String,String>	<p>CCE node label (not the native Kubernetes label).</p> <p>Labels are used to select objects that meet certain criteria. A label is a key-value pair.</p> <p>Example:  <pre>"labels": {   "key": "value" }</pre> </p>

Parameter	Mandatory	Type	Description
annotations	No	Map<String,String>	<p>CCE node annotations, in key-value pairs (not the native Kubernetes annotations).</p> <p>Example:</p> <pre>"annotations": {   "key1" : "value1",   "key2" : "value2" }</pre> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li><b>annotations</b> are not used to identify or select objects. The metadata in <b>annotations</b> may be small or large, structured or unstructured, and may include characters that are not allowed in labels.</li> <li>This parameter is used only for query and cannot be input through a request. Entered data of this parameter is invalid.</li> </ul>
creationTimestamp	No	String	Time when the object was created. The value is automatically generated after the object is created. A user-defined value will not take effect.
updateTimestamp	No	String	Time when the object was updated. The value is automatically generated after the object is created. A user-defined value will not take effect.
ownerReference	No	<b>ownerReference</b> object	Owner object

**Table 4-182** ownerReference

Parameter	Mandatory	Type	Description
nodepoolName	No	String	Node pool name
nodepoolID	No	String	Node pool UID

**Table 4-183** NodeSpec

Parameter	Mandatory	Type	Description
flavor	Yes	String	Node specifications. For details about the node specifications supported by CCE, see <a href="#">Node Flavor Description</a> .
az	Yes	String	AZ where the node to be created is located. You need to specify the AZ name. A node created using APIs does not support random AZs. <b>NOTE</b> <ul style="list-style-type: none"> <li>When you are creating a node pool and configuring a scaling group, this parameter cannot be set to <b>random</b>.</li> </ul>
os	No	String	Node OS. For details about the supported OSs, see <a href="#">Node OS</a> . <b>NOTE</b> <ul style="list-style-type: none"> <li>The system automatically selects the supported OS based on the cluster version. If the current cluster version does not support the OS, an error will be reported.</li> <li>If <b>alpha.cce/NodeImageID</b> in <b>extendParam</b> is specified during node creation, you do not need to configure this parameter.</li> <li>This parameter is mandatory when creating a node pool.</li> <li>If you are using shared disk space when creating a node, this parameter is mandatory. Specifically, the disk initialization configuration management parameter <b>storage</b> is used and the <b>name</b> field of <b>virtualSpaces</b> in <b>StorageGroups</b> is set to <b>share</b>.</li> </ul>
login	Yes	<a href="#">Login</a> object	Node login mode. Either the key pair or password must be used for login. <b>NOTE</b> This parameter cannot be modified when a node pool is being updated.



Parameter	Mandatory	Type	Description
rootVolume	Yes	<a href="#">Volume</a> object	Information about disks on the node
dataVolumes	Yes	Array of <a href="#">Volume</a> objects	Data disk parameters of a node. A VM can have up to 16 disks attached, while a BMS can only have up to 10 disks attached. The maximum number of disks that can be attached to either a VM or BMS is limited by their respective maximum disk limits. (You can add more data disks to each of your nodes through the CCE console and APIs.) The data disk used by the container runtime and kubelet cannot be detached, or the node will become unavailable. For DeC nodes, the parameter details are the same as those for <b>rootVolume</b> .
storage	No	<a href="#">Storage</a> object	Disk initialization management parameter. This parameter is complex to configure. For details, see <a href="#">Attaching Disks to a Node</a> . <b>NOTE</b> <ul style="list-style-type: none"> <li>If a node has both local and EVS disks attached, make sure to not keep the default value for this parameter, or it may result in unexpected disk partitions.</li> </ul> <b>NOTE</b> <ul style="list-style-type: none"> <li>If you want to change the value range of a data disk to 20 to 32768, do not use the default value.</li> </ul> <b>NOTE</b> <ul style="list-style-type: none"> <li>If you want to use the shared disk space (with the runtime and Kubernetes partitions cancelled), do not retain the default value of this parameter. For details about the shared disk space, see <a href="#">Data Disk Space Allocation</a>.</li> </ul>

Parameter	Mandatory	Type	Description
publicIP	No	<b>NodePublicIP</b> object	EIP of a node. <b>NOTE</b> This parameter is not supported when you add a node to a node pool.
nodeNicSpec	No	<b>NodeNicSpec</b> object	NIC of the node
count	No	Integer	Number of nodes to be created in a batch. The value must be a positive integer greater than or equal to 1 and less than or equal to the defined limit. This parameter can be left blank when it is used for a node pool.
billingMode	No	Integer	Node billing mode. <ul style="list-style-type: none"> <li>• <b>0</b>: pay-per-use</li> </ul>

Parameter	Mandatory	Type	Description
taints	No	Array of <b>Taint</b> objects	<p>You can add taints to created nodes to configure anti-affinity. A maximum of 20 taints can be added. Each taint contains the following parameters:</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain can be prefixed to a key.</li> <li>• <b>Value:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed.</li> <li>• <b>Effect:</b> Available options are <b>NoSchedule</b>, <b>PreferNoSchedule</b>, and <b>NoExecute</b>. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</li> </ul> <p>Example:</p> <pre>"taints": [{   "key": "status",   "value": "unavailable",   "effect": "NoSchedule" }, {   "key": "looks",   "value": "bad",   "effect": "NoSchedule" }]</pre>

Parameter	Mandatory	Type	Description
k8sTags	No	Map<String,String>	<p>Defined in key-value pairs. A maximum of 20 key-value pairs are allowed.</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain containing a maximum of 253 characters (such as <b>example.com/my-key</b>) can be prefixed to a key.</li> <li>• <b>Value:</b> The value can be left blank or contain 1 to 63 characters that start and end with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</li> </ul> <p>Example:</p> <pre>"k8sTags": {   "key": "value" }</pre>
ecsGroupId	No	String	<p>ECS group ID. If this parameter is configured, nodes will be created in the specified ECS group.</p> <p><b>NOTE</b> This configuration does not take effect when a node pool is created. To ensure that all nodes in a node pool are in the same ECS group, configure <b>nodeManagement</b> in the node pool.</p>

Parameter	Mandatory	Type	Description
dedicatedHostId	No	String	<p>ID of the DeH to which the node is scheduled.</p> <p><b>NOTE</b> This parameter is not supported when you add a node during node pool creation.</p>
userTags	No	Array of <b>UserTag</b> objects	<p>Cloud server tag (resource tag). The key of a tag must be unique. The maximum number of custom tags supported by CCE varies depending on regions and cannot exceed 8. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</p> <p><b>NOTE</b> A tag key can have a maximum of 36 characters. It can only contain letters, digits, Unicode characters, and special characters (-).</p>
runtime	No	<b>Runtime</b> object	<p>Container runtime. Default value:</p> <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>

Parameter	Mandatory	Type	Description
initializedConditions	No	Array of strings	<p>Custom initialization flag, which is left blank by default. Before CCE nodes are initialized, <b>node.cloudprovider.kubernetes.io/uninitialized</b> is added to the nodes to prevent pods from being scheduled to these nodes. When creating a node, you can configure the <b>initializedConditions</b> parameter to control the taint removal time. By default, the timeout period is not configured.</p> <p>Example:</p> <ol style="list-style-type: none"> <li>1. Create a node and add parameter <code>"initializedConditions": ["CCEInitial", "CustomedInitial"]</code>.</li> <li>2. After custom initialization is complete, CCE will call a Kubernetes API (for example, <b>PATCH /v1/nodes/{node_ip}/status</b>) to update the node <b>conditions</b> by adding two labels of types <b>CCEInitial</b> and <b>CustomedInitial</b> and setting the status to <b>True</b>, as shown in the following:</li> </ol> <pre>status:   conditions:   - type: CCEInitial     status: 'True'   - type: CustomedInitial     status: 'True'</pre> <ol style="list-style-type: none"> <li>1. CCE polls <b>status.Conditions</b> of nodes to check whether there are conditions of types <b>CCEInitial</b> and <b>CustomedInitial</b>. If such conditions exist and the status is <b>True</b>, the node initialization is complete and the initialization taint is removed.</li> </ol>

Parameter	Mandatory	Type	Description
			<p>2. <b>initializedConditions</b> allows you to configure a timeout period during node creation. For example, "initializedConditions": ["CCEInitial:15m", "CustomedInitial:15m"] indicates that the timeout period is 15 minutes. After the timeout period is reached, the initialization conditions will be automatically ignored and the initialization taint will be removed when CCE polls the node.</p> <ul style="list-style-type: none"> <li>• The timeout interval in <b>initializedConditions</b> ranges from 1 to 99 seconds.</li> <li>• Use only letters and digits. Max. characters: 20.</li> <li>• Max. flags: 2.</li> <li>• The unit of the timeout period is minute (m).</li> </ul>
extendParam	No	<b>NodeExtendParam</b> object	Extended parameters for creating a node.
hostnameConfig	No	<b>HostnameConfig</b> object	Kubernetes node name configuration parameter, which is supported by clusters of v1.23.6-r0 to v1.25 or clusters of v1.25.2-r0 or later versions.

Parameter	Mandatory	Type	Description
serverEnterpriseProjectID	No	String	Enterprise project ID of a server. CCE does not support features provided by EPS. This field is used only for synchronizing enterprise project IDs of servers. To create a node or node pool, you can specify an existing enterprise project for it. If this field is left blank, it inherits the enterprise project attributes of the cluster that runs the node or node pool. When updating a node pool, any configuration modifications you made will apply only to new nodes. You need to modify the enterprise projects of existing nodes on the EPS console. If no value is specified during the update, this field will not be updated. If this parameter is left blank, the cluster enterprise project is returned.

**Table 4-184** Login

Parameter	Mandatory	Type	Description
sshKey	No	String	Name of the key pair used for login.
userPassword	No	UserPassword object	Password used for node login.

**Table 4-185** UserPassword

Parameter	Mandatory	Type	Description
username	No	String	Login account. The default value is <b>root</b> .



Parameter	Mandatory	Type	Description
password	Yes	String	<p>Login password. If a username and password are used to create a node, this field is shielded in the response body. A password must meet the following complexity requirements:</p> <ul style="list-style-type: none"> <li>• Contains 8 to 26 characters.</li> <li>• Contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters !@\$%^_-=+[{ }];,./?</li> <li>• Cannot contain the username or the username spelled backwards. The <b>password</b> field must be salted during node creation. For details, see <a href="#">Adding a Salt in the password Field When Creating a Node</a>.</li> </ul>

**Table 4-186** Volume

Parameter	Mandatory	Type	Description
size	Yes	Integer	<p>Disk size, in GiB.</p> <ul style="list-style-type: none"> <li>• System disk: 40 to 1024</li> <li>• Value range for data disks: 100 to 32768</li> </ul>
volumetype	Yes	String	<p>Disk type. For details about possible values, see the description of the <b>root_volume</b> parameter in the API used to create an ECS.</p> <ul style="list-style-type: none"> <li>• <b>SAS</b>: high I/O SAS disk</li> <li>• <b>SSD</b>: ultra-high I/O SSD disk</li> <li>• <b>SATA</b>: common I/O SATA disk SATA disks have been removed from EVS. You can find them attached only on existing nodes.</li> </ul>

Parameter	Mandatory	Type	Description
extendParam	No	Map<String,Object>	Extended disk parameters, defined in <b>extendparam</b> in the API used to create an ECS.
cluster_id	No	String	ID of the storage pool used by the ECS system disk. This field is used only for DeC clusters, which functions as <b>dssPoolID</b> , that is, the ID of the DSS storage pool.
cluster_type	No	String	Storage class of the cloud server system disk. The value is always <b>dss</b> . This field is used only for DeC clusters.
hw:passthrough	No	Boolean	<ul style="list-style-type: none"> <li>Pay attention to this field if your ECS is SDI-compliant. If the value of this field is <b>true</b>, the created disk is of SCSI type.</li> <li>If the node pool type is <b>ElasticBMS</b>, this field must be set to <b>true</b>.</li> <li>If a node specification involves both local and EVS disks, configure the disk initialization parameters. For details, see <a href="#">Attaching Disks to a Node</a>.</li> </ul>
metadata	No	<b>VolumeMetadata</b> object	EVS disk encryption information. This field is mandatory only when you need to encrypt the system disk or data disks of the node to be created.

**Table 4-187** VolumeMetadata

Parameter	Mandatory	Type	Description
<code>__system__encrypted</code>	No	String	Whether the EVS disk is encrypted. The value <b>0</b> indicates that the EVS disk is not encrypted, and the value <b>1</b> indicates that the EVS disk is encrypted.  If this parameter is not specified, EVS disks will not be encrypted by default.
<code>__system__cmkid</code>	No	String	CMK ID, which indicates encryption in <b>metadata</b> . This field is used with <b>__system__encrypted</b> .

**Table 4-188** Storage

Parameter	Mandatory	Type	Description
<code>storageSelectors</code>	Yes	Array of <b>StorageSelectors</b> objects	Disk selection. Matched disks are managed according to <b>matchLabels</b> and <b>storageType</b> .
<code>storageGroups</code>	Yes	Array of <b>StorageGroups</b> objects	A storage group consists of multiple storage devices. It is used to divide storage space.

**Table 4-189** StorageSelectors

Parameter	Mandatory	Type	Description
<code>name</code>	Yes	String	Selector name, used as the index of <b>selectorNames</b> in <b>storageGroup</b> . Therefore, the name of each selector must be unique.
<code>storageType</code>	Yes	String	Storage type. Only <b>evs</b> (EVS disks) and <b>local</b> (local disks) are supported. The local storage does not support disk selection. All local disks will form a VG, so only one storageSelector of the local type is allowed.

Parameter	Mandatory	Type	Description
matchLabels	No	<a href="#">matchLabels</a> object	Matching field of an EVS disk. The value can be <b>size</b> , <b>volumeType</b> , <b>metadataEncrypted</b> , <b>metadataCmkid</b> , or <b>count</b> in <b>DataVolume</b> .

**Table 4-190** matchLabels

Parameter	Mandatory	Type	Description
size	No	String	Matched disk size. If this parameter is left unspecified, the disk size is not limited. Example: 100
volumeType	No	String	EVS disk type. <b>SSD</b> , <b>GPSSD</b> , <b>SAS</b> , <b>ESSD</b> , and <b>SATA</b> are supported. If this parameter is left blank, there is no restriction on the disk type.
metadataEncrypted	No	String	Disk encryption identifier. If this parameter is set to <b>0</b> , it specifies that the disk is not encrypted. If this parameter is set to <b>1</b> , it specifies that the disk is encrypted. If this parameter is left blank, there is no restriction on the disk encryption identifier.
metadataCmkid	No	String	CMK ID of an encrypted disk. The value is a string of 36 bytes. If this parameter is left blank, there is no restriction on the disk key ID.
count	No	String	Number of disks to be selected. If this parameter is left blank, all disks of this type are selected.

**Table 4-191** StorageGroups

Parameter	Mandatory	Type	Description
name	Yes	String	Name of a virtual storage group, which must be unique. <b>NOTE</b> <ul style="list-style-type: none"> <li>If <b>cceManaged</b> is set to <b>true</b>, the name must be set to <b>vgpaas</b>.</li> <li>If the data disk is used as a temporary storage volume, the name must be <b>vg-everest-localvolume-ephemeral</b>.</li> <li>If the data disk is used as a persistent storage volume, the name must be <b>vg-everest-localvolume-persistent</b>.</li> </ul>
cceManaged	No	Boolean	Storage space for Kubernetes and runtime components. Only one group can be set to <b>true</b> . If this parameter is left blank, the default value <b>false</b> is used.
selectorNames	Yes	Array of strings	Corresponds to <b>name</b> in <b>storageSelectors</b> . One group can select multiple selectors, but one selector can be selected by only one group.
virtualSpaces	Yes	Array of <b>VirtualSpace</b> objects	Detailed management of space configuration in a group.

**Table 4-192** VirtualSpace

Parameter	Mandatory	Type	Description
name	Yes	String	<p>Name of a virtualSpace. Options: <b>share</b>, <b>kubernetes</b>, <b>runtime</b>, and <b>user</b></p> <ul style="list-style-type: none"> <li>• <b>share</b>: shared space configuration (with the runtime and Kubernetes partitions cancelled). <b>lvmConfig</b> needs to be configured.</li> <li>• <b>Kubernetes</b>: Kubernetes space configuration. <b>lvmConfig</b> needs to be configured.</li> <li>• <b>runtime</b>: runtime space configuration. <b>runtimeConfig</b> needs to be configured.</li> <li>• <b>user</b>: user space configuration. <b>lvmConfig</b> needs to be configured.</li> </ul>
size	Yes	String	<p>Size of a virtualSpace. The value must be an integer in percentage. Example: 90%. <b>NOTE</b> The value ranges from 10 to 90. The sum of percentages of all virtualSpaces in a group cannot exceed 100%.</p>
lvmConfig	No	<b>LVMConfig</b> object	<p>LVM configuration managements, applicable to the <i>share</i>*, <b>kubernetes</b>, and <b>user</b> spaces. Note that one virtual space supports only one config.</p>
runtimeConfig	No	<b>RuntimeConfig</b> object	<p>runtime configurations, applicable to the <b>runtime</b> space. Note that one virtual space supports only one config.</p>

**Table 4-193** LVMConfig

Parameter	Mandatory	Type	Description
lvType	Yes	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.
path	No	String	Disk mount path. This parameter takes effect only in user configuration. The value is an absolute path. Digits, letters, periods (.), hyphens (-), and underscores (_) are allowed.  <b>NOTE</b> Do not leave the path empty or set it to a key OS path such as the root directory. Key OS paths that cannot be specified: /, /home, /home/, /bin, /bin/, /lib, /lib/, /root, /root/, /boot, /boot/, /dev, /dev/, /etc, /etc/, /lost+found, /lost+found/, /mnt, /mnt/, /proc, /proc/, /sbin, /sbin/, /srv, /srv/, /tmp, /tmp/, /var, /var/, /media, /media/, /opt, /opt/, /selinux, /selinux/, /sys, /sys/, /usr, /usr/, /opt/cloud/, /mnt/paas/, /home/paas/, /var/paas/, /var/lib/, /var/script/

**Table 4-194** RuntimeConfig

Parameter	Mandatory	Type	Description
lvType	Yes	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.

**Table 4-195** NodePublicIP

Parameter	Mandatory	Type	Description
ids	No	Array of strings	IDs of existing EIPs. The quantity cannot be greater than the number of nodes to be created. <b>NOTE</b> If <b>ids</b> has been set, you do not need to set <b>count</b> and <b>eip</b> .
count	No	Integer	Number of EIPs to be dynamically created. <b>NOTE</b> <b>count</b> and <b>eip</b> must be set at the same time.
eip	No	<b>NodeEIPSpec</b> object	EIP configuration. <b>NOTE</b> This parameter is not supported when creating a node pool.

**Table 4-196** NodeEIPSpec

Parameter	Mandatory	Type	Description
iptype	Yes	String	EIP type, specified in <b>publicip.type</b> in the API for assigning an EIP.
bandwidth	No	<b>NodeBandwidth</b> object	Bandwidth parameters of the EIP



**Table 4-197** NodeBandwidth

Parameter	Mandatory	Type	Description
chargemode	No	String	<p>Bandwidth billing mode.</p> <ul style="list-style-type: none"> <li>If this field is not specified, the billing is based on bandwidth.</li> <li>If the field is null, the billing is based on bandwidth.</li> <li>If the field value is <b>traffic</b>, the billing is based on traffic.</li> <li>If the value is out of the preceding options, the cloud server will fail to be created.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>Billed by bandwidth: The billing will be based on the data transmission rate (in Mbps) of public networks. This billing mode is recommended if your bandwidth usage is higher than 10%.</li> <li>Billed by traffic: The billing is based on the total amount of data (in GB) transmitted over the public network. This mode is available only when you are creating a pay-per-use node. This billing mode is recommended if your bandwidth usage is lower than 10%.</li> </ul>
size	No	Integer	Bandwidth size, specified in <b>bandwidth.size</b> in the API for assigning an EIP.
sharetype	No	String	Bandwidth sharing type. Value options: <b>PER</b> (exclusive bandwidth)

**Table 4-198** NodeNicSpec

Parameter	Mandatory	Type	Description
primaryNic	No	<a href="#">NicSpec</a> object	Description of the primary NIC.

Parameter	Mandatory	Type	Description
extNics	No	Array of <b>NicSpec</b> objects	Extension NIC <b>NOTE</b> This parameter is not supported when you add a node to a node pool.

**Table 4-199** NicSpec

Parameter	Mandatory	Type	Description
subnetId	No	String	Network ID of the subnet to which a NIC belongs. If <b>subnetId</b> is not specified when a primary NIC is creating, the cluster subnet will be used. If <b>subnetList</b> is also configured for a node pool, the <b>subnetList</b> field is used for adding subnets to the node pool. When creating an extension NIC, you must specify <b>subnetId</b> .
fixedIps	No	Array of strings	The IP address of the primary ENI is specified using <b>fixedIps</b> . The number of IP addresses cannot be greater than the number of created nodes. Either <b>fixedIps</b> or <b>ipBlock</b> can be specified. <b>fixedIps</b> cannot be specified for extension ENIs.
ipBlock	No	String	CIDR format of the primary NIC IP range. The IP address of the created node falls in this range. <b>fixedIps</b> and <b>ipBlock</b> cannot be specified at the same time.
subnetList	No	Array of strings	Network ID list of the subnet where the NIC resides. A maximum of 20 subnets can be configured for a node pool.

**Table 4-200** Taint

Parameter	Mandatory	Type	Description
key	Yes	String	Key.
value	No	String	Value.
effect	Yes	String	Effect.

**Table 4-201** UserTag

Parameter	Mandatory	Type	Description
key	No	String	Key of the cloud server label. The value cannot start with <b>CCE-</b> or <b>__type_baremetal</b> .
value	No	String	Value of the cloud server label.

**Table 4-202** Runtime

Parameter	Mandatory	Type	Description
name	No	String	Container runtime. Default value: <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>

**Table 4-203** NodeExtendParam

Parameter	Mandatory	Type	Description
ecs:performancetype	No	String	ECS flavor types. This field is returned in the response.

Parameter	Mandatory	Type	Description
maxPods	No	Integer	<p>Maximum number of pods that can be created on a node, including the default system pods. Value range: 16 to 256.</p> <p>This limit prevents the node from being overloaded of pods.</p> <p>The number of pods that can be created on a node is determined by multiple parameters. For details, see <a href="#">Maximum Number of Pods That Can Be Created on a Node</a>.</p>
DockerLVMConfigOverride	No	String	<p>Docker data disk configuration item. (This parameter has been discarded. Use the <b>storage</b> field instead.)</p> <p>Example default configuration:  <code>"DockerLVMConfigOverride":"dockerThinpool=vgpaas/90%VG;kubernetesLV=vgpaas/10%VG;diskType=evs;lvType=linear"</code></p> <p>By default, if no VD disk is available, an error occurs because the data disk fails to be found. Set <b>diskType</b> based on the actual drive letter type. The following fields are included:</p> <ul style="list-style-type: none"> <li>• <b>userLV</b> (optional): size of the user space, for example, <b>vgpaas/20%VG</b>.</li> <li>• <b>userPath</b> (optional): mount path of the user space, for example, <b>/home/wqt-test</b>.</li> <li>• <b>diskType</b>: disk type. Currently, only <b>evs</b>, <b>hdd</b>, and <b>ssd</b> are supported.</li> <li>• <b>lvType</b>: type of a logic volume. The value can be <b>linear</b> or <b>striped</b>.</li> <li>• <b>dockerThinpool</b>: Docker space size, for example, <b>vgpaas/60%VG</b>.</li> <li>• <b>kubernetesLV</b>: kubelet space size, for example, <b>vgpaas/20%VG</b>.</li> </ul>

Parameter	Mandatory	Type	Description
dockerBaseSize	No	Integer	<p>Available disk space of a single container on a node, in GB. This parameter has been discarded. Use <b>containerBaseSize</b>.</p> <p>If this parameter is left blank or is set to <b>0</b>, the default value is used. In Device Mapper mode, the default value is <b>10</b>. In OverlayFS mode, the available space of a single container is not limited by default, and the <b>dockerBaseSize</b> setting takes effect only on nodes running EulerOS in clusters of the new versions.</p> <p>For details about how to allocate the space for the container runtime, see <a href="#">Data Disk Space Allocation</a>.</p> <p>When Device Mapper is used, it is recommended that <b>dockerBaseSize</b> be set to a value less than or equal to 80 GB. If the value is too large, the container runtime may fail to be started due to long initialization. If there are special requirements for the container disk space, you can mount an external or local storage device.</p>

Parameter	Mandatory	Type	Description
containerBaseSize	No	Integer	<p>Available disk space of a single container on a node, in GiB.</p> <p>If this parameter is left blank or is set to <b>0</b>, the default value is used. In OverlayFS mode, the available space of a single container is not limited by default. In Device Mapper mode, the default value is <b>10</b>, and the <b>containerBaseSize</b> setting takes effect only on nodes running EulerOS in clusters v1.23.14-r0, v1.25.9-r0, v1.27.6-r0, v1.28.4-r0 or later.</p> <p>For details about how to allocate the space for the container runtime, see <a href="#">Data Disk Space Allocation</a>.</p> <p>When Device Mapper is used, it is recommended that <b>containerBaseSize</b> be set to a value less than or equal to 80 GB. If the value is too large, the container runtime may fail to be started due to long initialization. If there are special requirements for the container disk space, you can mount an external or local storage device. In new versions, the Device Mapper mode is used only by BMSs in the same resource pool and is being discarded.</p> <p><b>NOTE</b> This parameter cannot be updated when a node pool is being updated.</p>
publicKey	No	String	Public key of a node.

Parameter	Mandatory	Type	Description
alpha.cce/ preInstall	No	String	Pre-installation scripts. The total number of characters after transcoding cannot exceed 10240. <b>NOTE</b> The input value must be encoded using Base64. The method is as follows: <pre>echo -n "Content to be encoded"   base64</pre>
alpha.cce/ postInstall	No	String	Post-installation scripts. The total number of characters after transcoding cannot exceed 10240. <b>NOTE</b> The input value must be encoded using Base64. The method is as follows: <pre>echo -n "Content to be encoded"   base64</pre>
alpha.cce/ NodeImageID	No	String	This parameter is required when a custom image is used to create a BMS node.
chargingMode	No	Integer	Billing mode of a node. This parameter has been deprecated. Use the <b>billingMode</b> parameter in <b>NodeSpec</b> .
agency_name	No	String	Name of an agency An agency is created by a tenant administrator on Identity and Access Management (IAM) to provide temporary credentials for CCE nodes to access cloud servers. This parameter is returned only when it is transferred during node creation.
kubeReservedMem	No	Integer	Reserved node memory, which is reserved for Kubernetes components.
systemReservedMem	No	Integer	Reserved node memory, which is reserved for system components.

Parameter	Mandatory	Type	Description
init-node-password	No	String	Node password. If this parameter is used as a response parameter, asterisks (*) will be displayed.

**Table 4-204** HostnameConfig

Parameter	Mandatory	Type	Description
type	Yes	String	<p>Configuration type of the Kubernetes node name. The default value is <b>privatelp</b>.</p> <ul style="list-style-type: none"> <li>• <b>privatelp</b>: The Kubernetes node is named after its IP address.</li> <li>• <b>cceNodeName</b>: The Kubernetes node is named after the CCE node.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• For a node which is configured using <b>cceNodeName</b>, the name is the same as the Kubernetes node name and the ECS name. The node name cannot be changed. If the ECS name is changed on the ECS console, the node name will retain unchanged after ECS synchronization.</li> <li>• For a node which is configured using <b>cceNodeName</b>, to avoid a conflict between Kubernetes nodes, the system automatically adds a suffix to each node name. The suffix is in the format of <b>A hyphen (-) Five random characters</b>. The value of the random characters is a lowercase letter or a digit ranging from 0 to 9.</li> </ul>

## Response Parameters

Status code: 201



**Table 4-205** Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at <b>Node</b> and cannot be changed.
apiVersion	String	API version. The value is fixed at <b>v3</b> and cannot be changed.
metadata	<b>NodeMetadata</b> object	Node metadata, which is a collection of attributes.
spec	<b>NodeSpec</b> object	Detailed description of the node. CCE creates or updates objects by defining or updating <b>spec</b> .
status	<b>NodeStatus</b> object	Node status, which is dynamically recorded. A user-defined value will not function when a node is being created or modified.

**Table 4-206** NodeMetadata

Parameter	Type	Description
name	String	Node name. <b>NOTE</b> Enter 1 to 56 characters starting with a lowercase letter and not ending with a hyphen (-). Only lowercase letters, digits, hyphens (-), and periods (.) are allowed. If <b>name</b> is not configured or is left blank, the node name is generated based on the default rule, which is <i>Cluster name-Random characters</i> . Only the first 36 characters are kept for a long cluster name. If the number of nodes ( <b>count</b> ) is greater than <b>1</b> , some random characters are added to the end of the entered node name according to the default rule, which is <i>User-defined name-Random characters</i> . If the length of the entered node name exceeds 50 characters, the system uses the first 50 characters and adds a random character string to the end of the name.
uid	String	Node ID, which is unique and automatically generated after the resource is created. A user-defined ID will not take effect.
labels	Map<String,String>	CCE node label (not the native Kubernetes label). Labels are used to select objects that meet certain criteria. A label is a key-value pair. Example: <pre>"labels": {   "key" : "value" }</pre>

Parameter	Type	Description
annotations	Map<String,String>	<p>CCE node annotations, in key-value pairs (not the native Kubernetes annotations). Example:</p> <pre>"annotations": {   "key1" : "value1",   "key2" : "value2" }</pre> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li><b>annotations</b> are not used to identify or select objects. The metadata in <b>annotations</b> may be small or large, structured or unstructured, and may include characters that are not allowed in labels.</li> <li>This parameter is used only for query and cannot be input through a request. Entered data of this parameter is invalid.</li> </ul>
creationTimesamp	String	Time when the object was created. The value is automatically generated after the object is created. A user-defined value will not take effect.
updateTimestamp	String	Time when the object was updated. The value is automatically generated after the object is created. A user-defined value will not take effect.
ownerReference	<a href="#">ownerReference</a> object	Owner object

**Table 4-207** ownerReference

Parameter	Type	Description
nodepoolName	String	Node pool name
nodepoolID	String	Node pool UID

**Table 4-208** NodeSpec

Parameter	Type	Description
flavor	String	Node specifications. For details about the node specifications supported by CCE, see <a href="#">Node Flavor Description</a> .

Parameter	Type	Description
az	String	AZ where the node to be created is located. You need to specify the AZ name. A node created using APIs does not support random AZs. <b>NOTE</b> <ul style="list-style-type: none"> <li>When you are creating a node pool and configuring a scaling group, this parameter cannot be set to <b>random</b>.</li> </ul>
os	String	Node OS. For details about the supported OSs, see <a href="#">Node OS</a> . <b>NOTE</b> <ul style="list-style-type: none"> <li>The system automatically selects the supported OS based on the cluster version. If the current cluster version does not support the OS, an error will be reported.</li> <li>If <b>alpha.cce/NodeImageID</b> in <b>extendParam</b> is specified during node creation, you do not need to configure this parameter.</li> <li>This parameter is mandatory when creating a node pool.</li> <li>If you are using shared disk space when creating a node, this parameter is mandatory. Specifically, the disk initialization configuration management parameter <b>storage</b> is used and the <b>name</b> field of <b>virtualSpaces</b> in <b>StorageGroups</b> is set to <b>share</b>.</li> </ul>
login	<a href="#">Login</a> object	Node login mode. Either the key pair or password must be used for login. <b>NOTE</b> This parameter cannot be modified when a node pool is being updated.
rootVolume	<a href="#">Volume</a> object	Information about disks on the node
dataVolumes	Array of <a href="#">Volume</a> objects	Data disk parameters of a node. A VM can have up to 16 disks attached, while a BMS can only have up to 10 disks attached. The maximum number of disks that can be attached to either a VM or BMS is limited by their respective maximum disk limits. (You can add more data disks to each of your nodes through the CCE console and APIs.) The data disk used by the container runtime and kubelet cannot be detached, or the node will become unavailable. For DeC nodes, the parameter details are the same as those for <b>rootVolume</b> .

Parameter	Type	Description
storage	<a href="#">Storage</a> object	<p>Disk initialization management parameter. This parameter is complex to configure. For details, see <a href="#">Attaching Disks to a Node</a>.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If a node has both local and EVS disks attached, make sure to not keep the default value for this parameter, or it may result in unexpected disk partitions.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If you want to change the value range of a data disk to 20 to 32768, do not use the default value.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If you want to use the shared disk space (with the runtime and Kubernetes partitions cancelled), do not retain the default value of this parameter. For details about the shared disk space, see <a href="#">Data Disk Space Allocation</a>.</li> </ul>
publicIP	<a href="#">NodePublicIP</a> object	<p>EIP of a node.</p> <p><b>NOTE</b></p> <p>This parameter is not supported when you add a node to a node pool.</p>
nodeNicSpec	<a href="#">NodeNicSpec</a> object	NIC of the node
count	Integer	Number of nodes to be created in a batch. The value must be a positive integer greater than or equal to 1 and less than or equal to the defined limit. This parameter can be left blank when it is used for a node pool.
billingMode	Integer	<p>Node billing mode.</p> <ul style="list-style-type: none"> <li><b>0</b>: pay-per-use</li> </ul>

Parameter	Type	Description
taints	Array of <b>Taint</b> objects	<p>You can add taints to created nodes to configure anti-affinity. A maximum of 20 taints can be added. Each taint contains the following parameters:</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain can be prefixed to a key.</li> <li>• <b>Value:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed.</li> <li>• <b>Effect:</b> Available options are <b>NoSchedule</b>, <b>PreferNoSchedule</b>, and <b>NoExecute</b>. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</li> </ul> <p>Example:</p> <pre> taints: [{   "key": "status",   "value": "unavailable",   "effect": "NoSchedule" }, {   "key": "looks",   "value": "bad",   "effect": "NoSchedule" }] </pre>

Parameter	Type	Description
k8sTags	Map<String,String>	<p>Defined in key-value pairs. A maximum of 20 key-value pairs are allowed.</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain containing a maximum of 253 characters (such as <b>example.com/my-key</b>) can be prefixed to a key.</li> <li>• <b>Value:</b> The value can be left blank or contain 1 to 63 characters that start and end with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</li> </ul> <p>Example:</p> <pre>"k8sTags": {   "key": "value" }</pre>
ecsGroupId	String	<p>ECS group ID. If this parameter is configured, nodes will be created in the specified ECS group.</p> <p><b>NOTE</b> This configuration does not take effect when a node pool is created. To ensure that all nodes in a node pool are in the same ECS group, configure <b>nodeManagement</b> in the node pool.</p>
dedicatedHostId	String	<p>ID of the DeH to which the node is scheduled.</p> <p><b>NOTE</b> This parameter is not supported when you add a node during node pool creation.</p>

Parameter	Type	Description
userTags	Array of <a href="#">UserTag</a> objects	<p>Cloud server tag (resource tag). The key of a tag must be unique. The maximum number of custom tags supported by CCE varies depending on regions and cannot exceed 8. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</p> <p><b>NOTE</b> A tag key can have a maximum of 36 characters. It can only contain letters, digits, Unicode characters, and special characters (-_).</p>
runtime	<a href="#">Runtime</a> object	<p>Container runtime. Default value:</p> <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>

Parameter	Type	Description
initializedConditions	Array of strings	<p>Custom initialization flag, which is left blank by default.</p> <p>Before CCE nodes are initialized, <b>node.cloudprovider.kubernetes.io/uninitialized</b> is added to the nodes to prevent pods from being scheduled to these nodes. When creating a node, you can configure the <b>initializedConditions</b> parameter to control the taint removal time. By default, the timeout period is not configured.</p> <p>Example:</p> <ol style="list-style-type: none"> <li>1. Create a node and add parameter "initializedConditions": ["CCEInitial", "CustomedInitial"].</li> <li>2. After custom initialization is complete, CCE will call a Kubernetes API (for example, <b>PATCH /v1/nodes/{node_ip}/status</b>) to update the node <b>conditions</b> by adding two labels of types <b>CCEInitial</b> and <b>CustomedInitial</b> and setting the status to <b>True</b>, as shown in the following:</li> </ol> <pre data-bbox="815 1055 1428 1205"> status:   conditions:   - type: CCEInitial     status: 'True'   - type: CustomedInitial     status: 'True'                     </pre> <ol style="list-style-type: none"> <li>1. CCE polls <b>status.Conditions</b> of nodes to check whether there are conditions of types <b>CCEInitial</b> and <b>CustomedInitial</b>. If such conditions exist and the status is <b>True</b>, the node initialization is complete and the initialization taint is removed.</li> <li>2. <b>initializedConditions</b> allows you to configure a timeout period during node creation. For example, "initializedConditions": ["CCEInitial:15m", "CustomedInitial:15m"] indicates that the timeout period is 15 minutes. After the timeout period is reached, the initialization conditions will be automatically ignored and the initialization taint will be removed when CCE polls the node. <ul style="list-style-type: none"> <li>● The timeout interval in <b>initializedConditions</b> ranges from 1 to 99 seconds.</li> <li>● Use only letters and digits. Max. characters: 20.</li> <li>● Max. flags: 2.</li> </ul> </li> </ol>



Parameter	Type	Description
		<ul style="list-style-type: none"> <li>The unit of the timeout period is minute (m).</li> </ul>
extendParam	<a href="#">NodeExtendParam</a> object	Extended parameters for creating a node.
hostnameConfig	<a href="#">HostnameConfig</a> object	Kubernetes node name configuration parameter, which is supported by clusters of v1.23.6-r0 to v1.25 or clusters of v1.25.2-r0 or later versions.
serverEnterpriseProjectID	String	Enterprise project ID of a server. CCE does not support features provided by EPS. This field is used only for synchronizing enterprise project IDs of servers. To create a node or node pool, you can specify an existing enterprise project for it. If this field is left blank, it inherits the enterprise project attributes of the cluster that runs the node or node pool. When updating a node pool, any configuration modifications you made will apply only to new nodes. You need to modify the enterprise projects of existing nodes on the EPS console. If no value is specified during the update, this field will not be updated. If this parameter is left blank, the cluster enterprise project is returned.

**Table 4-209** Login

Parameter	Type	Description
sshKey	String	Name of the key pair used for login.
userPassword	<a href="#">UserPassword</a> object	Password used for node login.

**Table 4-210** UserPassword

Parameter	Type	Description
username	String	Login account. The default value is <b>root</b> .

Parameter	Type	Description
password	String	<p>Login password. If a username and password are used to create a node, this field is shielded in the response body. A password must meet the following complexity requirements:</p> <ul style="list-style-type: none"> <li>• Contains 8 to 26 characters.</li> <li>• Contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters !@\$%^-_=+[{ }];,./?</li> <li>• Cannot contain the username or the username spelled backwards. The <b>password</b> field must be salted during node creation. For details, see <a href="#">Adding a Salt in the password Field When Creating a Node</a>.</li> </ul>

Table 4-211 Volume

Parameter	Type	Description
size	Integer	<p>Disk size, in GiB.</p> <ul style="list-style-type: none"> <li>• System disk: 40 to 1024</li> <li>• Value range for data disks: 100 to 32768</li> </ul>
volumetype	String	<p>Disk type. For details about possible values, see the description of the <b>root_volume</b> parameter in the API used to create an ECS.</p> <ul style="list-style-type: none"> <li>• <b>SAS</b>: high I/O SAS disk</li> <li>• <b>SSD</b>: ultra-high I/O SSD disk</li> <li>• <b>SATA</b>: common I/O SATA disk SATA disks have been removed from EVS. You can find them attached only on existing nodes.</li> </ul>
extendParam	Map<String, Object>	<p>Extended disk parameters, defined in <b>extendparam</b> in the API used to create an ECS.</p>
cluster_id	String	<p>ID of the storage pool used by the ECS system disk. This field is used only for DeC clusters, which functions as <b>dssPoolID</b>, that is, the ID of the DSS storage pool.</p>
cluster_type	String	<p>Storage class of the cloud server system disk. The value is always <b>dss</b>. This field is used only for DeC clusters.</p>

Parameter	Type	Description
hw:passthrough	Boolean	<ul style="list-style-type: none"> <li>Pay attention to this field if your ECS is SDI-compliant. If the value of this field is <b>true</b>, the created disk is of SCSI type.</li> <li>If the node pool type is <b>ElasticBMS</b>, this field must be set to <b>true</b>.</li> <li>If a node specification involves both local and EVS disks, configure the disk initialization parameters. For details, see <a href="#">Attaching Disks to a Node</a>.</li> </ul>
metadata	<b>VolumeMetadata</b> object	EVS disk encryption information. This field is mandatory only when you need to encrypt the system disk or data disks of the node to be created.

**Table 4-212** VolumeMetadata

Parameter	Type	Description
__system__encrypted	String	Whether the EVS disk is encrypted. The value <b>0</b> indicates that the EVS disk is not encrypted, and the value <b>1</b> indicates that the EVS disk is encrypted.  If this parameter is not specified, EVS disks will not be encrypted by default.
__system__cmkid	String	CMK ID, which indicates encryption in <b>metadata</b> . This field is used with <b>__system__encrypted</b> .

**Table 4-213** Storage

Parameter	Type	Description
storageSelectors	Array of <b>StorageSelectors</b> objects	Disk selection. Matched disks are managed according to <b>matchLabels</b> and <b>storageType</b> .
storageGroups	Array of <b>StorageGroups</b> objects	A storage group consists of multiple storage devices. It is used to divide storage space.

**Table 4-214** StorageSelectors

Parameter	Type	Description
name	String	Selector name, used as the index of <b>selectorNames</b> in <b>storageGroup</b> . Therefore, the name of each selector must be unique.
storageType	String	Storage type. Only <b>evs</b> (EVS disks) and <b>local</b> (local disks) are supported. The local storage does not support disk selection. All local disks will form a VG, so only one storageSelector of the local type is allowed.
matchLabels	<a href="#">matchLabels</a> object	Matching field of an EVS disk. The value can be <b>size</b> , <b>volumeType</b> , <b>metadataEncrypted</b> , <b>metadataCmkid</b> , or <b>count</b> in <b>DataVolume</b> .

**Table 4-215** matchLabels

Parameter	Type	Description
size	String	Matched disk size. If this parameter is left unspecified, the disk size is not limited. Example: 100
volumeType	String	EVS disk type. <b>SSD</b> , <b>GPSSD</b> , <b>SAS</b> , <b>ESSD</b> , and <b>SATA</b> are supported. If this parameter is left blank, there is no restriction on the disk type.
metadataEncrypted	String	Disk encryption identifier. If this parameter is set to <b>0</b> , it specifies that the disk is not encrypted. If this parameter is set to <b>1</b> , it specifies that the disk is encrypted. If this parameter is left blank, there is no restriction on the disk encryption identifier.
metadataCmkid	String	CMK ID of an encrypted disk. The value is a string of 36 bytes. If this parameter is left blank, there is no restriction on the disk key ID.
count	String	Number of disks to be selected. If this parameter is left blank, all disks of this type are selected.

**Table 4-216** StorageGroups

Parameter	Type	Description
name	String	Name of a virtual storage group, which must be unique. <b>NOTE</b> <ul style="list-style-type: none"> <li>If <b>cceManaged</b> is set to <b>true</b>, the name must be set to <b>vgpaas</b>.</li> <li>If the data disk is used as a temporary storage volume, the name must be <b>vg-everest-localvolume-ephemeral</b>.</li> <li>If the data disk is used as a persistent storage volume, the name must be <b>vg-everest-localvolume-persistent</b>.</li> </ul>
cceManaged	Boolean	Storage space for Kubernetes and runtime components. Only one group can be set to <b>true</b> . If this parameter is left blank, the default value <b>false</b> is used.
selectorNames	Array of strings	Corresponds to <b>name</b> in <b>storageSelectors</b> . One group can select multiple selectors, but one selector can be selected by only one group.
virtualSpaces	Array of <a href="#">VirtualSpace</a> objects	Detailed management of space configuration in a group.

**Table 4-217** VirtualSpace

Parameter	Type	Description
name	String	Name of a virtualSpace. Options: <b>share</b> , <b>kubernetes</b> , <b>runtime</b> , and <b>user</b> <ul style="list-style-type: none"> <li><b>share</b>: shared space configuration (with the runtime and Kubernetes partitions cancelled). <b>lvmConfig</b> needs to be configured.</li> <li><b>Kubernetes</b>: Kubernetes space configuration. <b>lvmConfig</b> needs to be configured.</li> <li><b>runtime</b>: runtime space configuration. <b>runtimeConfig</b> needs to be configured.</li> <li><b>user</b>: user space configuration. <b>lvmConfig</b> needs to be configured.</li> </ul>

Parameter	Type	Description
size	String	Size of a virtualSpace. The value must be an integer in percentage. Example: 90%. <b>NOTE</b> The value ranges from 10 to 90. The sum of percentages of all virtualSpaces in a group cannot exceed 100%.
lvmConfig	<b>LVMConfig</b> object	LVM configuration managements, applicable to the <i>share*</i> , <b>kubernetes</b> , and <b>user</b> spaces. Note that one virtual space supports only one config.
runtimeConfig	<b>RuntimeConfig</b> object	runtime configurations, applicable to the <b>runtime</b> space. Note that one virtual space supports only one config.

**Table 4-218** LVMConfig

Parameter	Type	Description
lvType	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.
path	String	Disk mount path. This parameter takes effect only in user configuration. The value is an absolute path. Digits, letters, periods (.), hyphens (-), and underscores (_) are allowed. <b>NOTE</b> Do not leave the path empty or set it to a key OS path such as the root directory. Key OS paths that cannot be specified: /, /home, /home/, /bin, /bin/, /lib, /lib/, /root, /root/, /boot, /boot/, /dev, /dev/, /etc, /etc/, /lost+found, /lost+found/, /mnt, /mnt/, /proc, /proc/, /sbin, /sbin/, /srv, /srv/, /tmp, /tmp/, /var, /var/, /media, /media/, /opt, /opt/, /selinux, /selinux/, /sys, /sys/, /usr, /usr/, /opt/cloud/, /mnt/paas/, /home/paas/, /var/paas/, /var/lib/, /var/script/

**Table 4-219** RuntimeConfig

Parameter	Type	Description
lvType	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.

**Table 4-220** NodePublicIP

Parameter	Type	Description
ids	Array of strings	IDs of existing EIPs. The quantity cannot be greater than the number of nodes to be created. <b>NOTE</b> If <b>ids</b> has been set, you do not need to set <b>count</b> and <b>eip</b> .
count	Integer	Number of EIPs to be dynamically created. <b>NOTE</b> <b>count</b> and <b>eip</b> must be set at the same time.
eip	<b>NodeEIPSpec</b> object	EIP configuration. <b>NOTE</b> This parameter is not supported when creating a node pool.

**Table 4-221** NodeEIPSpec

Parameter	Type	Description
iptype	String	EIP type, specified in <b>publicip.type</b> in the API for assigning an EIP.
bandwidth	<b>NodeBandwidth</b> object	Bandwidth parameters of the EIP

**Table 4-222** NodeBandwidth

Parameter	Type	Description
chargemode	String	<p>Bandwidth billing mode.</p> <ul style="list-style-type: none"> <li>• If this field is not specified, the billing is based on bandwidth.</li> <li>• If the field is null, the billing is based on bandwidth.</li> <li>• If the field value is <b>traffic</b>, the billing is based on traffic.</li> <li>• If the value is out of the preceding options, the cloud server will fail to be created.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• Billed by bandwidth: The billing will be based on the data transmission rate (in Mbps) of public networks. This billing mode is recommended if your bandwidth usage is higher than 10%.</li> <li>• Billed by traffic: The billing is based on the total amount of data (in GB) transmitted over the public network. This mode is available only when you are creating a pay-per-use node. This billing mode is recommended if your bandwidth usage is lower than 10%.</li> </ul>
size	Integer	Bandwidth size, specified in <b>bandwidth.size</b> in the API for assigning an EIP.
sharetype	String	Bandwidth sharing type. Value options: <b>PER</b> (exclusive bandwidth)

**Table 4-223** NodeNicSpec

Parameter	Type	Description
primaryNic	<b>NicSpec</b> object	Description of the primary NIC.
extNics	Array of <b>NicSpec</b> objects	<p>Extension NIC</p> <p><b>NOTE</b> This parameter is not supported when you add a node to a node pool.</p>



**Table 4-224 NicSpec**

Parameter	Type	Description
subnetId	String	Network ID of the subnet to which a NIC belongs. If <b>subnetId</b> is not specified when a primary NIC is creating, the cluster subnet will be used. If <b>subnetList</b> is also configured for a node pool, the <b>subnetList</b> field is used for adding subnets to the node pool. When creating an extension NIC, you must specify <b>subnetId</b> .
fixedIps	Array of strings	The IP address of the primary ENI is specified using <b>fixedIps</b> . The number of IP addresses cannot be greater than the number of created nodes. Either <b>fixedIps</b> or <b>ipBlock</b> can be specified. <b>fixedIps</b> cannot be specified for extension ENIs.
ipBlock	String	CIDR format of the primary NIC IP range. The IP address of the created node falls in this range. <b>fixedIps</b> and <b>ipBlock</b> cannot be specified at the same time.
subnetList	Array of strings	Network ID list of the subnet where the NIC resides. A maximum of 20 subnets can be configured for a node pool.

**Table 4-225 Taint**

Parameter	Type	Description
key	String	Key.
value	String	Value.
effect	String	Effect.

**Table 4-226 UserTag**

Parameter	Type	Description
key	String	Key of the cloud server label. The value cannot start with <b>CCE-</b> or <b>__type_baremetal</b> .
value	String	Value of the cloud server label.

**Table 4-227** Runtime

Parameter	Type	Description
name	String	<p>Container runtime. Default value:</p> <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>

**Table 4-228** NodeExtendParam

Parameter	Type	Description
ecs:performancetype	String	ECS flavor types. This field is returned in the response.
maxPods	Integer	<p>Maximum number of pods that can be created on a node, including the default system pods. Value range: 16 to 256.</p> <p>This limit prevents the node from being overloaded of pods.</p> <p>The number of pods that can be created on a node is determined by multiple parameters. For details, see <a href="#">Maximum Number of Pods That Can Be Created on a Node</a>.</p>

Parameter	Type	Description
DockerLVMConfigOverride	String	<p>Docker data disk configuration item. (This parameter has been discarded. Use the <b>storage</b> field instead.) Example default configuration: "DockerLVMConfigOverride":"dockerThinpool=vgpaas/90%VG;kubernetesLV=vgpaas/10%VG;diskType=evs;lvType=linear"</p> <p>By default, if no VD disk is available, an error occurs because the data disk fails to be found. Set <b>diskType</b> based on the actual drive letter type. The following fields are included:</p> <ul style="list-style-type: none"> <li>• <b>userLV</b> (optional): size of the user space, for example, <b>vgpaas/20%VG</b>.</li> <li>• <b>userPath</b> (optional): mount path of the user space, for example, <b>/home/wqt-test</b>.</li> <li>• <b>diskType</b>: disk type. Currently, only <b>evs</b>, <b>hdd</b>, and <b>ssd</b> are supported.</li> <li>• <b>lvType</b>: type of a logic volume. The value can be <b>linear</b> or <b>striped</b>.</li> <li>• <b>dockerThinpool</b>: Docker space size, for example, <b>vgpaas/60%VG</b>.</li> <li>• <b>kubernetesLV</b>: kubelet space size, for example, <b>vgpaas/20%VG</b>.</li> </ul>
dockerBaseSize	Integer	<p>Available disk space of a single container on a node, in GB. This parameter has been discarded. Use <b>containerBaseSize</b>.</p> <p>If this parameter is left blank or is set to <b>0</b>, the default value is used. In Device Mapper mode, the default value is <b>10</b>. In OverlayFS mode, the available space of a single container is not limited by default, and the <b>dockerBaseSize</b> setting takes effect only on nodes running EulerOS in clusters of the new versions.</p> <p>For details about how to allocate the space for the container runtime, see <a href="#">Data Disk Space Allocation</a>.</p> <p>When Device Mapper is used, it is recommended that <b>dockerBaseSize</b> be set to a value less than or equal to 80 GB. If the value is too large, the container runtime may fail to be started due to long initialization. If there are special requirements for the container disk space, you can mount an external or local storage device.</p>

Parameter	Type	Description
containerBaseSize	Integer	<p>Available disk space of a single container on a node, in GiB.</p> <p>If this parameter is left blank or is set to <b>0</b>, the default value is used. In OverlayFS mode, the available space of a single container is not limited by default. In Device Mapper mode, the default value is <b>10</b>, and the <b>containerBaseSize</b> setting takes effect only on nodes running EulerOS in clusters v1.23.14-r0, v1.25.9-r0, v1.27.6-r0, v1.28.4-r0 or later.</p> <p>For details about how to allocate the space for the container runtime, see <a href="#">Data Disk Space Allocation</a>.</p> <p>When Device Mapper is used, it is recommended that <b>containerBaseSize</b> be set to a value less than or equal to 80 GB. If the value is too large, the container runtime may fail to be started due to long initialization. If there are special requirements for the container disk space, you can mount an external or local storage device. In new versions, the Device Mapper mode is used only by BMSs in the same resource pool and is being discarded.</p> <p><b>NOTE</b> This parameter cannot be updated when a node pool is being updated.</p>
publicKey	String	Public key of a node.
alpha.cce/preInstall	String	<p>Pre-installation scripts. The total number of characters after transcoding cannot exceed 10240.</p> <p><b>NOTE</b> The input value must be encoded using Base64. The method is as follows:</p> <pre>echo -n "Content to be encoded"   base64</pre>
alpha.cce/postInstall	String	<p>Post-installation scripts. The total number of characters after transcoding cannot exceed 10240.</p> <p><b>NOTE</b> The input value must be encoded using Base64. The method is as follows:</p> <pre>echo -n "Content to be encoded"   base64</pre>
alpha.cce/NodeImageID	String	This parameter is required when a custom image is used to create a BMS node.

Parameter	Type	Description
chargingMode	Integer	Billing mode of a node. This parameter has been deprecated. Use the <b>billingMode</b> parameter in <b>NodeSpec</b> .
agency_name	String	Name of an agency An agency is created by a tenant administrator on Identity and Access Management (IAM) to provide temporary credentials for CCE nodes to access cloud servers. This parameter is returned only when it is transferred during node creation.
kubeReservedMem	Integer	Reserved node memory, which is reserved for Kubernetes components.
systemReservedMem	Integer	Reserved node memory, which is reserved for system components.
init-node-password	String	Node password. If this parameter is used as a response parameter, asterisks (*) will be displayed.

**Table 4-229** HostnameConfig

Parameter	Type	Description
type	String	Configuration type of the Kubernetes node name. The default value is <b>privatelp</b> . <ul style="list-style-type: none"> <li>• <b>privatelp</b>: The Kubernetes node is named after its IP address.</li> <li>• <b>cceNodeName</b>: The Kubernetes node is named after the CCE node.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• For a node which is configured using <b>cceNodeName</b>, the name is the same as the Kubernetes node name and the ECS name. The node name cannot be changed. If the ECS name is changed on the ECS console, the node name will retain unchanged after ECS synchronization.</li> <li>• For a node which is configured using <b>cceNodeName</b>, to avoid a conflict between Kubernetes nodes, the system automatically adds a suffix to each node name. The suffix is in the format of <b>A hyphen (-) Five random characters</b>. The value of the random characters is a lowercase letter or a digit ranging from 0 to 9.</li> </ul>

**Table 4-230** NodeStatus

Parameter	Type	Description
phase	String	Node status, which is the status of a node during its lifecycle (such as installation and uninstallation) or the status of a node in a Kubernetes cluster. Options: <ul style="list-style-type: none"> <li>● <b>Build</b>: The node is being created.</li> <li>● <b>Installing</b>: The node is being managed.</li> <li>● <b>Upgrading</b>: The node is being upgraded.</li> <li>● <b>Active</b>: The node is running properly.</li> <li>● <b>Abnormal</b>: The node is abnormal.</li> <li>● <b>Deleting</b>: The node is being deleted.</li> <li>● <b>Error</b>: The node is faulty.</li> </ul>
lastProbeTime	String	Last time when the node status was checked. If the cluster is in the abnormal, frozen, or intermediate state (for example, creating), the node status check may be affected. The node status that takes more than five minutes to check has no reference value.
jobID	String	ID of a creation or deletion job
serverId	String	ID of the underlying ECS or BMS node
privateIP	String	IP address in the private network segment of the primary NIC on the node
privateIPv6IP	String	IPv6 address in the private network segment of the primary NIC on the node
publicIP	String	Node EIP. If the ECS data is not synchronized in real time, you can click <b>Sync Node Data</b> on the console to manually update the data.
deleteStatus	<a href="#">DeleteStatus</a> object	Resource status during resource deletion.

**Table 4-231** DeleteStatus

Parameter	Type	Description
previous_total	Integer	Total number of existing cluster resource records when the cluster is deleted.
current_total	Integer	Latest number of resource records, which is generated based on the current cluster resource records.

Parameter	Type	Description
updated	Integer	Total number of resource records updated when the cluster is deleted.
added	Integer	Total number of resource records updated when the cluster is deleted.
deleted	Integer	Total number of resource records deleted when the cluster is deleted.

## Example Requests

- Create a node with the flavor of 2 vCPUs and 4 GiB of memory. The node runs EulerOS 2.5 for Docker containers. The sizes of the system disk and data disk of the node are 50 GB and 100 GB, respectively. Both the disks are of the high I/O type.

POST /api/v3/projects/{project\_id}/clusters/{cluster\_id}/nodes

```
{
  "kind": "Node",
  "apiVersion": "v3",
  "metadata": {
    "name": "test-83790"
  },
  "spec": {
    "flavor": "c7.large.2",
    "az": "*****",
    "os": "EulerOS 2.5",
    "dataVolumes": [ {
      "size": 100,
      "volumetype": "SAS"
    } ],
    "billingMode": 0,
    "extendParam": {
      "maxPods": 110
    },
    "nodeNicSpec": {
      "primaryNic": {
        "subnetId": "ca964acf-8468-4735-8229-97940ef6c881"
      }
    },
    "rootVolume": {
      "size": 50,
      "volumetype": "SAS"
    },
    "runtime": {
      "name": "docker"
    },
    "login": {
      "sshKey": "KeyPair-001"
    },
    "storage": {
      "storageSelectors": [ {
        "name": "cceUse",
        "storageType": "evs",
        "matchLabels": {
          "size": "100",
          "volumeType": "SAS",
          "count": "1"
        }
      } ],
      "storageGroups": [ {
```

```

    "name" : "vgpaas",
    "selectorNames" : [ "cceUse" ],
    "cceManaged" : true,
    "virtualSpaces" : [ {
      "name" : "runtime",
      "size" : "90%"
    }, {
      "name" : "kubernetes",
      "size" : "10%"
    } ]
  } ]
}
"count" : 1
}
}

```

- Create a node with the flavor of 2 vCPUs and 4 GiB of memory. The node uses Docker. The sizes of the system disk and data disk of the node are 50 GB and 100 GB, respectively. The data disk uses the shared disk space mode.

POST /api/v3/projects/{project\_id}/clusters/{cluster\_id}/nodes

```

{
  "kind" : "Node",
  "apiVersion" : "v3",
  "metadata" : {
    "name" : "test-66909"
  },
  "spec" : {
    "flavor" : "c7.large.2",
    "az" : "*****",
    "os" : ,
    "dataVolumes" : [ {
      "size" : 100,
      "volumetype" : "SAS"
    } ],
    "billingMode" : 0,
    "extendParam" : {
      "maxPods" : 110
    },
    "nodeNicSpec" : {
      "primaryNic" : {
        "subnetId" : "ca964acf-8468-4735-8229-97940ef6c881"
      }
    },
    "rootVolume" : {
      "size" : 50,
      "volumetype" : "SAS"
    },
    "runtime" : {
      "name" : "docker"
    },
    "login" : {
      "sshKey" : "KeyPair-001"
    },
    "storage" : {
      "storageSelectors" : [ {
        "name" : "cceUse",
        "storageType" : "evs",
        "matchLabels" : {
          "size" : "100",
          "volumeType" : "SAS",
          "count" : "1"
        }
      } ],
      "storageGroups" : [ {
        "name" : "vgpaas",
        "selectorNames" : [ "cceUse" ],
        "cceManaged" : true,
        "virtualSpaces" : [ {

```



```
    "name" : "share",
    "size" : "100%"
  }
}
},
"count" : 1
}
}
```

## Example Responses

### Status code: 201

The job for creating a node in a specified cluster is successfully delivered.

```
{
  "kind" : "Node",
  "apiVersion" : "v3",
  "metadata" : {
    "name" : "test-83790",
    "uid" : "5ecfddfe-87db-11ec-b5e5-0255ac101514",
    "annotations" : {
      "jobid" : "5ec1518c-87db-11ec-b5e5-0255ac101514",
      "resourceJobId" : "5ed0d692-87db-11ec-b5e5-0255ac101514"
    }
  },
  "spec" : {
    "flavor" : "c7.large.2",
    "az" : "*****",
    "os" : "EulerOS 2.5",
    "login" : {
      "sshKey" : "KeyPair-001"
    },
    "rootVolume" : {
      "volumetype" : "SAS",
      "size" : 50
    },
    "dataVolumes" : [ {
      "volumetype" : "SAS",
      "size" : 100
    } ],
    "storage" : {
      "storageSelectors" : [ {
        "name" : "cceUse",
        "storageType" : "evs",
        "matchLabels" : {
          "count" : "1",
          "size" : "100",
          "volumeType" : "SAS"
        }
      } ],
      "storageGroups" : [ {
        "name" : "vgpaas",
        "cceManaged" : true,
        "selectorNames" : [ "cceUse" ],
        "virtualSpaces" : [ {
          "name" : "runtime",
          "size" : "90%"
        }, {
          "name" : "kubernetes",
          "size" : "10%"
        } ]
      } ]
    },
    "publicIP" : {
      "eip" : {
        "bandwidth" : { }
      }
    }
  },
}
```

```
"nodeNicSpec" : {  
  "primaryNic" : {  
    "subnetId" : "ca964acf-8468-4735-8229-97940ef6c881"  
  }  
},  
"count" : 1,  
"billingMode" : 0,  
"runtime" : {  
  "name" : "docker"  
},  
"extendParam" : {  
  "chargingMode" : 0,  
  "ecs:performancetype" : "computingv3",  
  "init-node-password" : "*****",  
  "maxPods" : 110,  
  "publicKey" : ""  
}  
},  
"status" : {  
  "jobID" : "5ec1518c-87db-11ec-b5e5-0255ac101514"  
}  
}
```

## Status Codes

Status Code	Description
201	The job for creating a node in a specified cluster is successfully delivered.

## Error Codes

See [Error Codes](#).

### 4.3.2 Reading a Specified Node

#### Function

This API is used to obtain details about a specified node via the node ID.

#### NOTE

The URL for cluster management is in the format of `https://Endpoint/uri`. In the URL, **uri** indicates the resource path, that is, the path for API access.

#### URI

GET `/api/v3/projects/{project_id}/clusters/{cluster_id}/nodes/{node_id}`

**Table 4-232** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
node_id	Yes	String	Node ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

## Request Parameters

**Table 4-233** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

## Response Parameters

Status code: 200

**Table 4-234** Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at <b>Node</b> and cannot be changed.
apiVersion	String	API version. The value is fixed at <b>v3</b> and cannot be changed.

Parameter	Type	Description
metadata	<a href="#">NodeMetadata</a> object	Node metadata, which is a collection of attributes.
spec	<a href="#">NodeSpec</a> object	Detailed description of the node. CCE creates or updates objects by defining or updating <b>spec</b> .
status	<a href="#">NodeStatus</a> object	Node status, which is dynamically recorded. A user-defined value will not function when a node is being created or modified.

**Table 4-235** NodeMetadata

Parameter	Type	Description
name	String	Node name. <b>NOTE</b> Enter 1 to 56 characters starting with a lowercase letter and not ending with a hyphen (-). Only lowercase letters, digits, hyphens (-), and periods (.) are allowed. If <b>name</b> is not configured or is left blank, the node name is generated based on the default rule, which is <i>Cluster name-Random characters</i> . Only the first 36 characters are kept for a long cluster name. If the number of nodes ( <b>count</b> ) is greater than <b>1</b> , some random characters are added to the end of the entered node name according to the default rule, which is <i>User-defined name-Random characters</i> . If the length of the entered node name exceeds 50 characters, the system uses the first 50 characters and adds a random character string to the end of the name.
uid	String	Node ID, which is unique and automatically generated after the resource is created. A user-defined ID will not take effect.
labels	Map<String,String>	CCE node label (not the native Kubernetes label). Labels are used to select objects that meet certain criteria. A label is a key-value pair. Example: <pre>"labels": {   "key" : "value" }</pre>

Parameter	Type	Description
annotations	Map<String,String>	<p>CCE node annotations, in key-value pairs (not the native Kubernetes annotations). Example:</p> <pre>"annotations": {   "key1" : "value1",   "key2" : "value2" }</pre> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• <b>annotations</b> are not used to identify or select objects. The metadata in <b>annotations</b> may be small or large, structured or unstructured, and may include characters that are not allowed in labels.</li> <li>• This parameter is used only for query and cannot be input through a request. Entered data of this parameter is invalid.</li> </ul>
creationTimes tamp	String	Time when the object was created. The value is automatically generated after the object is created. A user-defined value will not take effect.
updateTimest amp	String	Time when the object was updated. The value is automatically generated after the object is created. A user-defined value will not take effect.
ownerReferen ce	<a href="#">ownerRefere nce</a> object	Owner object

**Table 4-236** ownerReference

Parameter	Type	Description
nodepoolName	String	Node pool name
nodepoolID	String	Node pool UID

**Table 4-237** NodeSpec

Parameter	Type	Description
flavor	String	Node specifications. For details about the node specifications supported by CCE, see <a href="#">Node Flavor Description</a> .

Parameter	Type	Description
az	String	<p>AZ where the node to be created is located. You need to specify the AZ name. A node created using APIs does not support random AZs.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>When you are creating a node pool and configuring a scaling group, this parameter cannot be set to <b>random</b>.</li> </ul>
os	String	<p>Node OS. For details about the supported OSs, see <a href="#">Node OS</a>.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>The system automatically selects the supported OS based on the cluster version. If the current cluster version does not support the OS, an error will be reported.</li> <li>If <b>alpha.cce/NodeImageID</b> in <b>extendParam</b> is specified during node creation, you do not need to configure this parameter.</li> <li>This parameter is mandatory when creating a node pool.</li> <li>If you are using shared disk space when creating a node, this parameter is mandatory. Specifically, the disk initialization configuration management parameter <b>storage</b> is used and the <b>name</b> field of <b>virtualSpaces</b> in <b>StorageGroups</b> is set to <b>share</b>.</li> </ul>
login	<a href="#">Login</a> object	<p>Node login mode. Either the key pair or password must be used for login.</p> <p><b>NOTE</b> This parameter cannot be modified when a node pool is being updated.</p>
rootVolume	<a href="#">Volume</a> object	Information about disks on the node
dataVolumes	Array of <a href="#">Volume</a> objects	<p>Data disk parameters of a node. A VM can have up to 16 disks attached, while a BMS can only have up to 10 disks attached. The maximum number of disks that can be attached to either a VM or BMS is limited by their respective maximum disk limits. (You can add more data disks to each of your nodes through the CCE console and APIs.) The data disk used by the container runtime and kubelet cannot be detached, or the node will become unavailable. For DeC nodes, the parameter details are the same as those for <b>rootVolume</b>.</p>

Parameter	Type	Description
storage	<a href="#">Storage</a> object	<p>Disk initialization management parameter. This parameter is complex to configure. For details, see <a href="#">Attaching Disks to a Node</a>.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If a node has both local and EVS disks attached, make sure to not keep the default value for this parameter, or it may result in unexpected disk partitions.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If you want to change the value range of a data disk to 20 to 32768, do not use the default value.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If you want to use the shared disk space (with the runtime and Kubernetes partitions cancelled), do not retain the default value of this parameter. For details about the shared disk space, see <a href="#">Data Disk Space Allocation</a>.</li> </ul>
publicIP	<a href="#">NodePublicIP</a> object	<p>EIP of a node.</p> <p><b>NOTE</b></p> <p>This parameter is not supported when you add a node to a node pool.</p>
nodeNicSpec	<a href="#">NodeNicSpec</a> object	NIC of the node
count	Integer	Number of nodes to be created in a batch. The value must be a positive integer greater than or equal to 1 and less than or equal to the defined limit. This parameter can be left blank when it is used for a node pool.
billingMode	Integer	<p>Node billing mode.</p> <ul style="list-style-type: none"> <li><b>0</b>: pay-per-use</li> </ul>

Parameter	Type	Description
taints	Array of <b>Taint</b> objects	<p>You can add taints to created nodes to configure anti-affinity. A maximum of 20 taints can be added. Each taint contains the following parameters:</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain can be prefixed to a key.</li> <li>• <b>Value:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed.</li> <li>• <b>Effect:</b> Available options are <b>NoSchedule</b>, <b>PreferNoSchedule</b>, and <b>NoExecute</b>. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</li> </ul> <p>Example:</p> <pre data-bbox="815 1108 1430 1339">"taints": [{   "key": "status",   "value": "unavailable",   "effect": "NoSchedule" }, {   "key": "looks",   "value": "bad",   "effect": "NoSchedule" }]</pre>



Parameter	Type	Description
k8sTags	Map<String,String>	<p>Defined in key-value pairs. A maximum of 20 key-value pairs are allowed.</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain containing a maximum of 253 characters (such as <b>example.com/my-key</b>) can be prefixed to a key.</li> <li>• <b>Value:</b> The value can be left blank or contain 1 to 63 characters that start and end with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</li> </ul> <p>Example:</p> <pre>"k8sTags": {   "key": "value" }</pre>
ecsGroupId	String	<p>ECS group ID. If this parameter is configured, nodes will be created in the specified ECS group.</p> <p><b>NOTE</b> This configuration does not take effect when a node pool is created. To ensure that all nodes in a node pool are in the same ECS group, configure <b>nodeManagement</b> in the node pool.</p>
dedicatedHostId	String	<p>ID of the DeH to which the node is scheduled.</p> <p><b>NOTE</b> This parameter is not supported when you add a node during node pool creation.</p>

Parameter	Type	Description
userTags	Array of <a href="#">UserTag</a> objects	<p>Cloud server tag (resource tag). The key of a tag must be unique. The maximum number of custom tags supported by CCE varies depending on regions and cannot exceed 8. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</p> <p><b>NOTE</b> A tag key can have a maximum of 36 characters. It can only contain letters, digits, Unicode characters, and special characters (-_).</p>
runtime	<a href="#">Runtime</a> object	<p>Container runtime. Default value:</p> <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>

Parameter	Type	Description
initializedConditions	Array of strings	<p>Custom initialization flag, which is left blank by default.</p> <p>Before CCE nodes are initialized, <b>node.cloudprovider.kubernetes.io/uninitialized</b> is added to the nodes to prevent pods from being scheduled to these nodes. When creating a node, you can configure the <b>initializedConditions</b> parameter to control the taint removal time. By default, the timeout period is not configured.</p> <p>Example:</p> <ol style="list-style-type: none"> <li>1. Create a node and add parameter "initializedConditions": ["CCEInitial", "CustomedInitial"].</li> <li>2. After custom initialization is complete, CCE will call a Kubernetes API (for example, <b>PATCH /v1/nodes/{node_ip}/status</b>) to update the node <b>conditions</b> by adding two labels of types <b>CCEInitial</b> and <b>CustomedInitial</b> and setting the status to <b>True</b>, as shown in the following:</li> </ol> <pre data-bbox="815 1055 1430 1205"> status:   conditions:   - type: CCEInitial     status: 'True'   - type: CustomedInitial     status: 'True' </pre> <ol style="list-style-type: none"> <li>1. CCE polls <b>status.Conditions</b> of nodes to check whether there are conditions of types <b>CCEInitial</b> and <b>CustomedInitial</b>. If such conditions exist and the status is <b>True</b>, the node initialization is complete and the initialization taint is removed.</li> <li>2. <b>initializedConditions</b> allows you to configure a timeout period during node creation. For example, "initializedConditions": ["CCEInitial:15m", "CustomedInitial:15m"] indicates that the timeout period is 15 minutes. After the timeout period is reached, the initialization conditions will be automatically ignored and the initialization taint will be removed when CCE polls the node. <ul style="list-style-type: none"> <li>• The timeout interval in <b>initializedConditions</b> ranges from 1 to 99 seconds.</li> <li>• Use only letters and digits. Max. characters: 20.</li> <li>• Max. flags: 2.</li> </ul> </li> </ol>

Parameter	Type	Description
		<ul style="list-style-type: none"> <li>The unit of the timeout period is minute (m).</li> </ul>
extendParam	<b>NodeExtendParam</b> object	Extended parameters for creating a node.
hostnameConfig	<b>HostnameConfig</b> object	Kubernetes node name configuration parameter, which is supported by clusters of v1.23.6-r0 to v1.25 or clusters of v1.25.2-r0 or later versions.
serverEnterpriseProjectID	String	Enterprise project ID of a server. CCE does not support features provided by EPS. This field is used only for synchronizing enterprise project IDs of servers. To create a node or node pool, you can specify an existing enterprise project for it. If this field is left blank, it inherits the enterprise project attributes of the cluster that runs the node or node pool. When updating a node pool, any configuration modifications you made will apply only to new nodes. You need to modify the enterprise projects of existing nodes on the EPS console. If no value is specified during the update, this field will not be updated. If this parameter is left blank, the cluster enterprise project is returned.

**Table 4-238** Login

Parameter	Type	Description
sshKey	String	Name of the key pair used for login.
userPassword	<b>UserPassword</b> object	Password used for node login.

**Table 4-239** UserPassword

Parameter	Type	Description
username	String	Login account. The default value is <b>root</b> .

Parameter	Type	Description
password	String	<p>Login password. If a username and password are used to create a node, this field is shielded in the response body. A password must meet the following complexity requirements:</p> <ul style="list-style-type: none"> <li>• Contains 8 to 26 characters.</li> <li>• Contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters !@\$%^-_=+[{ }];,./?</li> <li>• Cannot contain the username or the username spelled backwards. The <b>password</b> field must be salted during node creation. For details, see <a href="#">Adding a Salt in the password Field When Creating a Node</a>.</li> </ul>

**Table 4-240** Volume

Parameter	Type	Description
size	Integer	<p>Disk size, in GiB.</p> <ul style="list-style-type: none"> <li>• System disk: 40 to 1024</li> <li>• Value range for data disks: 100 to 32768</li> </ul>
volumetype	String	<p>Disk type. For details about possible values, see the description of the <b>root_volume</b> parameter in the API used to create an ECS.</p> <ul style="list-style-type: none"> <li>• <b>SAS</b>: high I/O SAS disk</li> <li>• <b>SSD</b>: ultra-high I/O SSD disk</li> <li>• <b>SATA</b>: common I/O SATA disk. SATA disks have been removed from EVS. You can find them attached only on existing nodes.</li> </ul>
extendParam	Map<String, Object>	<p>Extended disk parameters, defined in <b>extendparam</b> in the API used to create an ECS.</p>
cluster_id	String	<p>ID of the storage pool used by the ECS system disk. This field is used only for DeC clusters, which functions as <b>dssPoolID</b>, that is, the ID of the DSS storage pool.</p>
cluster_type	String	<p>Storage class of the cloud server system disk. The value is always <b>dss</b>. This field is used only for DeC clusters.</p>

Parameter	Type	Description
hw:passthrough	Boolean	<ul style="list-style-type: none"> <li>Pay attention to this field if your ECS is SDI-compliant. If the value of this field is <b>true</b>, the created disk is of SCSI type.</li> <li>If the node pool type is <b>ElasticBMS</b>, this field must be set to <b>true</b>.</li> <li>If a node specification involves both local and EVS disks, configure the disk initialization parameters. For details, see <a href="#">Attaching Disks to a Node</a>.</li> </ul>
metadata	<b>VolumeMetadata</b> object	EVS disk encryption information. This field is mandatory only when you need to encrypt the system disk or data disks of the node to be created.

**Table 4-241** VolumeMetadata

Parameter	Type	Description
__system__encrypted	String	Whether the EVS disk is encrypted. The value <b>0</b> indicates that the EVS disk is not encrypted, and the value <b>1</b> indicates that the EVS disk is encrypted. If this parameter is not specified, EVS disks will not be encrypted by default.
__system__cmkid	String	CMK ID, which indicates encryption in <b>metadata</b> . This field is used with <b>__system__encrypted</b> .

**Table 4-242** Storage

Parameter	Type	Description
storageSelectors	Array of <b>StorageSelectors</b> objects	Disk selection. Matched disks are managed according to <b>matchLabels</b> and <b>storageType</b> .
storageGroups	Array of <b>StorageGroups</b> objects	A storage group consists of multiple storage devices. It is used to divide storage space.

**Table 4-243** StorageSelectors

Parameter	Type	Description
name	String	Selector name, used as the index of <b>selectorNames</b> in <b>storageGroup</b> . Therefore, the name of each selector must be unique.
storageType	String	Storage type. Only <b>evs</b> (EVS disks) and <b>local</b> (local disks) are supported. The local storage does not support disk selection. All local disks will form a VG, so only one storageSelector of the local type is allowed.
matchLabels	<b>matchLabels</b> object	Matching field of an EVS disk. The value can be <b>size</b> , <b>volumeType</b> , <b>metadataEncrypted</b> , <b>metadataCmkid</b> , or <b>count</b> in <b>DataVolume</b> .

**Table 4-244** matchLabels

Parameter	Type	Description
size	String	Matched disk size. If this parameter is left unspecified, the disk size is not limited. Example: 100
volumeType	String	EVS disk type. <b>SSD</b> , <b>GPSSD</b> , <b>SAS</b> , <b>ESSD</b> , and <b>SATA</b> are supported. If this parameter is left blank, there is no restriction on the disk type.
metadataEncrypted	String	Disk encryption identifier. If this parameter is set to <b>0</b> , it specifies that the disk is not encrypted. If this parameter is set to <b>1</b> , it specifies that the disk is encrypted. If this parameter is left blank, there is no restriction on the disk encryption identifier.
metadataCmkid	String	CMK ID of an encrypted disk. The value is a string of 36 bytes. If this parameter is left blank, there is no restriction on the disk key ID.
count	String	Number of disks to be selected. If this parameter is left blank, all disks of this type are selected.

**Table 4-245** StorageGroups

Parameter	Type	Description
name	String	Name of a virtual storage group, which must be unique. <b>NOTE</b> <ul style="list-style-type: none"> <li>If <b>cceManaged</b> is set to <b>true</b>, the name must be set to <b>vgpaas</b>.</li> <li>If the data disk is used as a temporary storage volume, the name must be <b>vg-everest-localvolume-ephemeral</b>.</li> <li>If the data disk is used as a persistent storage volume, the name must be <b>vg-everest-localvolume-persistent</b>.</li> </ul>
cceManaged	Boolean	Storage space for Kubernetes and runtime components. Only one group can be set to <b>true</b> . If this parameter is left blank, the default value <b>false</b> is used.
selectorNames	Array of strings	Corresponds to <b>name</b> in <b>storageSelectors</b> . One group can select multiple selectors, but one selector can be selected by only one group.
virtualSpaces	Array of <a href="#">VirtualSpace</a> objects	Detailed management of space configuration in a group.

**Table 4-246** VirtualSpace

Parameter	Type	Description
name	String	Name of a virtualSpace. Options: <b>share</b> , <b>kubernetes</b> , <b>runtime</b> , and <b>user</b> <ul style="list-style-type: none"> <li><b>share</b>: shared space configuration (with the runtime and Kubernetes partitions cancelled). <b>lvmConfig</b> needs to be configured.</li> <li><b>Kubernetes</b>: Kubernetes space configuration. <b>lvmConfig</b> needs to be configured.</li> <li><b>runtime</b>: runtime space configuration. <b>runtimeConfig</b> needs to be configured.</li> <li><b>user</b>: user space configuration. <b>lvmConfig</b> needs to be configured.</li> </ul>



Parameter	Type	Description
size	String	Size of a virtualSpace. The value must be an integer in percentage. Example: 90%. <b>NOTE</b> The value ranges from 10 to 90. The sum of percentages of all virtualSpaces in a group cannot exceed 100%.
lvmConfig	<b>LVMConfig</b> object	LVM configuration managements, applicable to the <i>share*</i> , <b>kubernetes</b> , and <b>user</b> spaces. Note that one virtual space supports only one config.
runtimeConfig	<b>RuntimeConfig</b> object	runtime configurations, applicable to the <b>runtime</b> space. Note that one virtual space supports only one config.

**Table 4-247** LVMConfig

Parameter	Type	Description
lvType	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.
path	String	Disk mount path. This parameter takes effect only in user configuration. The value is an absolute path. Digits, letters, periods (.), hyphens (-), and underscores (_) are allowed. <b>NOTE</b> Do not leave the path empty or set it to a key OS path such as the root directory. Key OS paths that cannot be specified: /, /home, /home/, /bin, /bin/, /lib, /lib/, /root, /root/, /boot, /boot/, /dev, /dev/, /etc, /etc/, /lost+found, /lost+found/, /mnt, /mnt/, /proc, /proc/, /sbin, /sbin/, /srv, /srv/, /tmp, /tmp/, /var, /var/, /media, /media/, /opt, /opt/, /selinux, /selinux/, /sys, /sys/, /usr, /usr/, /opt/cloud/, /mnt/paas/, /home/paas/, /var/paas/, /var/lib/, /var/script/

**Table 4-248** RuntimeConfig

Parameter	Type	Description
lvType	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.

**Table 4-249** NodePublicIP

Parameter	Type	Description
ids	Array of strings	IDs of existing EIPs. The quantity cannot be greater than the number of nodes to be created. <b>NOTE</b> If <b>ids</b> has been set, you do not need to set <b>count</b> and <b>eip</b> .
count	Integer	Number of EIPs to be dynamically created. <b>NOTE</b> <b>count</b> and <b>eip</b> must be set at the same time.
eip	<b>NodeEIPSpec</b> object	EIP configuration. <b>NOTE</b> This parameter is not supported when creating a node pool.

**Table 4-250** NodeEIPSpec

Parameter	Type	Description
iptype	String	EIP type, specified in <b>publicip.type</b> in the API for assigning an EIP.
bandwidth	<b>NodeBandwidth</b> object	Bandwidth parameters of the EIP

**Table 4-251** NodeBandwidth

Parameter	Type	Description
chargemode	String	<p>Bandwidth billing mode.</p> <ul style="list-style-type: none"> <li>• If this field is not specified, the billing is based on bandwidth.</li> <li>• If the field is null, the billing is based on bandwidth.</li> <li>• If the field value is <b>traffic</b>, the billing is based on traffic.</li> <li>• If the value is out of the preceding options, the cloud server will fail to be created.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• Billed by bandwidth: The billing will be based on the data transmission rate (in Mbps) of public networks. This billing mode is recommended if your bandwidth usage is higher than 10%.</li> <li>• Billed by traffic: The billing is based on the total amount of data (in GB) transmitted over the public network. This mode is available only when you are creating a pay-per-use node. This billing mode is recommended if your bandwidth usage is lower than 10%.</li> </ul>
size	Integer	Bandwidth size, specified in <b>bandwidth.size</b> in the API for assigning an EIP.
sharetype	String	Bandwidth sharing type. Value options: <b>PER</b> (exclusive bandwidth)

**Table 4-252** NodeNicSpec

Parameter	Type	Description
primaryNic	<b>NicSpec</b> object	Description of the primary NIC.
extNics	Array of <b>NicSpec</b> objects	<p>Extension NIC</p> <p><b>NOTE</b> This parameter is not supported when you add a node to a node pool.</p>

**Table 4-253 NicSpec**

Parameter	Type	Description
subnetId	String	Network ID of the subnet to which a NIC belongs. If <b>subnetId</b> is not specified when a primary NIC is creating, the cluster subnet will be used. If <b>subnetList</b> is also configured for a node pool, the <b>subnetList</b> field is used for adding subnets to the node pool. When creating an extension NIC, you must specify <b>subnetId</b> .
fixedIps	Array of strings	The IP address of the primary ENI is specified using <b>fixedIps</b> . The number of IP addresses cannot be greater than the number of created nodes. Either <b>fixedIps</b> or <b>ipBlock</b> can be specified. <b>fixedIps</b> cannot be specified for extension ENIs.
ipBlock	String	CIDR format of the primary NIC IP range. The IP address of the created node falls in this range. <b>fixedIps</b> and <b>ipBlock</b> cannot be specified at the same time.
subnetList	Array of strings	Network ID list of the subnet where the NIC resides. A maximum of 20 subnets can be configured for a node pool.

**Table 4-254 Taint**

Parameter	Type	Description
key	String	Key.
value	String	Value.
effect	String	Effect.

**Table 4-255 UserTag**

Parameter	Type	Description
key	String	Key of the cloud server label. The value cannot start with <b>CCE-</b> or <b>__type_baremetal</b> .
value	String	Value of the cloud server label.

**Table 4-256** Runtime

Parameter	Type	Description
name	String	<p>Container runtime. Default value:</p> <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>

**Table 4-257** NodeExtendParam

Parameter	Type	Description
ecs:performancetype	String	ECS flavor types. This field is returned in the response.
maxPods	Integer	<p>Maximum number of pods that can be created on a node, including the default system pods. Value range: 16 to 256.</p> <p>This limit prevents the node from being overloaded of pods.</p> <p>The number of pods that can be created on a node is determined by multiple parameters. For details, see <a href="#">Maximum Number of Pods That Can Be Created on a Node</a>.</p>

Parameter	Type	Description
DockerLVMConfigOverride	String	<p>Docker data disk configuration item. (This parameter has been discarded. Use the <b>storage</b> field instead.) Example default configuration: "DockerLVMConfigOverride":"dockerThinpool=vgpaas/90%VG;kubernetesLV=vgpaas/10%VG;diskType=evs;lvType=linear"</p> <p>By default, if no VD disk is available, an error occurs because the data disk fails to be found. Set <b>diskType</b> based on the actual drive letter type. The following fields are included:</p> <ul style="list-style-type: none"> <li>• <b>userLV</b> (optional): size of the user space, for example, <b>vgpaas/20%VG</b>.</li> <li>• <b>userPath</b> (optional): mount path of the user space, for example, <b>/home/wqt-test</b>.</li> <li>• <b>diskType</b>: disk type. Currently, only <b>evs</b>, <b>hdd</b>, and <b>ssd</b> are supported.</li> <li>• <b>lvType</b>: type of a logic volume. The value can be <b>linear</b> or <b>striped</b>.</li> <li>• <b>dockerThinpool</b>: Docker space size, for example, <b>vgpaas/60%VG</b>.</li> <li>• <b>kubernetesLV</b>: kubelet space size, for example, <b>vgpaas/20%VG</b>.</li> </ul>
dockerBaseSize	Integer	<p>Available disk space of a single container on a node, in GB. This parameter has been discarded. Use <b>containerBaseSize</b>.</p> <p>If this parameter is left blank or is set to <b>0</b>, the default value is used. In Device Mapper mode, the default value is <b>10</b>. In OverlayFS mode, the available space of a single container is not limited by default, and the <b>dockerBaseSize</b> setting takes effect only on nodes running EulerOS in clusters of the new versions.</p> <p>For details about how to allocate the space for the container runtime, see <a href="#">Data Disk Space Allocation</a>.</p> <p>When Device Mapper is used, it is recommended that <b>dockerBaseSize</b> be set to a value less than or equal to 80 GB. If the value is too large, the container runtime may fail to be started due to long initialization. If there are special requirements for the container disk space, you can mount an external or local storage device.</p>

Parameter	Type	Description
containerBaseSize	Integer	<p>Available disk space of a single container on a node, in GiB.</p> <p>If this parameter is left blank or is set to <b>0</b>, the default value is used. In OverlayFS mode, the available space of a single container is not limited by default. In Device Mapper mode, the default value is <b>10</b>, and the <b>containerBaseSize</b> setting takes effect only on nodes running EulerOS in clusters v1.23.14-r0, v1.25.9-r0, v1.27.6-r0, v1.28.4-r0 or later.</p> <p>For details about how to allocate the space for the container runtime, see <a href="#">Data Disk Space Allocation</a>.</p> <p>When Device Mapper is used, it is recommended that <b>containerBaseSize</b> be set to a value less than or equal to 80 GB. If the value is too large, the container runtime may fail to be started due to long initialization. If there are special requirements for the container disk space, you can mount an external or local storage device. In new versions, the Device Mapper mode is used only by BMSs in the same resource pool and is being discarded.</p> <p><b>NOTE</b> This parameter cannot be updated when a node pool is being updated.</p>
publicKey	String	Public key of a node.
alpha.cce/preInstall	String	<p>Pre-installation scripts. The total number of characters after transcoding cannot exceed 10240.</p> <p><b>NOTE</b> The input value must be encoded using Base64. The method is as follows:</p> <pre>echo -n "Content to be encoded"   base64</pre>
alpha.cce/postInstall	String	<p>Post-installation scripts. The total number of characters after transcoding cannot exceed 10240.</p> <p><b>NOTE</b> The input value must be encoded using Base64. The method is as follows:</p> <pre>echo -n "Content to be encoded"   base64</pre>
alpha.cce/NodeImageID	String	This parameter is required when a custom image is used to create a BMS node.

Parameter	Type	Description
chargingMode	Integer	Billing mode of a node. This parameter has been deprecated. Use the <b>billingMode</b> parameter in <b>NodeSpec</b> .
agency_name	String	Name of an agency An agency is created by a tenant administrator on Identity and Access Management (IAM) to provide temporary credentials for CCE nodes to access cloud servers. This parameter is returned only when it is transferred during node creation.
kubeReservedMem	Integer	Reserved node memory, which is reserved for Kubernetes components.
systemReservedMem	Integer	Reserved node memory, which is reserved for system components.
init-node-password	String	Node password. If this parameter is used as a response parameter, asterisks (*) will be displayed.

**Table 4-258** HostnameConfig

Parameter	Type	Description
type	String	Configuration type of the Kubernetes node name. The default value is <b>privatelp</b> . <ul style="list-style-type: none"> <li>• <b>privatelp</b>: The Kubernetes node is named after its IP address.</li> <li>• <b>cceNodeName</b>: The Kubernetes node is named after the CCE node.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• For a node which is configured using <b>cceNodeName</b>, the name is the same as the Kubernetes node name and the ECS name. The node name cannot be changed. If the ECS name is changed on the ECS console, the node name will retain unchanged after ECS synchronization.</li> <li>• For a node which is configured using <b>cceNodeName</b>, to avoid a conflict between Kubernetes nodes, the system automatically adds a suffix to each node name. The suffix is in the format of <b>A hyphen (-) Five random characters</b>. The value of the random characters is a lowercase letter or a digit ranging from 0 to 9.</li> </ul>



**Table 4-259** NodeStatus

Parameter	Type	Description
phase	String	Node status, which is the status of a node during its lifecycle (such as installation and uninstallation) or the status of a node in a Kubernetes cluster. Options: <ul style="list-style-type: none"> <li>● <b>Build</b>: The node is being created.</li> <li>● <b>Installing</b>: The node is being managed.</li> <li>● <b>Upgrading</b>: The node is being upgraded.</li> <li>● <b>Active</b>: The node is running properly.</li> <li>● <b>Abnormal</b>: The node is abnormal.</li> <li>● <b>Deleting</b>: The node is being deleted.</li> <li>● <b>Error</b>: The node is faulty.</li> </ul>
lastProbeTime	String	Last time when the node status was checked. If the cluster is in the abnormal, frozen, or intermediate state (for example, creating), the node status check may be affected. The node status that takes more than five minutes to check has no reference value.
jobID	String	ID of a creation or deletion job
serverId	String	ID of the underlying ECS or BMS node
privateIP	String	IP address in the private network segment of the primary NIC on the node
privateIPv6IP	String	IPv6 address in the private network segment of the primary NIC on the node
publicIP	String	Node EIP. If the ECS data is not synchronized in real time, you can click <b>Sync Node Data</b> on the console to manually update the data.
deleteStatus	<a href="#">DeleteStatus</a> object	Resource status during resource deletion.

**Table 4-260** DeleteStatus

Parameter	Type	Description
previous_total	Integer	Total number of existing cluster resource records when the cluster is deleted.
current_total	Integer	Latest number of resource records, which is generated based on the current cluster resource records.

Parameter	Type	Description
updated	Integer	Total number of resource records updated when the cluster is deleted.
added	Integer	Total number of resource records updated when the cluster is deleted.
deleted	Integer	Total number of resource records deleted when the cluster is deleted.

## Example Requests

None

## Example Responses

**Status code: 200**

Information about the specified node is successfully obtained.

```
{
  "kind": "Node",
  "apiVersion": "v3",
  "metadata": {
    "name": "myhost",
    "uid": "4d1ecb2c-229a-11e8-9c75-0255ac100ceb",
    "creationTimestamp": "2018-08-02 08:12:40.124294439 +0000 UTC",
    "updateTimestamp": "2018-08-02 08:18:20.221871842 +0000 UTC",
    "annotations": {
      "kubernetes.io/node-pool.id": "az1.dc1#s1.medium#EulerOS 2.2"
    }
  },
  "spec": {
    "flavor": "s1.medium",
    "az": "*****",
    "os": "EulerOS 2.2",
    "login": {
      "sshKey": "KeyPair-001"
    },
    "rootVolume": {
      "volumetype": "SAS",
      "size": 40
    },
    "dataVolumes": [ {
      "volumetype": "SAS",
      "size": 100
    } ],
    "publicIP": {
      "eip": {
        "bandwidth": { }
      }
    },
    "billingMode": 0
  },
  "status": {
    "phase": "Active",
    "serverId": "456789abc-9368-46f3-8f29-d1a95622a568",
    "publicIP": "10.34.56.78",
    "privateIP": "192.168.1.23"
  }
}
```

## Status Codes

Status Code	Description
200	Information about the specified node is successfully obtained.

## Error Codes

See [Error Codes](#).

### 4.3.3 Listing All Nodes in a Cluster

#### Function

This API is used to obtain details about all nodes in a specified cluster with the cluster ID.

#### NOTE

The URL for cluster management is in the format of `https://Endpoint/uri`. In the URL, **uri** indicates the resource path, that is, the path for API access.

#### URI

GET `/api/v3/projects/{project_id}/clusters/{cluster_id}/nodes`

**Table 4-261** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

#### Request Parameters

**Table 4-262** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

## Response Parameters

Status code: 200

**Table 4-263** Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at <b>List</b> .
apiVersion	String	API version. The value is fixed at <b>v3</b> .
items	Array of <a href="#">Node</a> objects	List of details for all nodes in the current cluster. You can filter nodes by <b>items.metadata.name</b> .

**Table 4-264** Node

Parameter	Type	Description
kind	String	API type. The value is fixed at <b>Node</b> and cannot be changed.
apiVersion	String	API version. The value is fixed at <b>v3</b> and cannot be changed.
metadata	<a href="#">NodeMetadata</a> object	Node metadata, which is a collection of attributes.
spec	<a href="#">NodeSpec</a> object	Detailed description of the node. CCE creates or updates objects by defining or updating <b>spec</b> .
status	<a href="#">NodeStatus</a> object	Node status, which is dynamically recorded. A user-defined value will not function when a node is being created or modified.

**Table 4-265** NodeMetadata

Parameter	Type	Description
name	String	<p>Node name.</p> <p><b>NOTE</b></p> <p>Enter 1 to 56 characters starting with a lowercase letter and not ending with a hyphen (-). Only lowercase letters, digits, hyphens (-), and periods (.) are allowed. If <b>name</b> is not configured or is left blank, the node name is generated based on the default rule, which is <i>Cluster name-Random characters</i>. Only the first 36 characters are kept for a long cluster name. If the number of nodes (<b>count</b>) is greater than <b>1</b>, some random characters are added to the end of the entered node name according to the default rule, which is <i>User-defined name-Random characters</i>. If the length of the entered node name exceeds 50 characters, the system uses the first 50 characters and adds a random character string to the end of the name.</p>
uid	String	<p>Node ID, which is unique and automatically generated after the resource is created. A user-defined ID will not take effect.</p>
labels	Map<String,String>	<p>CCE node label (not the native Kubernetes label).</p> <p>Labels are used to select objects that meet certain criteria. A label is a key-value pair.</p> <p>Example:</p> <pre>"labels": {   "key" : "value" }</pre>
annotations	Map<String,String>	<p>CCE node annotations, in key-value pairs (not the native Kubernetes annotations). Example:</p> <pre>"annotations": {   "key1" : "value1",   "key2" : "value2" }</pre> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li><b>annotations</b> are not used to identify or select objects. The metadata in <b>annotations</b> may be small or large, structured or unstructured, and may include characters that are not allowed in labels.</li> <li>This parameter is used only for query and cannot be input through a request. Entered data of this parameter is invalid.</li> </ul>
creationTimes tamp	String	<p>Time when the object was created. The value is automatically generated after the object is created. A user-defined value will not take effect.</p>

Parameter	Type	Description
updateTimestamp	String	Time when the object was updated. The value is automatically generated after the object is created. A user-defined value will not take effect.
ownerReference	<a href="#">ownerReference</a> object	Owner object

**Table 4-266** ownerReference

Parameter	Type	Description
nodepoolName	String	Node pool name
nodepoolID	String	Node pool UID

**Table 4-267** NodeSpec

Parameter	Type	Description
flavor	String	Node specifications. For details about the node specifications supported by CCE, see <a href="#">Node Flavor Description</a> .
az	String	AZ where the node to be created is located. You need to specify the AZ name. A node created using APIs does not support random AZs. <b>NOTE</b> <ul style="list-style-type: none"> <li>When you are creating a node pool and configuring a scaling group, this parameter cannot be set to <b>random</b>.</li> </ul>

Parameter	Type	Description
os	String	<p>Node OS. For details about the supported OSs, see <a href="#">Node OS</a>.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>The system automatically selects the supported OS based on the cluster version. If the current cluster version does not support the OS, an error will be reported.</li> <li>If <code>alpha.cce/NodeImageID</code> in <code>extendParam</code> is specified during node creation, you do not need to configure this parameter.</li> <li>This parameter is mandatory when creating a node pool.</li> <li>If you are using shared disk space when creating a node, this parameter is mandatory. Specifically, the disk initialization configuration management parameter <code>storage</code> is used and the <code>name</code> field of <code>virtualSpaces</code> in <code>StorageGroups</code> is set to <code>share</code>.</li> </ul>
login	<a href="#">Login</a> object	<p>Node login mode. Either the key pair or password must be used for login.</p> <p><b>NOTE</b> This parameter cannot be modified when a node pool is being updated.</p>
rootVolume	<a href="#">Volume</a> object	Information about disks on the node
dataVolumes	Array of <a href="#">Volume</a> objects	<p>Data disk parameters of a node. A VM can have up to 16 disks attached, while a BMS can only have up to 10 disks attached. The maximum number of disks that can be attached to either a VM or BMS is limited by their respective maximum disk limits. (You can add more data disks to each of your nodes through the CCE console and APIs.) The data disk used by the container runtime and kubelet cannot be detached, or the node will become unavailable. For DeC nodes, the parameter details are the same as those for <b>rootVolume</b>.</p>

Parameter	Type	Description
storage	<a href="#">Storage</a> object	<p>Disk initialization management parameter. This parameter is complex to configure. For details, see <a href="#">Attaching Disks to a Node</a>.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If a node has both local and EVS disks attached, make sure to not keep the default value for this parameter, or it may result in unexpected disk partitions.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If you want to change the value range of a data disk to 20 to 32768, do not use the default value.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If you want to use the shared disk space (with the runtime and Kubernetes partitions cancelled), do not retain the default value of this parameter. For details about the shared disk space, see <a href="#">Data Disk Space Allocation</a>.</li> </ul>
publicIP	<a href="#">NodePublicIP</a> object	<p>EIP of a node.</p> <p><b>NOTE</b></p> <p>This parameter is not supported when you add a node to a node pool.</p>
nodeNicSpec	<a href="#">NodeNicSpec</a> object	NIC of the node
count	Integer	Number of nodes to be created in a batch. The value must be a positive integer greater than or equal to 1 and less than or equal to the defined limit. This parameter can be left blank when it is used for a node pool.
billingMode	Integer	<p>Node billing mode.</p> <ul style="list-style-type: none"> <li><b>0</b>: pay-per-use</li> </ul>



Parameter	Type	Description
taints	Array of <b>Taint</b> objects	<p>You can add taints to created nodes to configure anti-affinity. A maximum of 20 taints can be added. Each taint contains the following parameters:</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain can be prefixed to a key.</li> <li>• <b>Value:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed.</li> <li>• <b>Effect:</b> Available options are <b>NoSchedule</b>, <b>PreferNoSchedule</b>, and <b>NoExecute</b>. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</li> </ul> <p>Example:</p> <pre data-bbox="815 1108 1430 1339">"taints": [{   "key": "status",   "value": "unavailable",   "effect": "NoSchedule" }, {   "key": "looks",   "value": "bad",   "effect": "NoSchedule" }]</pre>

Parameter	Type	Description
k8sTags	Map<String,String>	<p>Defined in key-value pairs. A maximum of 20 key-value pairs are allowed.</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain containing a maximum of 253 characters (such as <b>example.com/my-key</b>) can be prefixed to a key.</li> <li>• <b>Value:</b> The value can be left blank or contain 1 to 63 characters that start and end with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</li> </ul> <p>Example:</p> <pre>"k8sTags": {   "key": "value" }</pre>
ecsGroupId	String	<p>ECS group ID. If this parameter is configured, nodes will be created in the specified ECS group.</p> <p><b>NOTE</b> This configuration does not take effect when a node pool is created. To ensure that all nodes in a node pool are in the same ECS group, configure <b>nodeManagement</b> in the node pool.</p>
dedicatedHostId	String	<p>ID of the DeH to which the node is scheduled.</p> <p><b>NOTE</b> This parameter is not supported when you add a node during node pool creation.</p>

Parameter	Type	Description
userTags	Array of <a href="#">UserTag</a> objects	<p>Cloud server tag (resource tag). The key of a tag must be unique. The maximum number of custom tags supported by CCE varies depending on regions and cannot exceed 8. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</p> <p><b>NOTE</b> A tag key can have a maximum of 36 characters. It can only contain letters, digits, Unicode characters, and special characters (-_).</p>
runtime	<a href="#">Runtime</a> object	<p>Container runtime. Default value:</p> <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>

Parameter	Type	Description
initializedConditions	Array of strings	<p>Custom initialization flag, which is left blank by default.</p> <p>Before CCE nodes are initialized, <b>node.cloudprovider.kubernetes.io/uninitialized</b> is added to the nodes to prevent pods from being scheduled to these nodes. When creating a node, you can configure the <b>initializedConditions</b> parameter to control the taint removal time. By default, the timeout period is not configured.</p> <p>Example:</p> <ol style="list-style-type: none"> <li>1. Create a node and add parameter "initializedConditions": ["CCEInitial", "CustomedInitial"].</li> <li>2. After custom initialization is complete, CCE will call a Kubernetes API (for example, <b>PATCH /v1/nodes/{node_ip}/status</b>) to update the node <b>conditions</b> by adding two labels of types <b>CCEInitial</b> and <b>CustomedInitial</b> and setting the status to <b>True</b>, as shown in the following:</li> </ol> <pre data-bbox="815 1055 1430 1205"> status:   conditions:   - type: CCEInitial     status: 'True'   - type: CustomedInitial     status: 'True'                     </pre> <ol style="list-style-type: none"> <li>1. CCE polls <b>status.Conditions</b> of nodes to check whether there are conditions of types <b>CCEInitial</b> and <b>CustomedInitial</b>. If such conditions exist and the status is <b>True</b>, the node initialization is complete and the initialization taint is removed.</li> <li>2. <b>initializedConditions</b> allows you to configure a timeout period during node creation. For example, "initializedConditions": ["CCEInitial:15m", "CustomedInitial:15m"] indicates that the timeout period is 15 minutes. After the timeout period is reached, the initialization conditions will be automatically ignored and the initialization taint will be removed when CCE polls the node. <ul style="list-style-type: none"> <li>● The timeout interval in <b>initializedConditions</b> ranges from 1 to 99 seconds.</li> <li>● Use only letters and digits. Max. characters: 20.</li> <li>● Max. flags: 2.</li> </ul> </li> </ol>

Parameter	Type	Description
		<ul style="list-style-type: none"> <li>The unit of the timeout period is minute (m).</li> </ul>
extendParam	<a href="#">NodeExtendParam</a> object	Extended parameters for creating a node.
hostnameConfig	<a href="#">HostnameConfig</a> object	Kubernetes node name configuration parameter, which is supported by clusters of v1.23.6-r0 to v1.25 or clusters of v1.25.2-r0 or later versions.
serverEnterpriseProjectID	String	Enterprise project ID of a server. CCE does not support features provided by EPS. This field is used only for synchronizing enterprise project IDs of servers. To create a node or node pool, you can specify an existing enterprise project for it. If this field is left blank, it inherits the enterprise project attributes of the cluster that runs the node or node pool. When updating a node pool, any configuration modifications you made will apply only to new nodes. You need to modify the enterprise projects of existing nodes on the EPS console. If no value is specified during the update, this field will not be updated. If this parameter is left blank, the cluster enterprise project is returned.

**Table 4-268** Login

Parameter	Type	Description
sshKey	String	Name of the key pair used for login.
userPassword	<a href="#">UserPassword</a> object	Password used for node login.

**Table 4-269** UserPassword

Parameter	Type	Description
username	String	Login account. The default value is <b>root</b> .

Parameter	Type	Description
password	String	<p>Login password. If a username and password are used to create a node, this field is shielded in the response body. A password must meet the following complexity requirements:</p> <ul style="list-style-type: none"> <li>• Contains 8 to 26 characters.</li> <li>• Contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters !@\$%^-_=+[{ }];,./?</li> <li>• Cannot contain the username or the username spelled backwards. The <b>password</b> field must be salted during node creation. For details, see <a href="#">Adding a Salt in the password Field When Creating a Node</a>.</li> </ul>

Table 4-270 Volume

Parameter	Type	Description
size	Integer	<p>Disk size, in GiB.</p> <ul style="list-style-type: none"> <li>• System disk: 40 to 1024</li> <li>• Value range for data disks: 100 to 32768</li> </ul>
volumetype	String	<p>Disk type. For details about possible values, see the description of the <b>root_volume</b> parameter in the API used to create an ECS.</p> <ul style="list-style-type: none"> <li>• <b>SAS</b>: high I/O SAS disk</li> <li>• <b>SSD</b>: ultra-high I/O SSD disk</li> <li>• <b>SATA</b>: common I/O SATA disk SATA disks have been removed from EVS. You can find them attached only on existing nodes.</li> </ul>
extendParam	Map<String, Object>	<p>Extended disk parameters, defined in <b>extendparam</b> in the API used to create an ECS.</p>
cluster_id	String	<p>ID of the storage pool used by the ECS system disk. This field is used only for DeC clusters, which functions as <b>dssPoolID</b>, that is, the ID of the DSS storage pool.</p>
cluster_type	String	<p>Storage class of the cloud server system disk. The value is always <b>dss</b>. This field is used only for DeC clusters.</p>

Parameter	Type	Description
hw:passthrough	Boolean	<ul style="list-style-type: none"> <li>Pay attention to this field if your ECS is SDI-compliant. If the value of this field is <b>true</b>, the created disk is of SCSI type.</li> <li>If the node pool type is <b>ElasticBMS</b>, this field must be set to <b>true</b>.</li> <li>If a node specification involves both local and EVS disks, configure the disk initialization parameters. For details, see <a href="#">Attaching Disks to a Node</a>.</li> </ul>
metadata	<b>VolumeMetadata</b> object	EVS disk encryption information. This field is mandatory only when you need to encrypt the system disk or data disks of the node to be created.

**Table 4-271** VolumeMetadata

Parameter	Type	Description
__system__encrypted	String	Whether the EVS disk is encrypted. The value <b>0</b> indicates that the EVS disk is not encrypted, and the value <b>1</b> indicates that the EVS disk is encrypted. If this parameter is not specified, EVS disks will not be encrypted by default.
__system__cmkid	String	CMK ID, which indicates encryption in <b>metadata</b> . This field is used with <b>__system__encrypted</b> .

**Table 4-272** Storage

Parameter	Type	Description
storageSelectors	Array of <b>StorageSelectors</b> objects	Disk selection. Matched disks are managed according to <b>matchLabels</b> and <b>storageType</b> .
storageGroups	Array of <b>StorageGroups</b> objects	A storage group consists of multiple storage devices. It is used to divide storage space.

**Table 4-273** StorageSelectors

Parameter	Type	Description
name	String	Selector name, used as the index of <b>selectorNames</b> in <b>storageGroup</b> . Therefore, the name of each selector must be unique.
storageType	String	Storage type. Only <b>evs</b> (EVS disks) and <b>local</b> (local disks) are supported. The local storage does not support disk selection. All local disks will form a VG, so only one storageSelector of the local type is allowed.
matchLabels	<b>matchLabels</b> object	Matching field of an EVS disk. The value can be <b>size</b> , <b>volumeType</b> , <b>metadataEncrypted</b> , <b>metadataCmkid</b> , or <b>count</b> in <b>DataVolume</b> .

**Table 4-274** matchLabels

Parameter	Type	Description
size	String	Matched disk size. If this parameter is left unspecified, the disk size is not limited. Example: 100
volumeType	String	EVS disk type. <b>SSD</b> , <b>GPSSD</b> , <b>SAS</b> , <b>ESSD</b> , and <b>SATA</b> are supported. If this parameter is left blank, there is no restriction on the disk type.
metadataEncrypted	String	Disk encryption identifier. If this parameter is set to <b>0</b> , it specifies that the disk is not encrypted. If this parameter is set to <b>1</b> , it specifies that the disk is encrypted. If this parameter is left blank, there is no restriction on the disk encryption identifier.
metadataCmkid	String	CMK ID of an encrypted disk. The value is a string of 36 bytes. If this parameter is left blank, there is no restriction on the disk key ID.
count	String	Number of disks to be selected. If this parameter is left blank, all disks of this type are selected.



**Table 4-275** StorageGroups

Parameter	Type	Description
name	String	Name of a virtual storage group, which must be unique. <b>NOTE</b> <ul style="list-style-type: none"> <li>If <b>cceManaged</b> is set to <b>true</b>, the name must be set to <b>vgpaas</b>.</li> <li>If the data disk is used as a temporary storage volume, the name must be <b>vg-everest-localvolume-ephemeral</b>.</li> <li>If the data disk is used as a persistent storage volume, the name must be <b>vg-everest-localvolume-persistent</b>.</li> </ul>
cceManaged	Boolean	Storage space for Kubernetes and runtime components. Only one group can be set to <b>true</b> . If this parameter is left blank, the default value <b>false</b> is used.
selectorNames	Array of strings	Corresponds to <b>name</b> in <b>storageSelectors</b> . One group can select multiple selectors, but one selector can be selected by only one group.
virtualSpaces	Array of <a href="#">VirtualSpace</a> objects	Detailed management of space configuration in a group.

**Table 4-276** VirtualSpace

Parameter	Type	Description
name	String	Name of a virtualSpace. Options: <b>share</b> , <b>kubernetes</b> , <b>runtime</b> , and <b>user</b> <ul style="list-style-type: none"> <li><b>share</b>: shared space configuration (with the runtime and Kubernetes partitions cancelled). <b>lvmConfig</b> needs to be configured.</li> <li><b>Kubernetes</b>: Kubernetes space configuration. <b>lvmConfig</b> needs to be configured.</li> <li><b>runtime</b>: runtime space configuration. <b>runtimeConfig</b> needs to be configured.</li> <li><b>user</b>: user space configuration. <b>lvmConfig</b> needs to be configured.</li> </ul>

Parameter	Type	Description
size	String	Size of a virtualSpace. The value must be an integer in percentage. Example: 90%. <b>NOTE</b> The value ranges from 10 to 90. The sum of percentages of all virtualSpaces in a group cannot exceed 100%.
lvmConfig	<b>LVMConfig</b> object	LVM configuration managements, applicable to the <i>share*</i> , <b>kubernetes</b> , and <b>user</b> spaces. Note that one virtual space supports only one config.
runtimeConfig	<b>RuntimeConfig</b> object	runtime configurations, applicable to the <b>runtime</b> space. Note that one virtual space supports only one config.

**Table 4-277** LVMConfig

Parameter	Type	Description
lvType	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.
path	String	Disk mount path. This parameter takes effect only in user configuration. The value is an absolute path. Digits, letters, periods (.), hyphens (-), and underscores (_) are allowed. <b>NOTE</b> Do not leave the path empty or set it to a key OS path such as the root directory. Key OS paths that cannot be specified: /, /home, /home/, /bin, /bin/, /lib, /lib/, /root, /root/, /boot, /boot/, /dev, /dev/, /etc, /etc/, /lost+found, /lost+found/, /mnt, /mnt/, /proc, /proc/, /sbin, /sbin/, /srv, /srv/, /tmp, /tmp/, /var, /var/, /media, /media/, /opt, /opt/, /selinux, /selinux/, /sys, /sys/, /usr, /usr/, /opt/cloud/, /mnt/paas/, /home/paas/, /var/paas/, /var/lib/, /var/script/

**Table 4-278** RuntimeConfig

Parameter	Type	Description
lvType	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.

**Table 4-279** NodePublicIP

Parameter	Type	Description
ids	Array of strings	IDs of existing EIPs. The quantity cannot be greater than the number of nodes to be created. <b>NOTE</b> If <b>ids</b> has been set, you do not need to set <b>count</b> and <b>eip</b> .
count	Integer	Number of EIPs to be dynamically created. <b>NOTE</b> <b>count</b> and <b>eip</b> must be set at the same time.
eip	<b>NodeEIPSpec</b> object	EIP configuration. <b>NOTE</b> This parameter is not supported when creating a node pool.

**Table 4-280** NodeEIPSpec

Parameter	Type	Description
iptype	String	EIP type, specified in <b>publicip.type</b> in the API for assigning an EIP.
bandwidth	<b>NodeBandwidth</b> object	Bandwidth parameters of the EIP

**Table 4-281** NodeBandwidth

Parameter	Type	Description
chargemode	String	<p>Bandwidth billing mode.</p> <ul style="list-style-type: none"> <li>• If this field is not specified, the billing is based on bandwidth.</li> <li>• If the field is null, the billing is based on bandwidth.</li> <li>• If the field value is <b>traffic</b>, the billing is based on traffic.</li> <li>• If the value is out of the preceding options, the cloud server will fail to be created.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• Billed by bandwidth: The billing will be based on the data transmission rate (in Mbps) of public networks. This billing mode is recommended if your bandwidth usage is higher than 10%.</li> <li>• Billed by traffic: The billing is based on the total amount of data (in GB) transmitted over the public network. This mode is available only when you are creating a pay-per-use node. This billing mode is recommended if your bandwidth usage is lower than 10%.</li> </ul>
size	Integer	Bandwidth size, specified in <b>bandwidth.size</b> in the API for assigning an EIP.
sharetype	String	Bandwidth sharing type. Value options: <b>PER</b> (exclusive bandwidth)

**Table 4-282** NodeNicSpec

Parameter	Type	Description
primaryNic	<b>NicSpec</b> object	Description of the primary NIC.
extNics	Array of <b>NicSpec</b> objects	<p>Extension NIC</p> <p><b>NOTE</b> This parameter is not supported when you add a node to a node pool.</p>

**Table 4-283 NicSpec**

Parameter	Type	Description
subnetId	String	Network ID of the subnet to which a NIC belongs. If <b>subnetId</b> is not specified when a primary NIC is creating, the cluster subnet will be used. If <b>subnetList</b> is also configured for a node pool, the <b>subnetList</b> field is used for adding subnets to the node pool. When creating an extension NIC, you must specify <b>subnetId</b> .
fixedIps	Array of strings	The IP address of the primary ENI is specified using <b>fixedIps</b> . The number of IP addresses cannot be greater than the number of created nodes. Either <b>fixedIps</b> or <b>ipBlock</b> can be specified. <b>fixedIps</b> cannot be specified for extension ENIs.
ipBlock	String	CIDR format of the primary NIC IP range. The IP address of the created node falls in this range. <b>fixedIps</b> and <b>ipBlock</b> cannot be specified at the same time.
subnetList	Array of strings	Network ID list of the subnet where the NIC resides. A maximum of 20 subnets can be configured for a node pool.

**Table 4-284 Taint**

Parameter	Type	Description
key	String	Key.
value	String	Value.
effect	String	Effect.

**Table 4-285 UserTag**

Parameter	Type	Description
key	String	Key of the cloud server label. The value cannot start with <b>CCE-</b> or <b>__type_baremetal</b> .
value	String	Value of the cloud server label.

**Table 4-286** Runtime

Parameter	Type	Description
name	String	<p>Container runtime. Default value:</p> <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>

**Table 4-287** NodeExtendParam

Parameter	Type	Description
ecs:performancetype	String	ECS flavor types. This field is returned in the response.
maxPods	Integer	<p>Maximum number of pods that can be created on a node, including the default system pods. Value range: 16 to 256.</p> <p>This limit prevents the node from being overloaded of pods.</p> <p>The number of pods that can be created on a node is determined by multiple parameters. For details, see <a href="#">Maximum Number of Pods That Can Be Created on a Node</a>.</p>

Parameter	Type	Description
DockerLVMConfigOverride	String	<p>Docker data disk configuration item. (This parameter has been discarded. Use the <b>storage</b> field instead.) Example default configuration: "DockerLVMConfigOverride":"dockerThinpool=vgpaas/90%VG;kubernetesLV=vgpaas/10%VG;diskType=evs;lvType=linear"</p> <p>By default, if no VD disk is available, an error occurs because the data disk fails to be found. Set <b>diskType</b> based on the actual drive letter type. The following fields are included:</p> <ul style="list-style-type: none"> <li>• <b>userLV</b> (optional): size of the user space, for example, <b>vgpaas/20%VG</b>.</li> <li>• <b>userPath</b> (optional): mount path of the user space, for example, <b>/home/wqt-test</b>.</li> <li>• <b>diskType</b>: disk type. Currently, only <b>evs</b>, <b>hdd</b>, and <b>ssd</b> are supported.</li> <li>• <b>lvType</b>: type of a logic volume. The value can be <b>linear</b> or <b>striped</b>.</li> <li>• <b>dockerThinpool</b>: Docker space size, for example, <b>vgpaas/60%VG</b>.</li> <li>• <b>kubernetesLV</b>: kubelet space size, for example, <b>vgpaas/20%VG</b>.</li> </ul>
dockerBaseSize	Integer	<p>Available disk space of a single container on a node, in GB. This parameter has been discarded. Use <b>containerBaseSize</b>.</p> <p>If this parameter is left blank or is set to <b>0</b>, the default value is used. In Device Mapper mode, the default value is <b>10</b>. In OverlayFS mode, the available space of a single container is not limited by default, and the <b>dockerBaseSize</b> setting takes effect only on nodes running EulerOS in clusters of the new versions.</p> <p>For details about how to allocate the space for the container runtime, see <a href="#">Data Disk Space Allocation</a>.</p> <p>When Device Mapper is used, it is recommended that <b>dockerBaseSize</b> be set to a value less than or equal to 80 GB. If the value is too large, the container runtime may fail to be started due to long initialization. If there are special requirements for the container disk space, you can mount an external or local storage device.</p>

Parameter	Type	Description
containerBaseSize	Integer	<p>Available disk space of a single container on a node, in GiB.</p> <p>If this parameter is left blank or is set to <b>0</b>, the default value is used. In OverlayFS mode, the available space of a single container is not limited by default. In Device Mapper mode, the default value is <b>10</b>, and the <b>containerBaseSize</b> setting takes effect only on nodes running EulerOS in clusters v1.23.14-r0, v1.25.9-r0, v1.27.6-r0, v1.28.4-r0 or later.</p> <p>For details about how to allocate the space for the container runtime, see <a href="#">Data Disk Space Allocation</a>.</p> <p>When Device Mapper is used, it is recommended that <b>containerBaseSize</b> be set to a value less than or equal to 80 GB. If the value is too large, the container runtime may fail to be started due to long initialization. If there are special requirements for the container disk space, you can mount an external or local storage device. In new versions, the Device Mapper mode is used only by BMSs in the same resource pool and is being discarded.</p> <p><b>NOTE</b> This parameter cannot be updated when a node pool is being updated.</p>
publicKey	String	Public key of a node.
alpha.cce/preInstall	String	<p>Pre-installation scripts. The total number of characters after transcoding cannot exceed 10240.</p> <p><b>NOTE</b> The input value must be encoded using Base64. The method is as follows:</p> <pre>echo -n "Content to be encoded"   base64</pre>
alpha.cce/postInstall	String	<p>Post-installation scripts. The total number of characters after transcoding cannot exceed 10240.</p> <p><b>NOTE</b> The input value must be encoded using Base64. The method is as follows:</p> <pre>echo -n "Content to be encoded"   base64</pre>
alpha.cce/NodeImageID	String	This parameter is required when a custom image is used to create a BMS node.



Parameter	Type	Description
chargingMode	Integer	Billing mode of a node. This parameter has been deprecated. Use the <b>billingMode</b> parameter in <b>NodeSpec</b> .
agency_name	String	Name of an agency An agency is created by a tenant administrator on Identity and Access Management (IAM) to provide temporary credentials for CCE nodes to access cloud servers. This parameter is returned only when it is transferred during node creation.
kubeReservedMem	Integer	Reserved node memory, which is reserved for Kubernetes components.
systemReservedMem	Integer	Reserved node memory, which is reserved for system components.
init-node-password	String	Node password. If this parameter is used as a response parameter, asterisks (*) will be displayed.

**Table 4-288** HostnameConfig

Parameter	Type	Description
type	String	Configuration type of the Kubernetes node name. The default value is <b>privatelp</b> . <ul style="list-style-type: none"> <li>• <b>privatelp</b>: The Kubernetes node is named after its IP address.</li> <li>• <b>cceNodeName</b>: The Kubernetes node is named after the CCE node.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• For a node which is configured using <b>cceNodeName</b>, the name is the same as the Kubernetes node name and the ECS name. The node name cannot be changed. If the ECS name is changed on the ECS console, the node name will retain unchanged after ECS synchronization.</li> <li>• For a node which is configured using <b>cceNodeName</b>, to avoid a conflict between Kubernetes nodes, the system automatically adds a suffix to each node name. The suffix is in the format of <b>A hyphen (-) Five random characters</b>. The value of the random characters is a lowercase letter or a digit ranging from 0 to 9.</li> </ul>

**Table 4-289** NodeStatus

Parameter	Type	Description
phase	String	Node status, which is the status of a node during its lifecycle (such as installation and uninstallation) or the status of a node in a Kubernetes cluster. Options: <ul style="list-style-type: none"> <li>• <b>Build</b>: The node is being created.</li> <li>• <b>Installing</b>: The node is being managed.</li> <li>• <b>Upgrading</b>: The node is being upgraded.</li> <li>• <b>Active</b>: The node is running properly.</li> <li>• <b>Abnormal</b>: The node is abnormal.</li> <li>• <b>Deleting</b>: The node is being deleted.</li> <li>• <b>Error</b>: The node is faulty.</li> </ul>
lastProbeTime	String	Last time when the node status was checked. If the cluster is in the abnormal, frozen, or intermediate state (for example, creating), the node status check may be affected. The node status that takes more than five minutes to check has no reference value.
jobID	String	ID of a creation or deletion job
serverId	String	ID of the underlying ECS or BMS node
privateIP	String	IP address in the private network segment of the primary NIC on the node
privateIPv6IP	String	IPv6 address in the private network segment of the primary NIC on the node
publicIP	String	Node EIP. If the ECS data is not synchronized in real time, you can click <b>Sync Node Data</b> on the console to manually update the data.
deleteStatus	<a href="#">DeleteStatus</a> object	Resource status during resource deletion.

**Table 4-290** DeleteStatus

Parameter	Type	Description
previous_total	Integer	Total number of existing cluster resource records when the cluster is deleted.
current_total	Integer	Latest number of resource records, which is generated based on the current cluster resource records.

Parameter	Type	Description
updated	Integer	Total number of resource records updated when the cluster is deleted.
added	Integer	Total number of resource records updated when the cluster is deleted.
deleted	Integer	Total number of resource records deleted when the cluster is deleted.

## Example Requests

None

## Example Responses

**Status code: 200**

Information about all the nodes in the specified cluster is successfully obtained.

```
{
  "kind": "List",
  "apiVersion": "v3",
  "items": [ {
    "kind": "Node",
    "apiVersion": "v3",
    "metadata": {
      "name": "myhost",
      "uid": "4d1ecb2c-229a-11e8-9c75-0255ac100ceb",
      "creationTimestamp": "2018-08-02 07:37:24.005071325 +0000 UTC",
      "updateTimestamp": "2018-08-02 07:44:04.965500815 +0000 UTC",
      "annotations": {
        "kubernetes.io/node-pool.id": "az1.dc1#s1.medium#EulerOS 2.2"
      }
    },
    "spec": {
      "flavor": "s1.medium",
      "az": "az1.dc1",
      "os": "EulerOS 2.2",
      "login": {
        "sshKey": "KeyPair-001"
      },
      "rootVolume": {
        "volumetype": "SAS",
        "size": 40
      },
      "dataVolumes": [ {
        "volumetype": "SAS",
        "size": 100
      } ],
      "publicIP": {
        "eip": {
          "bandwidth": { }
        }
      }
    },
    "billingMode": 0
  },
  "status": {
    "phase": "Active",
    "serverId": "456789abc-9368-46f3-8f29-d1a95622a568",
    "publicIP": "10.34.56.78",
  }
}
```

```
"privateIP" : "192.168.1.23"
}
}]
}
```

## Status Codes

Status Code	Description
200	Information about all the nodes in the specified cluster is successfully obtained.

## Error Codes

See [Error Codes](#).

### 4.3.4 Updating a Specified Node

#### Function

This API is used to update information about a specified node.

#### NOTE

- Currently, only the **name** field in **metadata** can be updated. This field indicates the node name.
- The URL for cluster management is in the format of **https://Endpoint/uri**. In the URL, **uri** indicates the resource path, that is, the path for API access.

#### URI

PUT /api/v3/projects/{project\_id}/clusters/{cluster\_id}/nodes/{node\_id}

**Table 4-291** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
node_id	Yes	String	Node ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

## Request Parameters

**Table 4-292** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

**Table 4-293** Request body parameters

Parameter	Mandatory	Type	Description
metadata	Yes	<a href="#">ClusterNodeInformationMetadata</a> object	Node metadata, which is a collection of attributes.

**Table 4-294** ClusterNodeInformationMetadata

Parameter	Mandatory	Type	Description
name	Yes	String	Node name <b>NOTE</b> After the node name is changed, the ECS name (VM name) will be changed accordingly. Enter 1 to 56 characters starting with a lowercase letter and not ending with a hyphen (-). Only lowercase letters, digits, hyphens (-), and periods (.) are allowed.

## Response Parameters

Status code: 200

**Table 4-295** Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at <b>Node</b> and cannot be changed.
apiVersion	String	API version. The value is fixed at <b>v3</b> and cannot be changed.
metadata	<b>NodeMetadata</b> object	Node metadata, which is a collection of attributes.
spec	<b>NodeSpec</b> object	Detailed description of the node. CCE creates or updates objects by defining or updating <b>spec</b> .
status	<b>NodeStatus</b> object	Node status, which is dynamically recorded. A user-defined value will not function when a node is being created or modified.

**Table 4-296** NodeMetadata

Parameter	Type	Description
name	String	Node name. <b>NOTE</b> Enter 1 to 56 characters starting with a lowercase letter and not ending with a hyphen (-). Only lowercase letters, digits, hyphens (-), and periods (.) are allowed. If <b>name</b> is not configured or is left blank, the node name is generated based on the default rule, which is <i>Cluster name-Random characters</i> . Only the first 36 characters are kept for a long cluster name. If the number of nodes ( <b>count</b> ) is greater than <b>1</b> , some random characters are added to the end of the entered node name according to the default rule, which is <i>User-defined name-Random characters</i> . If the length of the entered node name exceeds 50 characters, the system uses the first 50 characters and adds a random character string to the end of the name.
uid	String	Node ID, which is unique and automatically generated after the resource is created. A user-defined ID will not take effect.
labels	Map<String,String>	CCE node label (not the native Kubernetes label). Labels are used to select objects that meet certain criteria. A label is a key-value pair. Example: <pre>"labels": {   "key": "value" }</pre>

Parameter	Type	Description
annotations	Map<String,String>	<p>CCE node annotations, in key-value pairs (not the native Kubernetes annotations). Example:</p> <pre>"annotations": {   "key1" : "value1",   "key2" : "value2" }</pre> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• <b>annotations</b> are not used to identify or select objects. The metadata in <b>annotations</b> may be small or large, structured or unstructured, and may include characters that are not allowed in labels.</li> <li>• This parameter is used only for query and cannot be input through a request. Entered data of this parameter is invalid.</li> </ul>
creationTimes tamp	String	Time when the object was created. The value is automatically generated after the object is created. A user-defined value will not take effect.
updateTimest amp	String	Time when the object was updated. The value is automatically generated after the object is created. A user-defined value will not take effect.
ownerReferen ce	<a href="#">ownerRefere nce</a> object	Owner object

**Table 4-297** ownerReference

Parameter	Type	Description
nodepoolName	String	Node pool name
nodepoolID	String	Node pool UID

**Table 4-298** NodeSpec

Parameter	Type	Description
flavor	String	Node specifications. For details about the node specifications supported by CCE, see <a href="#">Node Flavor Description</a> .

Parameter	Type	Description
az	String	AZ where the node to be created is located. You need to specify the AZ name. A node created using APIs does not support random AZs. <b>NOTE</b> <ul style="list-style-type: none"> <li>When you are creating a node pool and configuring a scaling group, this parameter cannot be set to <b>random</b>.</li> </ul>
os	String	Node OS. For details about the supported OSs, see <a href="#">Node OS</a> . <b>NOTE</b> <ul style="list-style-type: none"> <li>The system automatically selects the supported OS based on the cluster version. If the current cluster version does not support the OS, an error will be reported.</li> <li>If <b>alpha.cce/NodeImageID</b> in <b>extendParam</b> is specified during node creation, you do not need to configure this parameter.</li> <li>This parameter is mandatory when creating a node pool.</li> <li>If you are using shared disk space when creating a node, this parameter is mandatory. Specifically, the disk initialization configuration management parameter <b>storage</b> is used and the <b>name</b> field of <b>virtualSpaces</b> in <b>StorageGroups</b> is set to <b>share</b>.</li> </ul>
login	<a href="#">Login</a> object	Node login mode. Either the key pair or password must be used for login. <b>NOTE</b> This parameter cannot be modified when a node pool is being updated.
rootVolume	<a href="#">Volume</a> object	Information about disks on the node
dataVolumes	Array of <a href="#">Volume</a> objects	Data disk parameters of a node. A VM can have up to 16 disks attached, while a BMS can only have up to 10 disks attached. The maximum number of disks that can be attached to either a VM or BMS is limited by their respective maximum disk limits. (You can add more data disks to each of your nodes through the CCE console and APIs.) The data disk used by the container runtime and kubelet cannot be detached, or the node will become unavailable. For DeC nodes, the parameter details are the same as those for <b>rootVolume</b> .



Parameter	Type	Description
storage	<a href="#">Storage</a> object	<p>Disk initialization management parameter. This parameter is complex to configure. For details, see <a href="#">Attaching Disks to a Node</a>.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If a node has both local and EVS disks attached, make sure to not keep the default value for this parameter, or it may result in unexpected disk partitions.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If you want to change the value range of a data disk to 20 to 32768, do not use the default value.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If you want to use the shared disk space (with the runtime and Kubernetes partitions cancelled), do not retain the default value of this parameter. For details about the shared disk space, see <a href="#">Data Disk Space Allocation</a>.</li> </ul>
publicIP	<a href="#">NodePublicIP</a> object	<p>EIP of a node.</p> <p><b>NOTE</b></p> <p>This parameter is not supported when you add a node to a node pool.</p>
nodeNicSpec	<a href="#">NodeNicSpec</a> object	NIC of the node
count	Integer	Number of nodes to be created in a batch. The value must be a positive integer greater than or equal to 1 and less than or equal to the defined limit. This parameter can be left blank when it is used for a node pool.
billingMode	Integer	<p>Node billing mode.</p> <ul style="list-style-type: none"> <li><b>0</b>: pay-per-use</li> </ul>

Parameter	Type	Description
taints	Array of <b>Taint</b> objects	<p>You can add taints to created nodes to configure anti-affinity. A maximum of 20 taints can be added. Each taint contains the following parameters:</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain can be prefixed to a key.</li> <li>• <b>Value:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed.</li> <li>• <b>Effect:</b> Available options are <b>NoSchedule</b>, <b>PreferNoSchedule</b>, and <b>NoExecute</b>. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</li> </ul> <p>Example:</p> <pre data-bbox="818 1111 1430 1339">"taints": [{   "key": "status",   "value": "unavailable",   "effect": "NoSchedule" }, {   "key": "looks",   "value": "bad",   "effect": "NoSchedule" }]</pre>

Parameter	Type	Description
k8sTags	Map<String,String>	<p>Defined in key-value pairs. A maximum of 20 key-value pairs are allowed.</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain containing a maximum of 253 characters (such as <b>example.com/my-key</b>) can be prefixed to a key.</li> <li>• <b>Value:</b> The value can be left blank or contain 1 to 63 characters that start and end with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</li> </ul> <p>Example:</p> <pre>"k8sTags": {   "key": "value" }</pre>
ecsGroupId	String	<p>ECS group ID. If this parameter is configured, nodes will be created in the specified ECS group.</p> <p><b>NOTE</b> This configuration does not take effect when a node pool is created. To ensure that all nodes in a node pool are in the same ECS group, configure <b>nodeManagement</b> in the node pool.</p>
dedicatedHostId	String	<p>ID of the DeH to which the node is scheduled.</p> <p><b>NOTE</b> This parameter is not supported when you add a node during node pool creation.</p>

Parameter	Type	Description
userTags	Array of <a href="#">UserTag</a> objects	<p>Cloud server tag (resource tag). The key of a tag must be unique. The maximum number of custom tags supported by CCE varies depending on regions and cannot exceed 8. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</p> <p><b>NOTE</b> A tag key can have a maximum of 36 characters. It can only contain letters, digits, Unicode characters, and special characters (-_).</p>
runtime	<a href="#">Runtime</a> object	<p>Container runtime. Default value:</p> <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>

Parameter	Type	Description
initializedConditions	Array of strings	<p>Custom initialization flag, which is left blank by default.</p> <p>Before CCE nodes are initialized, <b>node.cloudprovider.kubernetes.io/uninitialized</b> is added to the nodes to prevent pods from being scheduled to these nodes. When creating a node, you can configure the <b>initializedConditions</b> parameter to control the taint removal time. By default, the timeout period is not configured.</p> <p>Example:</p> <ol style="list-style-type: none"> <li>1. Create a node and add parameter "initializedConditions": ["CCEInitial", "CustomedInitial"].</li> <li>2. After custom initialization is complete, CCE will call a Kubernetes API (for example, <b>PATCH /v1/nodes/{node_ip}/status</b>) to update the node <b>conditions</b> by adding two labels of types <b>CCEInitial</b> and <b>CustomedInitial</b> and setting the status to <b>True</b>, as shown in the following:</li> </ol> <pre data-bbox="815 1055 1430 1205"> status:   conditions:   - type: CCEInitial     status: 'True'   - type: CustomedInitial     status: 'True'                     </pre> <ol style="list-style-type: none"> <li>1. CCE polls <b>status.Conditions</b> of nodes to check whether there are conditions of types <b>CCEInitial</b> and <b>CustomedInitial</b>. If such conditions exist and the status is <b>True</b>, the node initialization is complete and the initialization taint is removed.</li> <li>2. <b>initializedConditions</b> allows you to configure a timeout period during node creation. For example, "initializedConditions": ["CCEInitial:15m", "CustomedInitial:15m"] indicates that the timeout period is 15 minutes. After the timeout period is reached, the initialization conditions will be automatically ignored and the initialization taint will be removed when CCE polls the node. <ul style="list-style-type: none"> <li>● The timeout interval in <b>initializedConditions</b> ranges from 1 to 99 seconds.</li> <li>● Use only letters and digits. Max. characters: 20.</li> <li>● Max. flags: 2.</li> </ul> </li> </ol>

Parameter	Type	Description
		<ul style="list-style-type: none"> <li>The unit of the timeout period is minute (m).</li> </ul>
extendParam	<b>NodeExtendParam</b> object	Extended parameters for creating a node.
hostnameConfig	<b>HostnameConfig</b> object	Kubernetes node name configuration parameter, which is supported by clusters of v1.23.6-r0 to v1.25 or clusters of v1.25.2-r0 or later versions.
serverEnterpriseProjectID	String	Enterprise project ID of a server. CCE does not support features provided by EPS. This field is used only for synchronizing enterprise project IDs of servers. To create a node or node pool, you can specify an existing enterprise project for it. If this field is left blank, it inherits the enterprise project attributes of the cluster that runs the node or node pool. When updating a node pool, any configuration modifications you made will apply only to new nodes. You need to modify the enterprise projects of existing nodes on the EPS console. If no value is specified during the update, this field will not be updated. If this parameter is left blank, the cluster enterprise project is returned.

**Table 4-299** Login

Parameter	Type	Description
sshKey	String	Name of the key pair used for login.
userPassword	<b>UserPassword</b> object	Password used for node login.

**Table 4-300** UserPassword

Parameter	Type	Description
username	String	Login account. The default value is <b>root</b> .

Parameter	Type	Description
password	String	<p>Login password. If a username and password are used to create a node, this field is shielded in the response body. A password must meet the following complexity requirements:</p> <ul style="list-style-type: none"> <li>• Contains 8 to 26 characters.</li> <li>• Contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters !@\$%^-_=+[{ }];:./?</li> <li>• Cannot contain the username or the username spelled backwards. The <b>password</b> field must be salted during node creation. For details, see <a href="#">Adding a Salt in the password Field When Creating a Node</a>.</li> </ul>

**Table 4-301** Volume

Parameter	Type	Description
size	Integer	<p>Disk size, in GiB.</p> <ul style="list-style-type: none"> <li>• System disk: 40 to 1024</li> <li>• Value range for data disks: 100 to 32768</li> </ul>
volumetype	String	<p>Disk type. For details about possible values, see the description of the <b>root_volume</b> parameter in the API used to create an ECS.</p> <ul style="list-style-type: none"> <li>• <b>SAS</b>: high I/O SAS disk</li> <li>• <b>SSD</b>: ultra-high I/O SSD disk</li> <li>• <b>SATA</b>: common I/O SATA disk SATA disks have been removed from EVS. You can find them attached only on existing nodes.</li> </ul>
extendParam	Map<String, Object>	<p>Extended disk parameters, defined in <b>extendparam</b> in the API used to create an ECS.</p>
cluster_id	String	<p>ID of the storage pool used by the ECS system disk. This field is used only for DeC clusters, which functions as <b>dssPoolID</b>, that is, the ID of the DSS storage pool.</p>
cluster_type	String	<p>Storage class of the cloud server system disk. The value is always <b>dss</b>. This field is used only for DeC clusters.</p>

Parameter	Type	Description
hw:passthrough	Boolean	<ul style="list-style-type: none"> <li>Pay attention to this field if your ECS is SDI-compliant. If the value of this field is <b>true</b>, the created disk is of SCSI type.</li> <li>If the node pool type is <b>ElasticBMS</b>, this field must be set to <b>true</b>.</li> <li>If a node specification involves both local and EVS disks, configure the disk initialization parameters. For details, see <a href="#">Attaching Disks to a Node</a>.</li> </ul>
metadata	<b>VolumeMetadata</b> object	EVS disk encryption information. This field is mandatory only when you need to encrypt the system disk or data disks of the node to be created.

**Table 4-302** VolumeMetadata

Parameter	Type	Description
__system__encrypted	String	Whether the EVS disk is encrypted. The value <b>0</b> indicates that the EVS disk is not encrypted, and the value <b>1</b> indicates that the EVS disk is encrypted.  If this parameter is not specified, EVS disks will not be encrypted by default.
__system__cmkid	String	CMK ID, which indicates encryption in <b>metadata</b> . This field is used with <b>__system__encrypted</b> .

**Table 4-303** Storage

Parameter	Type	Description
storageSelectors	Array of <b>StorageSelectors</b> objects	Disk selection. Matched disks are managed according to <b>matchLabels</b> and <b>storageType</b> .
storageGroups	Array of <b>StorageGroups</b> objects	A storage group consists of multiple storage devices. It is used to divide storage space.



**Table 4-304** StorageSelectors

Parameter	Type	Description
name	String	Selector name, used as the index of <b>selectorNames</b> in <b>storageGroup</b> . Therefore, the name of each selector must be unique.
storageType	String	Storage type. Only <b>evs</b> (EVS disks) and <b>local</b> (local disks) are supported. The local storage does not support disk selection. All local disks will form a VG, so only one storageSelector of the local type is allowed.
matchLabels	<b>matchLabels</b> object	Matching field of an EVS disk. The value can be <b>size</b> , <b>volumeType</b> , <b>metadataEncrypted</b> , <b>metadataCmkid</b> , or <b>count</b> in <b>DataVolume</b> .

**Table 4-305** matchLabels

Parameter	Type	Description
size	String	Matched disk size. If this parameter is left unspecified, the disk size is not limited. Example: 100
volumeType	String	EVS disk type. <b>SSD</b> , <b>GPSSD</b> , <b>SAS</b> , <b>ESSD</b> , and <b>SATA</b> are supported. If this parameter is left blank, there is no restriction on the disk type.
metadataEncrypted	String	Disk encryption identifier. If this parameter is set to <b>0</b> , it specifies that the disk is not encrypted. If this parameter is set to <b>1</b> , it specifies that the disk is encrypted. If this parameter is left blank, there is no restriction on the disk encryption identifier.
metadataCmkid	String	CMK ID of an encrypted disk. The value is a string of 36 bytes. If this parameter is left blank, there is no restriction on the disk key ID.
count	String	Number of disks to be selected. If this parameter is left blank, all disks of this type are selected.

**Table 4-306** StorageGroups

Parameter	Type	Description
name	String	Name of a virtual storage group, which must be unique. <b>NOTE</b> <ul style="list-style-type: none"> <li>If <b>cceManaged</b> is set to <b>true</b>, the name must be set to <b>vgpaas</b>.</li> <li>If the data disk is used as a temporary storage volume, the name must be <b>vg-everest-localvolume-ephemeral</b>.</li> <li>If the data disk is used as a persistent storage volume, the name must be <b>vg-everest-localvolume-persistent</b>.</li> </ul>
cceManaged	Boolean	Storage space for Kubernetes and runtime components. Only one group can be set to <b>true</b> . If this parameter is left blank, the default value <b>false</b> is used.
selectorNames	Array of strings	Corresponds to <b>name</b> in <b>storageSelectors</b> . One group can select multiple selectors, but one selector can be selected by only one group.
virtualSpaces	Array of <a href="#">VirtualSpace</a> objects	Detailed management of space configuration in a group.

**Table 4-307** VirtualSpace

Parameter	Type	Description
name	String	Name of a virtualSpace. Options: <b>share</b> , <b>kubernetes</b> , <b>runtime</b> , and <b>user</b> <ul style="list-style-type: none"> <li><b>share</b>: shared space configuration (with the runtime and Kubernetes partitions cancelled). <b>lvmConfig</b> needs to be configured.</li> <li><b>Kubernetes</b>: Kubernetes space configuration. <b>lvmConfig</b> needs to be configured.</li> <li><b>runtime</b>: runtime space configuration. <b>runtimeConfig</b> needs to be configured.</li> <li><b>user</b>: user space configuration. <b>lvmConfig</b> needs to be configured.</li> </ul>

Parameter	Type	Description
size	String	Size of a virtualSpace. The value must be an integer in percentage. Example: 90%. <b>NOTE</b> The value ranges from 10 to 90. The sum of percentages of all virtualSpaces in a group cannot exceed 100%.
lvmConfig	<b>LVMConfig</b> object	LVM configuration managements, applicable to the <i>share*</i> , <b>kubernetes</b> , and <b>user</b> spaces. Note that one virtual space supports only one config.
runtimeConfig	<b>RuntimeConfig</b> object	runtime configurations, applicable to the <b>runtime</b> space. Note that one virtual space supports only one config.

**Table 4-308** LVMConfig

Parameter	Type	Description
lvType	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.
path	String	Disk mount path. This parameter takes effect only in user configuration. The value is an absolute path. Digits, letters, periods (.), hyphens (-), and underscores (_) are allowed. <b>NOTE</b> Do not leave the path empty or set it to a key OS path such as the root directory. Key OS paths that cannot be specified: /, /home, /home/, /bin, /bin/, /lib, /lib/, /root, /root/, /boot, /boot/, /dev, /dev/, /etc, /etc/, /lost+found, /lost+found/, /mnt, /mnt/, /proc, /proc/, /sbin, /sbin/, /srv, /srv/, /tmp, /tmp/, /var, /var/, /media, /media/, /opt, /opt/, /selinux, /selinux/, /sys, /sys/, /usr, /usr/, /opt/cloud/, /mnt/paas/, /home/paas/, /var/paas/, /var/lib/, /var/script/

**Table 4-309** RuntimeConfig

Parameter	Type	Description
lvType	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.

**Table 4-310** NodePublicIP

Parameter	Type	Description
ids	Array of strings	IDs of existing EIPs. The quantity cannot be greater than the number of nodes to be created. <b>NOTE</b> If <b>ids</b> has been set, you do not need to set <b>count</b> and <b>eip</b> .
count	Integer	Number of EIPs to be dynamically created. <b>NOTE</b> <b>count</b> and <b>eip</b> must be set at the same time.
eip	<b>NodeEIPSpec</b> object	EIP configuration. <b>NOTE</b> This parameter is not supported when creating a node pool.

**Table 4-311** NodeEIPSpec

Parameter	Type	Description
iptype	String	EIP type, specified in <b>publicip.type</b> in the API for assigning an EIP.
bandwidth	<b>NodeBandwidth</b> object	Bandwidth parameters of the EIP

**Table 4-312** NodeBandwidth

Parameter	Type	Description
chargemode	String	<p>Bandwidth billing mode.</p> <ul style="list-style-type: none"> <li>• If this field is not specified, the billing is based on bandwidth.</li> <li>• If the field is null, the billing is based on bandwidth.</li> <li>• If the field value is <b>traffic</b>, the billing is based on traffic.</li> <li>• If the value is out of the preceding options, the cloud server will fail to be created.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• Billed by bandwidth: The billing will be based on the data transmission rate (in Mbps) of public networks. This billing mode is recommended if your bandwidth usage is higher than 10%.</li> <li>• Billed by traffic: The billing is based on the total amount of data (in GB) transmitted over the public network. This mode is available only when you are creating a pay-per-use node. This billing mode is recommended if your bandwidth usage is lower than 10%.</li> </ul>
size	Integer	Bandwidth size, specified in <b>bandwidth.size</b> in the API for assigning an EIP.
sharetype	String	Bandwidth sharing type. Value options: <b>PER</b> (exclusive bandwidth)

**Table 4-313** NodeNicSpec

Parameter	Type	Description
primaryNic	<b>NicSpec</b> object	Description of the primary NIC.
extNics	Array of <b>NicSpec</b> objects	<p>Extension NIC</p> <p><b>NOTE</b> This parameter is not supported when you add a node to a node pool.</p>

**Table 4-314 NicSpec**

Parameter	Type	Description
subnetId	String	Network ID of the subnet to which a NIC belongs. If <b>subnetId</b> is not specified when a primary NIC is creating, the cluster subnet will be used. If <b>subnetList</b> is also configured for a node pool, the <b>subnetList</b> field is used for adding subnets to the node pool. When creating an extension NIC, you must specify <b>subnetId</b> .
fixedIps	Array of strings	The IP address of the primary ENI is specified using <b>fixedIps</b> . The number of IP addresses cannot be greater than the number of created nodes. Either <b>fixedIps</b> or <b>ipBlock</b> can be specified. <b>fixedIps</b> cannot be specified for extension ENIs.
ipBlock	String	CIDR format of the primary NIC IP range. The IP address of the created node falls in this range. <b>fixedIps</b> and <b>ipBlock</b> cannot be specified at the same time.
subnetList	Array of strings	Network ID list of the subnet where the NIC resides. A maximum of 20 subnets can be configured for a node pool.

**Table 4-315 Taint**

Parameter	Type	Description
key	String	Key.
value	String	Value.
effect	String	Effect.

**Table 4-316 UserTag**

Parameter	Type	Description
key	String	Key of the cloud server label. The value cannot start with <b>CCE-</b> or <b>__type_baremetal</b> .
value	String	Value of the cloud server label.

**Table 4-317 Runtime**

Parameter	Type	Description
name	String	<p>Container runtime. Default value:</p> <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>

**Table 4-318 NodeExtendParam**

Parameter	Type	Description
ecs:performancetype	String	ECS flavor types. This field is returned in the response.
maxPods	Integer	<p>Maximum number of pods that can be created on a node, including the default system pods. Value range: 16 to 256.</p> <p>This limit prevents the node from being overloaded of pods.</p> <p>The number of pods that can be created on a node is determined by multiple parameters. For details, see <a href="#">Maximum Number of Pods That Can Be Created on a Node</a>.</p>

Parameter	Type	Description
DockerLVMConfigOverride	String	<p>Docker data disk configuration item. (This parameter has been discarded. Use the <b>storage</b> field instead.) Example default configuration: "DockerLVMConfigOverride":"dockerThinpool=vgpaas/90%VG;kubernetesLV=vgpaas/10%VG;diskType=evs;lvType=linear"</p> <p>By default, if no VD disk is available, an error occurs because the data disk fails to be found. Set <b>diskType</b> based on the actual drive letter type. The following fields are included:</p> <ul style="list-style-type: none"> <li>• <b>userLV</b> (optional): size of the user space, for example, <b>vgpaas/20%VG</b>.</li> <li>• <b>userPath</b> (optional): mount path of the user space, for example, <b>/home/wqt-test</b>.</li> <li>• <b>diskType</b>: disk type. Currently, only <b>evs</b>, <b>hdd</b>, and <b>ssd</b> are supported.</li> <li>• <b>lvType</b>: type of a logic volume. The value can be <b>linear</b> or <b>striped</b>.</li> <li>• <b>dockerThinpool</b>: Docker space size, for example, <b>vgpaas/60%VG</b>.</li> <li>• <b>kubernetesLV</b>: kubelet space size, for example, <b>vgpaas/20%VG</b>.</li> </ul>
dockerBaseSize	Integer	<p>Available disk space of a single container on a node, in GB. This parameter has been discarded. Use <b>containerBaseSize</b>.</p> <p>If this parameter is left blank or is set to <b>0</b>, the default value is used. In Device Mapper mode, the default value is <b>10</b>. In OverlayFS mode, the available space of a single container is not limited by default, and the <b>dockerBaseSize</b> setting takes effect only on nodes running EulerOS in clusters of the new versions.</p> <p>For details about how to allocate the space for the container runtime, see <a href="#">Data Disk Space Allocation</a>.</p> <p>When Device Mapper is used, it is recommended that <b>dockerBaseSize</b> be set to a value less than or equal to 80 GB. If the value is too large, the container runtime may fail to be started due to long initialization. If there are special requirements for the container disk space, you can mount an external or local storage device.</p>



Parameter	Type	Description
containerBaseSize	Integer	<p>Available disk space of a single container on a node, in GiB.</p> <p>If this parameter is left blank or is set to <b>0</b>, the default value is used. In OverlayFS mode, the available space of a single container is not limited by default. In Device Mapper mode, the default value is <b>10</b>, and the <b>containerBaseSize</b> setting takes effect only on nodes running EulerOS in clusters v1.23.14-r0, v1.25.9-r0, v1.27.6-r0, v1.28.4-r0 or later.</p> <p>For details about how to allocate the space for the container runtime, see <a href="#">Data Disk Space Allocation</a>.</p> <p>When Device Mapper is used, it is recommended that <b>containerBaseSize</b> be set to a value less than or equal to 80 GB. If the value is too large, the container runtime may fail to be started due to long initialization. If there are special requirements for the container disk space, you can mount an external or local storage device. In new versions, the Device Mapper mode is used only by BMSs in the same resource pool and is being discarded.</p> <p><b>NOTE</b> This parameter cannot be updated when a node pool is being updated.</p>
publicKey	String	Public key of a node.
alpha.cce/preInstall	String	<p>Pre-installation scripts. The total number of characters after transcoding cannot exceed 10240.</p> <p><b>NOTE</b> The input value must be encoded using Base64. The method is as follows:</p> <pre>echo -n "Content to be encoded"   base64</pre>
alpha.cce/postInstall	String	<p>Post-installation scripts. The total number of characters after transcoding cannot exceed 10240.</p> <p><b>NOTE</b> The input value must be encoded using Base64. The method is as follows:</p> <pre>echo -n "Content to be encoded"   base64</pre>
alpha.cce/NodeImageID	String	This parameter is required when a custom image is used to create a BMS node.

Parameter	Type	Description
chargingMode	Integer	Billing mode of a node. This parameter has been deprecated. Use the <b>billingMode</b> parameter in <b>NodeSpec</b> .
agency_name	String	Name of an agency An agency is created by a tenant administrator on Identity and Access Management (IAM) to provide temporary credentials for CCE nodes to access cloud servers. This parameter is returned only when it is transferred during node creation.
kubeReservedMem	Integer	Reserved node memory, which is reserved for Kubernetes components.
systemReservedMem	Integer	Reserved node memory, which is reserved for system components.
init-node-password	String	Node password. If this parameter is used as a response parameter, asterisks (*) will be displayed.

**Table 4-319** HostnameConfig

Parameter	Type	Description
type	String	Configuration type of the Kubernetes node name. The default value is <b>privatelp</b> . <ul style="list-style-type: none"> <li>• <b>privatelp</b>: The Kubernetes node is named after its IP address.</li> <li>• <b>cceNodeName</b>: The Kubernetes node is named after the CCE node.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• For a node which is configured using <b>cceNodeName</b>, the name is the same as the Kubernetes node name and the ECS name. The node name cannot be changed. If the ECS name is changed on the ECS console, the node name will retain unchanged after ECS synchronization.</li> <li>• For a node which is configured using <b>cceNodeName</b>, to avoid a conflict between Kubernetes nodes, the system automatically adds a suffix to each node name. The suffix is in the format of <b>A hyphen (-) Five random characters</b>. The value of the random characters is a lowercase letter or a digit ranging from 0 to 9.</li> </ul>

**Table 4-320** NodeStatus

Parameter	Type	Description
phase	String	Node status, which is the status of a node during its lifecycle (such as installation and uninstallation) or the status of a node in a Kubernetes cluster. Options: <ul style="list-style-type: none"> <li>● <b>Build</b>: The node is being created.</li> <li>● <b>Installing</b>: The node is being managed.</li> <li>● <b>Upgrading</b>: The node is being upgraded.</li> <li>● <b>Active</b>: The node is running properly.</li> <li>● <b>Abnormal</b>: The node is abnormal.</li> <li>● <b>Deleting</b>: The node is being deleted.</li> <li>● <b>Error</b>: The node is faulty.</li> </ul>
lastProbeTime	String	Last time when the node status was checked. If the cluster is in the abnormal, frozen, or intermediate state (for example, creating), the node status check may be affected. The node status that takes more than five minutes to check has no reference value.
jobID	String	ID of a creation or deletion job
serverId	String	ID of the underlying ECS or BMS node
privateIP	String	IP address in the private network segment of the primary NIC on the node
privateIPv6IP	String	IPv6 address in the private network segment of the primary NIC on the node
publicIP	String	Node EIP. If the ECS data is not synchronized in real time, you can click <b>Sync Node Data</b> on the console to manually update the data.
deleteStatus	<a href="#">DeleteStatus</a> object	Resource status during resource deletion.

**Table 4-321** DeleteStatus

Parameter	Type	Description
previous_total	Integer	Total number of existing cluster resource records when the cluster is deleted.
current_total	Integer	Latest number of resource records, which is generated based on the current cluster resource records.

Parameter	Type	Description
updated	Integer	Total number of resource records updated when the cluster is deleted.
added	Integer	Total number of resource records updated when the cluster is deleted.
deleted	Integer	Total number of resource records deleted when the cluster is deleted.

## Example Requests

Update the name of a specified node.

```
{
  "metadata": {
    "name": "new-hostname"
  }
}
```

## Example Responses

**Status code: 200**

Information about the specified node is successfully updated.

```
{
  "kind": "Node",
  "apiVersion": "v3",
  "metadata": {
    "name": "new-hostname",
    "uid": "4d1ecb2c-229a-11e8-9c75-0255ac100ceb",
    "creationTimestamp": "2017-08-20T21:11:09Z",
    "updateTimestamp": "2017-08-20T21:11:09Z",
    "annotations": {
      "kubernetes.io/node-pool.id": "az1.dc1#s1.medium#EulerOS 2.2"
    }
  },
  "spec": {
    "flavor": "s1.medium",
    "az": "az1.dc1",
    "os": "EulerOS 2.2",
    "login": {
      "sshKey": "KeyPair-001"
    },
    "rootVolume": {
      "volumetype": "SAS",
      "size": 40
    },
    "dataVolumes": [ {
      "volumetype": "SAS",
      "size": 100
    } ],
    "publicIP": {
      "eip": { }
    },
    "billingMode": 0
  },
  "status": {
    "phase": "Active",
    "serverId": "456789abc-9368-46f3-8f29-d1a95622a568",
    "publicIP": "10.34.56.78",
  }
}
```

```
"privateIP" : "192.168.1.23"
}
}
```

## Status Codes

Status Code	Description
200	Information about the specified node is successfully updated.

## Error Codes

See [Error Codes](#).

### 4.3.5 Deleting a Node

#### Function

This API is used to delete a specified node.

#### NOTE

The URL for cluster management is in the format of `https://Endpoint/uri`. In the URL, **uri** indicates the resource path, that is, the path for API access.

#### URI

DELETE `/api/v3/projects/{project_id}/clusters/{cluster_id}/nodes/{node_id}`

**Table 4-322** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
node_id	Yes	String	Node ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

**Table 4-323** Query Parameters

Parameter	Mandatory	Type	Description
nodepoolScaleDown	No	String	Whether the request is delivered by the node pool. If the value is not <b>NoScaleDown</b> , the number of pods in the corresponding node pool is automatically updated.

## Request Parameters

**Table 4-324** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

## Response Parameters

**Status code: 200**

**Table 4-325** Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at <b>Node</b> and cannot be changed.
apiVersion	String	API version. The value is fixed at <b>v3</b> and cannot be changed.
metadata	<a href="#">NodeMetadata</a> object	Node metadata, which is a collection of attributes.
spec	<a href="#">NodeSpec</a> object	Detailed description of the node. CCE creates or updates objects by defining or updating <b>spec</b> .

Parameter	Type	Description
status	<b>NodeStatus</b> object	Node status, which is dynamically recorded. A user-defined value will not function when a node is being created or modified.

**Table 4-326** NodeMetadata

Parameter	Type	Description
name	String	Node name. <b>NOTE</b> Enter 1 to 56 characters starting with a lowercase letter and not ending with a hyphen (-). Only lowercase letters, digits, hyphens (-), and periods (.) are allowed. If <b>name</b> is not configured or is left blank, the node name is generated based on the default rule, which is <i>Cluster name-Random characters</i> . Only the first 36 characters are kept for a long cluster name. If the number of nodes ( <b>count</b> ) is greater than <b>1</b> , some random characters are added to the end of the entered node name according to the default rule, which is <i>User-defined name-Random characters</i> . If the length of the entered node name exceeds 50 characters, the system uses the first 50 characters and adds a random character string to the end of the name.
uid	String	Node ID, which is unique and automatically generated after the resource is created. A user-defined ID will not take effect.
labels	Map<String,String>	CCE node label (not the native Kubernetes label). Labels are used to select objects that meet certain criteria. A label is a key-value pair. Example: <pre>"labels": {   "key": "value" }</pre>

Parameter	Type	Description
annotations	Map<String,String>	<p>CCE node annotations, in key-value pairs (not the native Kubernetes annotations). Example:</p> <pre>"annotations": {   "key1" : "value1",   "key2" : "value2" }</pre> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• <b>annotations</b> are not used to identify or select objects. The metadata in <b>annotations</b> may be small or large, structured or unstructured, and may include characters that are not allowed in labels.</li> <li>• This parameter is used only for query and cannot be input through a request. Entered data of this parameter is invalid.</li> </ul>
creationTimes tamp	String	Time when the object was created. The value is automatically generated after the object is created. A user-defined value will not take effect.
updateTimest amp	String	Time when the object was updated. The value is automatically generated after the object is created. A user-defined value will not take effect.
ownerReferen ce	<a href="#">ownerRefere nce</a> object	Owner object

**Table 4-327** ownerReference

Parameter	Type	Description
nodepoolName	String	Node pool name
nodepoolID	String	Node pool UID

**Table 4-328** NodeSpec

Parameter	Type	Description
flavor	String	Node specifications. For details about the node specifications supported by CCE, see <a href="#">Node Flavor Description</a> .



Parameter	Type	Description
az	String	<p>AZ where the node to be created is located. You need to specify the AZ name. A node created using APIs does not support random AZs.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>When you are creating a node pool and configuring a scaling group, this parameter cannot be set to <b>random</b>.</li> </ul>
os	String	<p>Node OS. For details about the supported OSs, see <a href="#">Node OS</a>.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>The system automatically selects the supported OS based on the cluster version. If the current cluster version does not support the OS, an error will be reported.</li> <li>If <b>alpha.cce/NodeImageID</b> in <b>extendParam</b> is specified during node creation, you do not need to configure this parameter.</li> <li>This parameter is mandatory when creating a node pool.</li> <li>If you are using shared disk space when creating a node, this parameter is mandatory. Specifically, the disk initialization configuration management parameter <b>storage</b> is used and the <b>name</b> field of <b>virtualSpaces</b> in <b>StorageGroups</b> is set to <b>share</b>.</li> </ul>
login	<a href="#">Login</a> object	<p>Node login mode. Either the key pair or password must be used for login.</p> <p><b>NOTE</b> This parameter cannot be modified when a node pool is being updated.</p>
rootVolume	<a href="#">Volume</a> object	Information about disks on the node
dataVolumes	Array of <a href="#">Volume</a> objects	<p>Data disk parameters of a node. A VM can have up to 16 disks attached, while a BMS can only have up to 10 disks attached. The maximum number of disks that can be attached to either a VM or BMS is limited by their respective maximum disk limits. (You can add more data disks to each of your nodes through the CCE console and APIs.) The data disk used by the container runtime and kubelet cannot be detached, or the node will become unavailable. For DeC nodes, the parameter details are the same as those for <b>rootVolume</b>.</p>

Parameter	Type	Description
storage	<a href="#">Storage</a> object	<p>Disk initialization management parameter. This parameter is complex to configure. For details, see <a href="#">Attaching Disks to a Node</a>.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If a node has both local and EVS disks attached, make sure to not keep the default value for this parameter, or it may result in unexpected disk partitions.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If you want to change the value range of a data disk to 20 to 32768, do not use the default value.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If you want to use the shared disk space (with the runtime and Kubernetes partitions cancelled), do not retain the default value of this parameter. For details about the shared disk space, see <a href="#">Data Disk Space Allocation</a>.</li> </ul>
publicIP	<a href="#">NodePublicIP</a> object	<p>EIP of a node.</p> <p><b>NOTE</b></p> <p>This parameter is not supported when you add a node to a node pool.</p>
nodeNicSpec	<a href="#">NodeNicSpec</a> object	NIC of the node
count	Integer	Number of nodes to be created in a batch. The value must be a positive integer greater than or equal to 1 and less than or equal to the defined limit. This parameter can be left blank when it is used for a node pool.
billingMode	Integer	<p>Node billing mode.</p> <ul style="list-style-type: none"> <li><b>0</b>: pay-per-use</li> </ul>

Parameter	Type	Description
taints	Array of <b>Taint</b> objects	<p>You can add taints to created nodes to configure anti-affinity. A maximum of 20 taints can be added. Each taint contains the following parameters:</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain can be prefixed to a key.</li> <li>• <b>Value:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed.</li> <li>• <b>Effect:</b> Available options are <b>NoSchedule</b>, <b>PreferNoSchedule</b>, and <b>NoExecute</b>. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</li> </ul> <p>Example:</p> <pre data-bbox="815 1108 1430 1339">"taints": [{   "key": "status",   "value": "unavailable",   "effect": "NoSchedule" }, {   "key": "looks",   "value": "bad",   "effect": "NoSchedule" }]</pre>

Parameter	Type	Description
k8sTags	Map<String,String>	<p>Defined in key-value pairs. A maximum of 20 key-value pairs are allowed.</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain containing a maximum of 253 characters (such as <b>example.com/my-key</b>) can be prefixed to a key.</li> <li>• <b>Value:</b> The value can be left blank or contain 1 to 63 characters that start and end with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</li> </ul> <p>Example:</p> <pre>"k8sTags": {   "key": "value" }</pre>
ecsGroupId	String	<p>ECS group ID. If this parameter is configured, nodes will be created in the specified ECS group.</p> <p><b>NOTE</b> This configuration does not take effect when a node pool is created. To ensure that all nodes in a node pool are in the same ECS group, configure <b>nodeManagement</b> in the node pool.</p>
dedicatedHostId	String	<p>ID of the DeH to which the node is scheduled.</p> <p><b>NOTE</b> This parameter is not supported when you add a node during node pool creation.</p>

Parameter	Type	Description
userTags	Array of <a href="#">UserTag</a> objects	<p>Cloud server tag (resource tag). The key of a tag must be unique. The maximum number of custom tags supported by CCE varies depending on regions and cannot exceed 8. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</p> <p><b>NOTE</b> A tag key can have a maximum of 36 characters. It can only contain letters, digits, Unicode characters, and special characters (-_).</p>
runtime	<a href="#">Runtime</a> object	<p>Container runtime. Default value:</p> <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>

Parameter	Type	Description
initializedConditions	Array of strings	<p>Custom initialization flag, which is left blank by default.</p> <p>Before CCE nodes are initialized, <b>node.cloudprovider.kubernetes.io/uninitialized</b> is added to the nodes to prevent pods from being scheduled to these nodes. When creating a node, you can configure the <b>initializedConditions</b> parameter to control the taint removal time. By default, the timeout period is not configured.</p> <p>Example:</p> <ol style="list-style-type: none"> <li>1. Create a node and add parameter "initializedConditions": ["CCEInitial", "CustomedInitial"].</li> <li>2. After custom initialization is complete, CCE will call a Kubernetes API (for example, <b>PATCH /v1/nodes/{node_ip}/status</b>) to update the node <b>conditions</b> by adding two labels of types <b>CCEInitial</b> and <b>CustomedInitial</b> and setting the status to <b>True</b>, as shown in the following:</li> </ol> <pre data-bbox="815 1055 1430 1205"> status:   conditions:   - type: CCEInitial     status: 'True'   - type: CustomedInitial     status: 'True' </pre> <ol style="list-style-type: none"> <li>1. CCE polls <b>status.Conditions</b> of nodes to check whether there are conditions of types <b>CCEInitial</b> and <b>CustomedInitial</b>. If such conditions exist and the status is <b>True</b>, the node initialization is complete and the initialization taint is removed.</li> <li>2. <b>initializedConditions</b> allows you to configure a timeout period during node creation. For example, "initializedConditions": ["CCEInitial:15m", "CustomedInitial:15m"] indicates that the timeout period is 15 minutes. After the timeout period is reached, the initialization conditions will be automatically ignored and the initialization taint will be removed when CCE polls the node. <ul style="list-style-type: none"> <li>● The timeout interval in <b>initializedConditions</b> ranges from 1 to 99 seconds.</li> <li>● Use only letters and digits. Max. characters: 20.</li> <li>● Max. flags: 2.</li> </ul> </li> </ol>

Parameter	Type	Description
		<ul style="list-style-type: none"> <li>The unit of the timeout period is minute (m).</li> </ul>
extendParam	<a href="#">NodeExtendParam</a> object	Extended parameters for creating a node.
hostnameConfig	<a href="#">HostnameConfig</a> object	Kubernetes node name configuration parameter, which is supported by clusters of v1.23.6-r0 to v1.25 or clusters of v1.25.2-r0 or later versions.
serverEnterpriseProjectID	String	Enterprise project ID of a server. CCE does not support features provided by EPS. This field is used only for synchronizing enterprise project IDs of servers. To create a node or node pool, you can specify an existing enterprise project for it. If this field is left blank, it inherits the enterprise project attributes of the cluster that runs the node or node pool. When updating a node pool, any configuration modifications you made will apply only to new nodes. You need to modify the enterprise projects of existing nodes on the EPS console. If no value is specified during the update, this field will not be updated. If this parameter is left blank, the cluster enterprise project is returned.

**Table 4-329** Login

Parameter	Type	Description
sshKey	String	Name of the key pair used for login.
userPassword	<a href="#">UserPassword</a> object	Password used for node login.

**Table 4-330** UserPassword

Parameter	Type	Description
username	String	Login account. The default value is <b>root</b> .

Parameter	Type	Description
password	String	<p>Login password. If a username and password are used to create a node, this field is shielded in the response body. A password must meet the following complexity requirements:</p> <ul style="list-style-type: none"> <li>• Contains 8 to 26 characters.</li> <li>• Contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters !@\$%^-_=+[{ }];,./?</li> <li>• Cannot contain the username or the username spelled backwards. The <b>password</b> field must be salted during node creation. For details, see <a href="#">Adding a Salt in the password Field When Creating a Node</a>.</li> </ul>

**Table 4-331** Volume

Parameter	Type	Description
size	Integer	<p>Disk size, in GiB.</p> <ul style="list-style-type: none"> <li>• System disk: 40 to 1024</li> <li>• Value range for data disks: 100 to 32768</li> </ul>
volumetype	String	<p>Disk type. For details about possible values, see the description of the <b>root_volume</b> parameter in the API used to create an ECS.</p> <ul style="list-style-type: none"> <li>• <b>SAS</b>: high I/O SAS disk</li> <li>• <b>SSD</b>: ultra-high I/O SSD disk</li> <li>• <b>SATA</b>: common I/O SATA disk. SATA disks have been removed from EVS. You can find them attached only on existing nodes.</li> </ul>
extendParam	Map<String, Object>	<p>Extended disk parameters, defined in <b>extendparam</b> in the API used to create an ECS.</p>
cluster_id	String	<p>ID of the storage pool used by the ECS system disk. This field is used only for DeC clusters, which functions as <b>dssPoolID</b>, that is, the ID of the DSS storage pool.</p>
cluster_type	String	<p>Storage class of the cloud server system disk. The value is always <b>dss</b>. This field is used only for DeC clusters.</p>



Parameter	Type	Description
hw:passthrough	Boolean	<ul style="list-style-type: none"> <li>Pay attention to this field if your ECS is SDI-compliant. If the value of this field is <b>true</b>, the created disk is of SCSI type.</li> <li>If the node pool type is <b>ElasticBMS</b>, this field must be set to <b>true</b>.</li> <li>If a node specification involves both local and EVS disks, configure the disk initialization parameters. For details, see <a href="#">Attaching Disks to a Node</a>.</li> </ul>
metadata	<b>VolumeMetadata</b> object	EVS disk encryption information. This field is mandatory only when you need to encrypt the system disk or data disks of the node to be created.

**Table 4-332** VolumeMetadata

Parameter	Type	Description
__system__encrypted	String	Whether the EVS disk is encrypted. The value <b>0</b> indicates that the EVS disk is not encrypted, and the value <b>1</b> indicates that the EVS disk is encrypted. If this parameter is not specified, EVS disks will not be encrypted by default.
__system__cmkid	String	CMK ID, which indicates encryption in <b>metadata</b> . This field is used with <b>__system__encrypted</b> .

**Table 4-333** Storage

Parameter	Type	Description
storageSelectors	Array of <b>StorageSelectors</b> objects	Disk selection. Matched disks are managed according to <b>matchLabels</b> and <b>storageType</b> .
storageGroups	Array of <b>StorageGroups</b> objects	A storage group consists of multiple storage devices. It is used to divide storage space.

**Table 4-334** StorageSelectors

Parameter	Type	Description
name	String	Selector name, used as the index of <b>selectorNames</b> in <b>storageGroup</b> . Therefore, the name of each selector must be unique.
storageType	String	Storage type. Only <b>evs</b> (EVS disks) and <b>local</b> (local disks) are supported. The local storage does not support disk selection. All local disks will form a VG, so only one storageSelector of the local type is allowed.
matchLabels	<b>matchLabels</b> object	Matching field of an EVS disk. The value can be <b>size</b> , <b>volumeType</b> , <b>metadataEncrypted</b> , <b>metadataCmkid</b> , or <b>count</b> in <b>DataVolume</b> .

**Table 4-335** matchLabels

Parameter	Type	Description
size	String	Matched disk size. If this parameter is left unspecified, the disk size is not limited. Example: 100
volumeType	String	EVS disk type. <b>SSD</b> , <b>GPSSD</b> , <b>SAS</b> , <b>ESSD</b> , and <b>SATA</b> are supported. If this parameter is left blank, there is no restriction on the disk type.
metadataEncrypted	String	Disk encryption identifier. If this parameter is set to <b>0</b> , it specifies that the disk is not encrypted. If this parameter is set to <b>1</b> , it specifies that the disk is encrypted. If this parameter is left blank, there is no restriction on the disk encryption identifier.
metadataCmkid	String	CMK ID of an encrypted disk. The value is a string of 36 bytes. If this parameter is left blank, there is no restriction on the disk key ID.
count	String	Number of disks to be selected. If this parameter is left blank, all disks of this type are selected.

**Table 4-336** StorageGroups

Parameter	Type	Description
name	String	Name of a virtual storage group, which must be unique. <b>NOTE</b> <ul style="list-style-type: none"> <li>If <b>cceManaged</b> is set to <b>true</b>, the name must be set to <b>vgpaas</b>.</li> <li>If the data disk is used as a temporary storage volume, the name must be <b>vg-everest-localvolume-ephemeral</b>.</li> <li>If the data disk is used as a persistent storage volume, the name must be <b>vg-everest-localvolume-persistent</b>.</li> </ul>
cceManaged	Boolean	Storage space for Kubernetes and runtime components. Only one group can be set to <b>true</b> . If this parameter is left blank, the default value <b>false</b> is used.
selectorNames	Array of strings	Corresponds to <b>name</b> in <b>storageSelectors</b> . One group can select multiple selectors, but one selector can be selected by only one group.
virtualSpaces	Array of <a href="#">VirtualSpace</a> objects	Detailed management of space configuration in a group.

**Table 4-337** VirtualSpace

Parameter	Type	Description
name	String	Name of a virtualSpace. Options: <b>share</b> , <b>kubernetes</b> , <b>runtime</b> , and <b>user</b> <ul style="list-style-type: none"> <li><b>share</b>: shared space configuration (with the runtime and Kubernetes partitions cancelled). <b>lvmConfig</b> needs to be configured.</li> <li><b>Kubernetes</b>: Kubernetes space configuration. <b>lvmConfig</b> needs to be configured.</li> <li><b>runtime</b>: runtime space configuration. <b>runtimeConfig</b> needs to be configured.</li> <li><b>user</b>: user space configuration. <b>lvmConfig</b> needs to be configured.</li> </ul>

Parameter	Type	Description
size	String	Size of a virtualSpace. The value must be an integer in percentage. Example: 90%. <b>NOTE</b> The value ranges from 10 to 90. The sum of percentages of all virtualSpaces in a group cannot exceed 100%.
lvmConfig	<b>LVMConfig</b> object	LVM configuration managements, applicable to the <i>share*</i> , <b>kubernetes</b> , and <b>user</b> spaces. Note that one virtual space supports only one config.
runtimeConfig	<b>RuntimeConfig</b> object	runtime configurations, applicable to the <b>runtime</b> space. Note that one virtual space supports only one config.

**Table 4-338** LVMConfig

Parameter	Type	Description
lvType	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.
path	String	Disk mount path. This parameter takes effect only in user configuration. The value is an absolute path. Digits, letters, periods (.), hyphens (-), and underscores (_) are allowed. <b>NOTE</b> Do not leave the path empty or set it to a key OS path such as the root directory. Key OS paths that cannot be specified: /, /home, /home/, /bin, /bin/, /lib, /lib/, /root, /root/, /boot, /boot/, /dev, /dev/, /etc, /etc/, /lost+found, /lost+found/, /mnt, /mnt/, /proc, /proc/, /sbin, /sbin/, /srv, /srv/, /tmp, /tmp/, /var, /var/, /media, /media/, /opt, /opt/, /selinux, /selinux/, /sys, /sys/, /usr, /usr/, /opt/cloud/, /mnt/paas/, /home/paas/, /var/paas/, /var/lib/, /var/script/

**Table 4-339** RuntimeConfig

Parameter	Type	Description
lvType	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.

**Table 4-340** NodePublicIP

Parameter	Type	Description
ids	Array of strings	IDs of existing EIPs. The quantity cannot be greater than the number of nodes to be created. <b>NOTE</b> If <b>ids</b> has been set, you do not need to set <b>count</b> and <b>eip</b> .
count	Integer	Number of EIPs to be dynamically created. <b>NOTE</b> <b>count</b> and <b>eip</b> must be set at the same time.
eip	<b>NodeEIPSpec</b> object	EIP configuration. <b>NOTE</b> This parameter is not supported when creating a node pool.

**Table 4-341** NodeEIPSpec

Parameter	Type	Description
iptype	String	EIP type, specified in <b>publicip.type</b> in the API for assigning an EIP.
bandwidth	<b>NodeBandwidth</b> object	Bandwidth parameters of the EIP

**Table 4-342** NodeBandwidth

Parameter	Type	Description
chargemode	String	<p>Bandwidth billing mode.</p> <ul style="list-style-type: none"> <li>• If this field is not specified, the billing is based on bandwidth.</li> <li>• If the field is null, the billing is based on bandwidth.</li> <li>• If the field value is <b>traffic</b>, the billing is based on traffic.</li> <li>• If the value is out of the preceding options, the cloud server will fail to be created.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• Billed by bandwidth: The billing will be based on the data transmission rate (in Mbps) of public networks. This billing mode is recommended if your bandwidth usage is higher than 10%.</li> <li>• Billed by traffic: The billing is based on the total amount of data (in GB) transmitted over the public network. This mode is available only when you are creating a pay-per-use node. This billing mode is recommended if your bandwidth usage is lower than 10%.</li> </ul>
size	Integer	Bandwidth size, specified in <b>bandwidth.size</b> in the API for assigning an EIP.
sharetype	String	Bandwidth sharing type. Value options: <b>PER</b> (exclusive bandwidth)

**Table 4-343** NodeNicSpec

Parameter	Type	Description
primaryNic	<b>NicSpec</b> object	Description of the primary NIC.
extNics	Array of <b>NicSpec</b> objects	<p>Extension NIC</p> <p><b>NOTE</b> This parameter is not supported when you add a node to a node pool.</p>

**Table 4-344** NicSpec

Parameter	Type	Description
subnetId	String	Network ID of the subnet to which a NIC belongs. If <b>subnetId</b> is not specified when a primary NIC is creating, the cluster subnet will be used. If <b>subnetList</b> is also configured for a node pool, the <b>subnetList</b> field is used for adding subnets to the node pool. When creating an extension NIC, you must specify <b>subnetId</b> .
fixedIps	Array of strings	The IP address of the primary ENI is specified using <b>fixedIps</b> . The number of IP addresses cannot be greater than the number of created nodes. Either <b>fixedIps</b> or <b>ipBlock</b> can be specified. <b>fixedIps</b> cannot be specified for extension ENIs.
ipBlock	String	CIDR format of the primary NIC IP range. The IP address of the created node falls in this range. <b>fixedIps</b> and <b>ipBlock</b> cannot be specified at the same time.
subnetList	Array of strings	Network ID list of the subnet where the NIC resides. A maximum of 20 subnets can be configured for a node pool.

**Table 4-345** Taint

Parameter	Type	Description
key	String	Key.
value	String	Value.
effect	String	Effect.

**Table 4-346** UserTag

Parameter	Type	Description
key	String	Key of the cloud server label. The value cannot start with <b>CCE-</b> or <b>__type_baremetal</b> .
value	String	Value of the cloud server label.

**Table 4-347** Runtime

Parameter	Type	Description
name	String	<p>Container runtime. Default value:</p> <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>

**Table 4-348** NodeExtendParam

Parameter	Type	Description
ecs:performancetype	String	ECS flavor types. This field is returned in the response.
maxPods	Integer	<p>Maximum number of pods that can be created on a node, including the default system pods. Value range: 16 to 256.</p> <p>This limit prevents the node from being overloaded of pods.</p> <p>The number of pods that can be created on a node is determined by multiple parameters. For details, see <a href="#">Maximum Number of Pods That Can Be Created on a Node</a>.</p>



Parameter	Type	Description
DockerLVMConfigOverride	String	<p>Docker data disk configuration item. (This parameter has been discarded. Use the <b>storage</b> field instead.) Example default configuration: "DockerLVMConfigOverride":"dockerThinpool=vgpaas/90%VG;kubernetesLV=vgpaas/10%VG;diskType=evs;lvType=linear"</p> <p>By default, if no VD disk is available, an error occurs because the data disk fails to be found. Set <b>diskType</b> based on the actual drive letter type. The following fields are included:</p> <ul style="list-style-type: none"> <li>• <b>userLV</b> (optional): size of the user space, for example, <b>vgpaas/20%VG</b>.</li> <li>• <b>userPath</b> (optional): mount path of the user space, for example, <b>/home/wqt-test</b>.</li> <li>• <b>diskType</b>: disk type. Currently, only <b>evs</b>, <b>hdd</b>, and <b>ssd</b> are supported.</li> <li>• <b>lvType</b>: type of a logic volume. The value can be <b>linear</b> or <b>striped</b>.</li> <li>• <b>dockerThinpool</b>: Docker space size, for example, <b>vgpaas/60%VG</b>.</li> <li>• <b>kubernetesLV</b>: kubelet space size, for example, <b>vgpaas/20%VG</b>.</li> </ul>
dockerBaseSize	Integer	<p>Available disk space of a single container on a node, in GB. This parameter has been discarded. Use <b>containerBaseSize</b>.</p> <p>If this parameter is left blank or is set to <b>0</b>, the default value is used. In Device Mapper mode, the default value is <b>10</b>. In OverlayFS mode, the available space of a single container is not limited by default, and the <b>dockerBaseSize</b> setting takes effect only on nodes running EulerOS in clusters of the new versions.</p> <p>For details about how to allocate the space for the container runtime, see <a href="#">Data Disk Space Allocation</a>.</p> <p>When Device Mapper is used, it is recommended that <b>dockerBaseSize</b> be set to a value less than or equal to 80 GB. If the value is too large, the container runtime may fail to be started due to long initialization. If there are special requirements for the container disk space, you can mount an external or local storage device.</p>

Parameter	Type	Description
containerBaseSize	Integer	<p>Available disk space of a single container on a node, in GiB.</p> <p>If this parameter is left blank or is set to <b>0</b>, the default value is used. In OverlayFS mode, the available space of a single container is not limited by default. In Device Mapper mode, the default value is <b>10</b>, and the <b>containerBaseSize</b> setting takes effect only on nodes running EulerOS in clusters v1.23.14-r0, v1.25.9-r0, v1.27.6-r0, v1.28.4-r0 or later.</p> <p>For details about how to allocate the space for the container runtime, see <a href="#">Data Disk Space Allocation</a>.</p> <p>When Device Mapper is used, it is recommended that <b>containerBaseSize</b> be set to a value less than or equal to 80 GB. If the value is too large, the container runtime may fail to be started due to long initialization. If there are special requirements for the container disk space, you can mount an external or local storage device. In new versions, the Device Mapper mode is used only by BMSs in the same resource pool and is being discarded.</p> <p><b>NOTE</b> This parameter cannot be updated when a node pool is being updated.</p>
publicKey	String	Public key of a node.
alpha.cce/preInstall	String	<p>Pre-installation scripts. The total number of characters after transcoding cannot exceed 10240.</p> <p><b>NOTE</b> The input value must be encoded using Base64. The method is as follows:</p> <pre>echo -n "Content to be encoded"   base64</pre>
alpha.cce/postInstall	String	<p>Post-installation scripts. The total number of characters after transcoding cannot exceed 10240.</p> <p><b>NOTE</b> The input value must be encoded using Base64. The method is as follows:</p> <pre>echo -n "Content to be encoded"   base64</pre>
alpha.cce/NodeImageID	String	This parameter is required when a custom image is used to create a BMS node.

Parameter	Type	Description
chargingMode	Integer	Billing mode of a node. This parameter has been deprecated. Use the <b>billingMode</b> parameter in <b>NodeSpec</b> .
agency_name	String	Name of an agency An agency is created by a tenant administrator on Identity and Access Management (IAM) to provide temporary credentials for CCE nodes to access cloud servers. This parameter is returned only when it is transferred during node creation.
kubeReservedMem	Integer	Reserved node memory, which is reserved for Kubernetes components.
systemReservedMem	Integer	Reserved node memory, which is reserved for system components.
init-node-password	String	Node password. If this parameter is used as a response parameter, asterisks (*) will be displayed.

**Table 4-349** HostnameConfig

Parameter	Type	Description
type	String	Configuration type of the Kubernetes node name. The default value is <b>privatelp</b> . <ul style="list-style-type: none"> <li>• <b>privatelp</b>: The Kubernetes node is named after its IP address.</li> <li>• <b>cceNodeName</b>: The Kubernetes node is named after the CCE node.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• For a node which is configured using <b>cceNodeName</b>, the name is the same as the Kubernetes node name and the ECS name. The node name cannot be changed. If the ECS name is changed on the ECS console, the node name will retain unchanged after ECS synchronization.</li> <li>• For a node which is configured using <b>cceNodeName</b>, to avoid a conflict between Kubernetes nodes, the system automatically adds a suffix to each node name. The suffix is in the format of <b>A hyphen (-) Five random characters</b>. The value of the random characters is a lowercase letter or a digit ranging from 0 to 9.</li> </ul>

**Table 4-350** NodeStatus

Parameter	Type	Description
phase	String	Node status, which is the status of a node during its lifecycle (such as installation and uninstallation) or the status of a node in a Kubernetes cluster. Options: <ul style="list-style-type: none"> <li>● <b>Build</b>: The node is being created.</li> <li>● <b>Installing</b>: The node is being managed.</li> <li>● <b>Upgrading</b>: The node is being upgraded.</li> <li>● <b>Active</b>: The node is running properly.</li> <li>● <b>Abnormal</b>: The node is abnormal.</li> <li>● <b>Deleting</b>: The node is being deleted.</li> <li>● <b>Error</b>: The node is faulty.</li> </ul>
lastProbeTime	String	Last time when the node status was checked. If the cluster is in the abnormal, frozen, or intermediate state (for example, creating), the node status check may be affected. The node status that takes more than five minutes to check has no reference value.
jobID	String	ID of a creation or deletion job
serverId	String	ID of the underlying ECS or BMS node
privateIP	String	IP address in the private network segment of the primary NIC on the node
privateIPv6IP	String	IPv6 address in the private network segment of the primary NIC on the node
publicIP	String	Node EIP. If the ECS data is not synchronized in real time, you can click <b>Sync Node Data</b> on the console to manually update the data.
deleteStatus	<a href="#">DeleteStatus</a> object	Resource status during resource deletion.

**Table 4-351** DeleteStatus

Parameter	Type	Description
previous_total	Integer	Total number of existing cluster resource records when the cluster is deleted.
current_total	Integer	Latest number of resource records, which is generated based on the current cluster resource records.

Parameter	Type	Description
updated	Integer	Total number of resource records updated when the cluster is deleted.
added	Integer	Total number of resource records updated when the cluster is deleted.
deleted	Integer	Total number of resource records deleted when the cluster is deleted.

## Example Requests

None

## Example Responses

**Status code: 200**

The job for deleting a node is successfully delivered.

```
{
  "kind": "Node",
  "apiVersion": "v3",
  "metadata": {
    "name": "new-hostname",
    "uid": "cc697ad9-9563-11e8-8ea7-0255ac106311",
    "creationTimestamp": "2018-08-01 08:20:49.944664515 +0000 UTC",
    "updateTimestamp": "2018-08-01 09:20:05.644032347 +0000 UTC",
    "annotations": {
      "kubernetes.io/node-pool.id": "az1.dc1#s1.medium#EulerOS 2.2"
    }
  },
  "spec": {
    "flavor": "s1.medium",
    "az": "az1.dc1",
    "os": "EulerOS 2.2",
    "login": {
      "sshKey": "KeyPair-001"
    },
    "rootVolume": {
      "volumetype": "SAS",
      "size": 40
    },
    "dataVolumes": [ {
      "volumetype": "SAS",
      "size": 100
    } ],
    "publicIP": {
      "eip": {
        "bandwidth": { }
      }
    },
    "billingMode": 0
  },
  "status": {
    "phase": "Deleting",
    "jobID": "661f6f7d-956c-11e8-a916-0255ac10575d",
    "serverId": "5b504f8d-33f1-4ab7-a600-b62dac967d72",
    "privateIP": "192.168.0.69",
    "publicIP": "10.154.194.59"
  }
}
```

## Status Codes

Status Code	Description
200	The job for deleting a node is successfully delivered.

## Error Codes

See [Error Codes](#).

## 4.3.6 Enabling Scale-In Protection for a Node

### Function

This API is used to enable scale-in protection for a node. A node with scale-in protection enabled cannot be removed by simply changing the number of nodes in the node pool.

### URI

POST /api/v3/projects/{project\_id}/clusters/{cluster\_id}/nodes/locknodescaledown

**Table 4-352** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

### Request Parameters

**Table 4-353** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

**Table 4-354** Request body parameters

Parameter	Mandatory	Type	Description
apiVersion	Yes	String	API version. The value is fixed at <b>v3</b> .
kind	Yes	String	API type. The value is fixed at <b>List</b> .
nodeList	No	Array of strings	ID list of nodes for which scale-in protection needs to be enabled. For details about how to obtain node IDs, see <a href="#">How to Obtain Parameters in the API URI</a> .

## Response Parameters

None

## Example Requests

None

## Example Responses

None

## Status Codes

Status Code	Description
200	The scale-in protection is enabled for the node.

## Error Codes

See [Error Codes](#).

## 4.3.7 Disabling Scale-In Protection for a Node

### Function

This API is used to disable scale-in protection for a node. A node with scale-in protection disabled can be removed by changing the number of nodes in the node pool. Scale-in protection can be disabled only for pay-per-use nodes.

### URI

POST /api/v3/projects/{project\_id}/clusters/{cluster\_id}/nodes/  
unlocknodescaledown

**Table 4-355** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

### Request Parameters

**Table 4-356** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .



**Table 4-357** Request body parameters

Parameter	Mandatory	Type	Description
apiVersion	Yes	String	API version. The value is fixed at <b>v3</b> .
kind	Yes	String	API type. The value is fixed at <b>List</b> .
nodeList	No	Array of strings	ID list of nodes for which scale-in protection needs to be disabled. For details about how to obtain node IDs, see <a href="#">How to Obtain Parameters in the API URI</a> .

## Response Parameters

None

## Example Requests

None

## Example Responses

None

## Status Codes

Status Code	Description
200	The scale-in protection is disabled for the node.

## Error Codes

See [Error Codes](#).

## 4.3.8 Synchronizing Nodes

### Function

This API is used to synchronize nodes.

### URI

GET /api/v2/projects/{project\_id}/clusters/{cluster\_id}/nodes/{node\_id}/sync

**Table 4-358** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
node_id	Yes	String	Node ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

## Request Parameters

**Table 4-359** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

## Response Parameters

None

## Example Requests

None

## Example Responses

None

## Status Codes

Status Code	Description
200	Node configuration synchronized successfully.

## Error Codes

See [Error Codes](#).

## 4.3.9 Accepting a Node

### Function

This API is used to accept a node into a specified cluster.

 **NOTE**

The URL for cluster management is in the format of `https://Endpoint/uri`. In the URL, **uri** indicates the resource path, that is, the path for API access.

### Constraints

ECSs, DeHs, and BMSs can be managed.

### URI

POST `/api/v3/projects/{project_id}/clusters/{cluster_id}/nodes/add`

**Table 4-360** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

## Request Parameters

**Table 4-361** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

**Table 4-362** Request body parameters

Parameter	Mandatory	Type	Description
apiVersion	Yes	String	API version. The value is fixed at <b>v3</b> .
kind	Yes	String	API type. The value is fixed at <b>List</b> .
nodeList	Yes	Array of <a href="#">AddNode</a> objects	List of nodes to be managed. A maximum of 200 nodes can be managed simultaneously.

**Table 4-363** AddNode

Parameter	Mandatory	Type	Description
serverID	Yes	String	Server ID, which can be obtained from the ECS or BMS console
spec	Yes	<a href="#">ReinstallNodeSpec</a> object	Node reinstallation configuration parameters. Currently, accepted nodes cannot be added into node pools.

**Table 4-364** ReinstallNodeSpec

Parameter	Mandatory	Type	Description
os	Yes	String	OS. If you specify a custom image, the actual OS version of the IMS image is used. Select an OS version supported by the current cluster.
login	Yes	<a href="#">Login</a> object	Node login mode. Either the key pair or password must be used for login. <b>NOTE</b> This parameter cannot be modified when a node pool is being updated.
name	No	String	Node name. <b>NOTE</b> Specifying this field during reinstallation will change the node name, and the server name will change accordingly. By default, the current server name is used as the node name. Enter 1 to 56 characters starting with a lowercase letter. Only lowercase letters, digits, hyphens (-), and periods (.) are allowed.
serverConfig	No	<a href="#">ReinstallServerConfig</a> object	Server configuration.
volumeConfig	No	<a href="#">ReinstallVolumeConfig</a> object	Volume management configuration.
runtimeConfig	No	<a href="#">ReinstallRuntimeConfig</a> object	Container runtime configuration.
k8sOptions	No	<a href="#">ReinstallK8sOptionsConfig</a> object	Kubernetes node configuration.
lifecycle	No	<a href="#">NodeLifecycleConfig</a> object	Customized lifecycle configuration of a node.

Parameter	Mandatory	Type	Description
initializedConditions	No	Array of strings	<p>Custom initialization flag. Before CCE nodes are initialized, they are tainted with <b>node.cloudprovider.kubernetes.io/uninitialized</b> to prevent pods from being scheduled to them.</p> <p>CCE supports custom initialization flags. After receiving the <b>initializedConditions</b> parameter, CCE converts the parameter value into a node label and provisions the label with the node, for example, <b>cloudprovider.openvessel.io/inject-initialized-conditions=CCEInitial_CustomedInitial</b>.</p> <p>After the node is labeled, its <b>status.Conditions</b> is polled to check whether the <b>type</b> of <b>conditions</b> has a flag name, such as <b>CCEInitial</b> and <b>CustomedInitial</b>. If all input flags exist and their status is <b>True</b>, the node initialization is complete and the initialization taint is removed.</p> <ul style="list-style-type: none"> <li>• Use only letters and digits. Max. characters: 20.</li> <li>• Max. flags: 2.</li> </ul>
extendParam	No	<b>ReinstallExtendParam</b> object	Extended reinstallation parameter, which is discarded.
hostnameConfig	No	<b>HostnameConfig</b> object	Kubernetes node name configuration parameter, which is supported by clusters of v1.23.6-r0 to v1.25 or clusters of v1.25.2-r0 or later versions.

**Table 4-365** Login

Parameter	Mandatory	Type	Description
sshKey	No	String	Name of the key pair used for login.
userPassword	No	<a href="#">UserPassword</a> object	Password used for node login.

**Table 4-366** UserPassword

Parameter	Mandatory	Type	Description
username	No	String	Login account. The default value is <b>root</b> .
password	Yes	String	<p>Login password. If a username and password are used to create a node, this field is shielded in the response body. A password must meet the following complexity requirements:</p> <ul style="list-style-type: none"> <li>• Contains 8 to 26 characters.</li> <li>• Contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters <code>!@\$%^_-=+[{ }];,./?</code></li> <li>• Cannot contain the username or the username spelled backwards. The <b>password</b> field must be salted during node creation. For details, see <a href="#">Adding a Salt in the password Field When Creating a Node</a>.</li> </ul>

**Table 4-367** ReinstallServerConfig

Parameter	Mandatory	Type	Description
userTags	No	Array of <b>UserTag</b> objects	Cloud server tag (resource tag). The key of a tag must be unique. The maximum number of custom tags supported by CCE varies depending on regions and cannot exceed 8.
rootVolume	No	<b>ReinstallVolumeSpec</b> object	System disk configurations used in reinstallation.

**Table 4-368** UserTag

Parameter	Mandatory	Type	Description
key	No	String	Key of the cloud server label. The value cannot start with <b>CCE-</b> or <b>__type_baremetal</b> .
value	No	String	Value of the cloud server label.

**Table 4-369** ReinstallVolumeSpec

Parameter	Mandatory	Type	Description
imageID	No	String	Custom image ID.
cmkID	No	String	User master key ID. If this parameter is left blank by default, the EVS disk is not encrypted.



**Table 4-370** ReinstallVolumeConfig

Parameter	Mandatory	Type	Description
lvmConfig	No	String	<p>Configuration item of a Docker data disk (deprecated).</p> <p>The following is an example of the default configuration:</p> <pre>"lvmConfig": "dockerThinpool=vgpaas/90%VG;kubernetesLV=vgpaas/10%VG;diskType=evs;lvType=linear"</pre> <p>The following fields are included:</p> <ul style="list-style-type: none"> <li>• <b>userLV</b>: size of the user space, for example, <b>vgpaas/20%VG</b>.</li> <li>• <b>userPath</b>: mount path of the user space, for example, <b>/home/wqt-test</b>.</li> <li>• <b>diskType</b>: disk type. Currently, only <b>evs</b>, <b>hdd</b>, and <b>ssd</b> are supported.</li> <li>• <b>lvType</b>: type of a logic volume. The value can be <b>linear</b> or <b>striped</b>.</li> <li>• <b>dockerThinpool</b>: Docker space size, for example, <b>vgpaas/60%VG</b>.</li> <li>• <b>kubernetesLV</b>: kubelet space size, for example, <b>vgpaas/20%VG</b>.</li> </ul>

Parameter	Mandatory	Type	Description
storage	No	<a href="#">Storage object</a>	<p>Disk initialization management parameter.</p> <p>This parameter is complex to configure. For details, see <a href="#">Attaching Disks to a Node</a>.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If a node has both local and EVS disks attached, make sure to not keep the default value for this parameter, or it may result in unexpected disk partitions.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If you want to change the value range of a data disk to 20 to 32768, do not use the default value.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If you want to use the shared disk space (with the runtime and Kubernetes partitions cancelled), do not retain the default value of this parameter. For details about the shared disk space, see <a href="#">Data Disk Space Allocation</a>.</li> </ul>

**Table 4-371** Storage

Parameter	Mandatory	Type	Description
storageSelectors	Yes	Array of <a href="#">StorageSelectors</a> objects	Disk selection. Matched disks are managed according to <b>matchLabels</b> and <b>storageType</b> .
storageGroups	Yes	Array of <a href="#">StorageGroups</a> objects	A storage group consists of multiple storage devices. It is used to divide storage space.

**Table 4-372** StorageSelectors

Parameter	Mandatory	Type	Description
name	Yes	String	Selector name, used as the index of <b>selectorNames</b> in <b>storageGroup</b> . Therefore, the name of each selector must be unique.
storageType	Yes	String	Storage type. Only <b>evs</b> (EVS disks) and <b>local</b> (local disks) are supported. The local storage does not support disk selection. All local disks will form a VG, so only one storageSelector of the local type is allowed.
matchLabels	No	<b>matchLabels</b> object	Matching field of an EVS disk. The value can be <b>size</b> , <b>volumeType</b> , <b>metadataEncrypted</b> , <b>metadataCmkid</b> , or <b>count</b> in <b>DataVolume</b> .

**Table 4-373** matchLabels

Parameter	Mandatory	Type	Description
size	No	String	Matched disk size. If this parameter is left unspecified, the disk size is not limited. Example: 100
volumeType	No	String	EVS disk type. <b>SSD</b> , <b>GPSSD</b> , <b>SAS</b> , <b>ESSD</b> , and <b>SATA</b> are supported. If this parameter is left blank, there is no restriction on the disk type.
metadataEncrypted	No	String	Disk encryption identifier. If this parameter is set to <b>0</b> , it specifies that the disk is not encrypted. If this parameter is set to <b>1</b> , it specifies that the disk is encrypted. If this parameter is left blank, there is no restriction on the disk encryption identifier.

Parameter	Mandatory	Type	Description
metadataCmkid	No	String	CMK ID of an encrypted disk. The value is a string of 36 bytes. If this parameter is left blank, there is no restriction on the disk key ID.
count	No	String	Number of disks to be selected. If this parameter is left blank, all disks of this type are selected.

**Table 4-374** StorageGroups

Parameter	Mandatory	Type	Description
name	Yes	String	Name of a virtual storage group, which must be unique. <b>NOTE</b> <ul style="list-style-type: none"> <li>If <b>cceManaged</b> is set to <b>true</b>, the name must be set to <b>vgpaas</b>.</li> <li>If the data disk is used as a temporary storage volume, the name must be <b>vg-everest-localvolume-ephemeral</b>.</li> <li>If the data disk is used as a persistent storage volume, the name must be <b>vg-everest-localvolume-persistent</b>.</li> </ul>
cceManaged	No	Boolean	Storage space for Kubernetes and runtime components. Only one group can be set to <b>true</b> . If this parameter is left blank, the default value <b>false</b> is used.
selectorNames	Yes	Array of strings	Corresponds to <b>name</b> in <b>storageSelectors</b> . One group can select multiple selectors, but one selector can be selected by only one group.
virtualSpaces	Yes	Array of <b>VirtualSpace</b> objects	Detailed management of space configuration in a group.

**Table 4-375** VirtualSpace

Parameter	Mandatory	Type	Description
name	Yes	String	<p>Name of a virtualSpace. Options: <b>share</b>, <b>kubernetes</b>, <b>runtime</b>, and <b>user</b></p> <ul style="list-style-type: none"> <li>• <b>share</b>: shared space configuration (with the runtime and Kubernetes partitions cancelled). <b>lvmConfig</b> needs to be configured.</li> <li>• <b>Kubernetes</b>: Kubernetes space configuration. <b>lvmConfig</b> needs to be configured.</li> <li>• <b>runtime</b>: runtime space configuration. <b>runtimeConfig</b> needs to be configured.</li> <li>• <b>user</b>: user space configuration. <b>lvmConfig</b> needs to be configured.</li> </ul>
size	Yes	String	<p>Size of a virtualSpace. The value must be an integer in percentage. Example: 90%.</p> <p><b>NOTE</b> The value ranges from 10 to 90. The sum of percentages of all virtualSpaces in a group cannot exceed 100%.</p>
lvmConfig	No	<b>LVMConfig</b> object	<p>LVM configuration managements, applicable to the <i>share</i>*, <b>kubernetes</b>, and <b>user</b> spaces. Note that one virtual space supports only one config.</p>
runtimeConfig	No	<b>RuntimeConfig</b> object	<p>runtime configurations, applicable to the <b>runtime</b> space. Note that one virtual space supports only one config.</p>

**Table 4-376** LVMConfig

Parameter	Mandatory	Type	Description
lvType	Yes	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.
path	No	String	Disk mount path. This parameter takes effect only in user configuration. The value is an absolute path. Digits, letters, periods (.), hyphens (-), and underscores (_) are allowed.  <b>NOTE</b> Do not leave the path empty or set it to a key OS path such as the root directory. Key OS paths that cannot be specified: /, /home, /home/, /bin, /bin/, /lib, /lib/, /root, /root/, /boot, /boot/, /dev, /dev/, /etc, /etc/, /lost+found, /lost+found/, /mnt, /mnt/, /proc, /proc/, /sbin, /sbin/, /srv, /srv/, /tmp, /tmp/, /var, /var/, /media, /media/, /opt, /opt/, /selinux, /selinux/, /sys, /sys/, /usr, /usr/, /opt/cloud/, /mnt/paas/, /home/paas/, /var/paas/, /var/lib/, /var/script/

**Table 4-377** RuntimeConfig

Parameter	Mandatory	Type	Description
lvType	Yes	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.

**Table 4-378** ReinstallRuntimeConfig

Parameter	Mandatory	Type	Description
dockerBaseSize	No	Integer	<p>Available disk space of a single container on a node, in GB. This parameter has been discarded. Use <b>containerBaseSize</b>.</p> <p>If this parameter is left blank or is set to <b>0</b>, the default value is used. In Device Mapper mode, the default value is <b>10</b>. In OverlayFS mode, the available space of a single container is not limited by default, and the <b>dockerBaseSize</b> setting takes effect only on nodes running EulerOS in clusters of the new versions.</p> <p>For details about how to allocate the space for the container runtime, see <a href="#">Data Disk Space Allocation</a>.</p> <p>When Device Mapper is used, it is recommended that <b>dockerBaseSize</b> be set to a value less than or equal to 80 GB. If the value is too large, the container runtime may fail to be started due to long initialization. If there are special requirements for the container disk space, you can mount an external or local storage device.</p>

Parameter	Mandatory	Type	Description
containerBaseSize	No	Integer	<p>Available disk space of a single container on a node, in GiB.</p> <p>If this parameter is left blank or is set to <b>0</b>, the default value is used. In OverlayFS mode, the available space of a single container is not limited by default. In Device Mapper mode, the default value is <b>10</b>, and the <b>containerBaseSize</b> setting takes effect only on nodes running EulerOS in clusters v1.23.14-r0, v1.25.9-r0, v1.27.6-r0, v1.28.4-r0 or later.</p> <p>For details about how to allocate the space for the container runtime, see <a href="#">Data Disk Space Allocation</a>.</p> <p>When Device Mapper is used, it is recommended that <b>containerBaseSize</b> be set to a value less than or equal to 80 GB. If the value is too large, the container runtime may fail to be started due to long initialization. If there are special requirements for the container disk space, you can mount an external or local storage device. In new versions, the Device Mapper mode is used only by BMSs in the same resource pool and is being discarded.</p> <p><b>NOTE</b> This parameter cannot be updated when a node pool is being updated.</p>



Parameter	Mandatory	Type	Description
runtime	No	<b>Runtime</b> object	Container runtime. Default value: <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>

**Table 4-379** Runtime

Parameter	Mandatory	Type	Description
name	No	String	Container runtime. Default value: <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>

**Table 4-380** ReinstallK8sOptionsConfig

Parameter	Mandatory	Type	Description
labels	No	Map<String,String>	<p>Defined in key-value pairs. A maximum of 20 key-value pairs are allowed.</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain containing a maximum of 253 characters (such as <b>example.com/my-key</b>) can be prefixed to a key.</li> <li>• <b>Value:</b> The value can be left blank or contain 1 to 63 characters that start and end with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed.</li> </ul> <p>Example:</p> <pre>"k8sTags": {   "key": "value" }</pre>

Parameter	Mandatory	Type	Description
taints	No	Array of <b>Taint</b> objects	<p>You can add taints to created nodes to configure anti-affinity. A maximum of 20 taints can be added. Each taint contains the following parameters:</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain can be prefixed to a key.</li> <li>• <b>Value:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed.</li> <li>• <b>Effect:</b> Available options are <b>NoSchedule</b>, <b>PreferNoSchedule</b>, and <b>NoExecute</b>.</li> </ul> <p>Example:</p> <pre>"taints": [{   "key": "status",   "value": "unavailable",   "effect": "NoSchedule" }, {   "key": "looks",   "value": "bad",   "effect": "NoSchedule" }]</pre>
maxPods	No	Integer	Maximum number of pods that can be created on a node, including the default system pods. Value range: 16 to 256. This limit prevents the node from being overloaded of pods.

**Table 4-381** Taint

Parameter	Mandatory	Type	Description
key	Yes	String	Key.

Parameter	Mandatory	Type	Description
value	No	String	Value.
effect	Yes	String	Effect.

**Table 4-382** NodeLifecycleConfig

Parameter	Mandatory	Type	Description
preInstall	No	String	Pre-installation scripts. The total number of characters after transcoding cannot exceed 10240.  <b>NOTE</b> The input value must be encoded using Base64. The method is as follows: <pre>echo -n "Content to be encoded"   base64</pre>
postInstall	No	String	Post-installation scripts. The total number of characters after transcoding cannot exceed 10240.  <b>NOTE</b> The input value must be encoded using Base64. The method is as follows: <pre>echo -n "Content to be encoded"   base64</pre>

**Table 4-383** ReinstallExtendParam

Parameter	Mandatory	Type	Description
alpha.cce/ NodeImageID	No	String	(Discarded) ID of the user image to run the target OS. Specifying this parameter is equivalent to specifying <b>imageID</b> in <b>ReinstallVolumeSpec</b> . The original value will be overwritten.

**Table 4-384** HostnameConfig

Parameter	Mandatory	Type	Description
type	Yes	String	<p>Configuration type of the Kubernetes node name. The default value is <b>privatelp</b>.</p> <ul style="list-style-type: none"> <li>• <b>privatelp</b>: The Kubernetes node is named after its IP address.</li> <li>• <b>cceNodeName</b>: The Kubernetes node is named after the CCE node.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• For a node which is configured using <b>cceNodeName</b>, the name is the same as the Kubernetes node name and the ECS name. The node name cannot be changed. If the ECS name is changed on the ECS console, the node name will retain unchanged after ECS synchronization.</li> <li>• For a node which is configured using <b>cceNodeName</b>, to avoid a conflict between Kubernetes nodes, the system automatically adds a suffix to each node name. The suffix is in the format of <b>A hyphen (-) Five random characters</b>. The value of the random characters is a lowercase letter or a digit ranging from 0 to 9.</li> </ul>

## Response Parameters

Status code: 200

**Table 4-385** Response body parameters

Parameter	Type	Description
jobid	String	Job ID returned after the job is delivered. The job ID can be used to query the job execution status.

## Example Requests

Add a node running EulerOS 2.5 to a cluster.

```
POST /api/v3/projects/{project_id}/clusters/{cluster_id}/nodes/add
{
  "kind": "List",
  "apiVersion": "v3",
  "nodeList": [ {
    "serverID": "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
    "spec": {
      "name": "my-ecs-0001",
      "os": "EulerOS 2.5",
      "login": {
        "sshKey": "KeyPair-001"
      }
    }
  }
]
}
```

## Example Responses

**Status code: 200**

The job for accepting a node into a specified cluster is successfully delivered.

```
{
  "jobid": "2ec9b78d-9368-46f3-8f29-d1a95622a568"
}
```

## Status Codes

Status Code	Description
200	The job for accepting a node into a specified cluster is successfully delivered.

## Error Codes

See [Error Codes](#).

### 4.3.10 Managing a Node in a Customized Node Pool

#### Function

This API is used to manage a node in a customized node pool of a cluster. Spot instances are not supported.

#### NOTE

- ECSs, DeHs, and BMSs can be managed.
- The URL for cluster management is in the format of **https://Endpoint/uri**, where **uri** specifies the resource path for API access.

## URI

POST /api/v3/projects/{project\_id}/clusters/{cluster\_id}/nodepools/{nodepool\_id}/nodes/add

**Table 4-386** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
nodepool_id	Yes	String	Node pool ID.

## Request Parameters

**Table 4-387** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

**Table 4-388** Request body parameters

Parameter	Mandatory	Type	Description
apiVersion	Yes	String	API version. The value is fixed at <b>v3</b> .
kind	Yes	String	API type. The value is fixed at <b>List</b> .

Parameter	Mandatory	Type	Description
nodeList	Yes	Array of <a href="#">AddNodesToNodePool</a> objects	List of nodes to be managed. A maximum of 200 nodes can be managed simultaneously.

**Table 4-389** AddNodesToNodePool

Parameter	Mandatory	Type	Description
serverID	Yes	String	Server ID, which can be obtained from the ECS or BMS console

## Response Parameters

**Status code: 200**

**Table 4-390** Response body parameters

Parameter	Type	Description
jobid	String	Job ID returned after the job is delivered. The job ID can be used to query the job execution status.

## Example Requests

Manage a node in a customized node pool.

POST /api/v3/projects/{project\_id}/clusters/{cluster\_id}/nodepool/{nodepool\_id}/nodes/add

```
{
  "kind": "List",
  "apiVersion": "v3",
  "nodeList": [
    {
      "serverID": "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx1"
    },
    {
      "serverID": "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx2",
    }
  ]
}
```

## Example Responses

**Status code: 200**

The job for managing a node in the customized node pool in the cluster delivered.



```
{
  "jobid" : "2ec9b78d-9368-46f3-8f29-d1a95622a568"
}
```

## Status Codes

Status Code	Description
200	The job for managing a node in the customized node pool in the cluster delivered.

## Error Codes

See [Error Codes](#).

### 4.3.11 Resetting a Node

#### Function

This API is used to reset a node in a specified cluster.

#### NOTE

The URL for cluster management is in the format of `https://Endpoint/uri`. In the URL, **uri** indicates the resource path, that is, the path for API access.

#### URI

POST `/api/v3/projects/{project_id}/clusters/{cluster_id}/nodes/reset`

**Table 4-391** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

## Request Parameters

**Table 4-392** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

**Table 4-393** Request body parameters

Parameter	Mandatory	Type	Description
apiVersion	Yes	String	API version. The value is fixed at <b>v3</b> .
kind	Yes	String	API type. The value is fixed at <b>List</b> .
nodeList	Yes	Array of <a href="#">ResetNode</a> objects	List of nodes to be reset. A maximum of 200 nodes can be reset at a time.

**Table 4-394** ResetNode

Parameter	Mandatory	Type	Description
nodeID	Yes	String	Node ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
spec	Yes	<a href="#">ReinstallNodeSpec</a> object	Node reinstallation configuration parameters. Nodes in a node pool cannot be specified externally. These nodes will be reinstalled based on the node pool settings. By default, this parameter is mandatory for nodes in a node pool.

**Table 4-395** ReinstallNodeSpec

Parameter	Mandatory	Type	Description
os	Yes	String	OS. If you specify a custom image, the actual OS version of the IMS image is used. Select an OS version supported by the current cluster.
login	Yes	<a href="#">Login</a> object	Node login mode. Either the key pair or password must be used for login. <b>NOTE</b> This parameter cannot be modified when a node pool is being updated.
name	No	String	Node name. <b>NOTE</b> Specifying this field during reinstallation will change the node name, and the server name will change accordingly. By default, the current server name is used as the node name. Enter 1 to 56 characters starting with a lowercase letter. Only lowercase letters, digits, hyphens (-), and periods (.) are allowed.
serverConfig	No	<a href="#">ReinstallServerConfig</a> object	Server configuration.
volumeConfig	No	<a href="#">ReinstallVolumeConfig</a> object	Volume management configuration.
runtimeConfig	No	<a href="#">ReinstallRuntimeConfig</a> object	Container runtime configuration.
k8sOptions	No	<a href="#">ReinstallK8sOptionsConfig</a> object	Kubernetes node configuration.
lifecycle	No	<a href="#">NodeLifecycleConfig</a> object	Customized lifecycle configuration of a node.

Parameter	Mandatory	Type	Description
initializedConditions	No	Array of strings	<p>Custom initialization flag. Before CCE nodes are initialized, they are tainted with <b>node.cloudprovider.kubernetes.io/uninitialized</b> to prevent pods from being scheduled to them.</p> <p>CCE supports custom initialization flags. After receiving the <b>initializedConditions</b> parameter, CCE converts the parameter value into a node label and provisions the label with the node, for example, <b>cloudprovider.openvessel.io/inject-initialized-conditions=CCEInitial_CustomedInitial</b>.</p> <p>After the node is labeled, its <b>status.Conditions</b> is polled to check whether the <b>type</b> of <b>conditions</b> has a flag name, such as <b>CCEInitial</b> and <b>CustomedInitial</b>. If all input flags exist and their status is <b>True</b>, the node initialization is complete and the initialization taint is removed.</p> <ul style="list-style-type: none"> <li>• Use only letters and digits. Max. characters: 20.</li> <li>• Max. flags: 2.</li> </ul>
extendParam	No	<b>ReinstallExtendParam</b> object	Extended reinstallation parameter, which is discarded.
hostnameConfig	No	<b>HostnameConfig</b> object	Kubernetes node name configuration parameter, which is supported by clusters of v1.23.6-r0 to v1.25 or clusters of v1.25.2-r0 or later versions.

**Table 4-396** Login

Parameter	Mandatory	Type	Description
sshKey	No	String	Name of the key pair used for login.
userPassword	No	<a href="#">UserPassword</a> object	Password used for node login.

**Table 4-397** UserPassword

Parameter	Mandatory	Type	Description
username	No	String	Login account. The default value is <b>root</b> .
password	Yes	String	<p>Login password. If a username and password are used to create a node, this field is shielded in the response body. A password must meet the following complexity requirements:</p> <ul style="list-style-type: none"> <li>• Contains 8 to 26 characters.</li> <li>• Contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters !@\$%&amp;^_-=+[{ }];,./?</li> <li>• Cannot contain the username or the username spelled backwards. The <b>password</b> field must be salted during node creation. For details, see <a href="#">Adding a Salt in the password Field When Creating a Node</a>.</li> </ul>

**Table 4-398** ReinstallServerConfig

Parameter	Mandatory	Type	Description
userTags	No	Array of <b>UserTag</b> objects	Cloud server tag (resource tag). The key of a tag must be unique. The maximum number of custom tags supported by CCE varies depending on regions and cannot exceed 8.
rootVolume	No	<b>ReinstallVolumeSpec</b> object	System disk configurations used in reinstallation.

**Table 4-399** UserTag

Parameter	Mandatory	Type	Description
key	No	String	Key of the cloud server label. The value cannot start with <b>CCE-</b> or <b>__type_baremetal</b> .
value	No	String	Value of the cloud server label.

**Table 4-400** ReinstallVolumeSpec

Parameter	Mandatory	Type	Description
imageID	No	String	Custom image ID.
cmkID	No	String	User master key ID. If this parameter is left blank by default, the EVS disk is not encrypted.

**Table 4-401** ReinstallVolumeConfig

Parameter	Mandatory	Type	Description
lvmConfig	No	String	<p>Configuration item of a Docker data disk (deprecated).</p> <p>The following is an example of the default configuration:</p> <pre>"lvmConfig": "dockerThinpool=vgpaas/90%VG;kubernetesLV=vgpaas/10%VG;diskType=evs;lvType=linear"</pre> <p>The following fields are included:</p> <ul style="list-style-type: none"> <li>• <b>userLV</b>: size of the user space, for example, <b>vgpaas/20%VG</b>.</li> <li>• <b>userPath</b>: mount path of the user space, for example, <b>/home/wqt-test</b>.</li> <li>• <b>diskType</b>: disk type. Currently, only <b>evs</b>, <b>hdd</b>, and <b>ssd</b> are supported.</li> <li>• <b>lvType</b>: type of a logic volume. The value can be <b>linear</b> or <b>striped</b>.</li> <li>• <b>dockerThinpool</b>: Docker space size, for example, <b>vgpaas/60%VG</b>.</li> <li>• <b>kubernetesLV</b>: kubelet space size, for example, <b>vgpaas/20%VG</b>.</li> </ul>

Parameter	Mandatory	Type	Description
storage	No	<a href="#">Storage object</a>	<p>Disk initialization management parameter.</p> <p>This parameter is complex to configure. For details, see <a href="#">Attaching Disks to a Node</a>.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If a node has both local and EVS disks attached, make sure to not keep the default value for this parameter, or it may result in unexpected disk partitions.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If you want to change the value range of a data disk to 20 to 32768, do not use the default value.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If you want to use the shared disk space (with the runtime and Kubernetes partitions cancelled), do not retain the default value of this parameter. For details about the shared disk space, see <a href="#">Data Disk Space Allocation</a>.</li> </ul>

**Table 4-402** Storage

Parameter	Mandatory	Type	Description
storageSelectors	Yes	Array of <a href="#">StorageSelectors</a> objects	Disk selection. Matched disks are managed according to <b>matchLabels</b> and <b>storageType</b> .
storageGroups	Yes	Array of <a href="#">StorageGroups</a> objects	A storage group consists of multiple storage devices. It is used to divide storage space.



**Table 4-403** StorageSelectors

Parameter	Mandatory	Type	Description
name	Yes	String	Selector name, used as the index of <b>selectorNames</b> in <b>storageGroup</b> . Therefore, the name of each selector must be unique.
storageType	Yes	String	Storage type. Only <b>evs</b> (EVS disks) and <b>local</b> (local disks) are supported. The local storage does not support disk selection. All local disks will form a VG, so only one storageSelector of the local type is allowed.
matchLabels	No	<b>matchLabels</b> object	Matching field of an EVS disk. The value can be <b>size</b> , <b>volumeType</b> , <b>metadataEncrypted</b> , <b>metadataCmkid</b> , or <b>count</b> in <b>DataVolume</b> .

**Table 4-404** matchLabels

Parameter	Mandatory	Type	Description
size	No	String	Matched disk size. If this parameter is left unspecified, the disk size is not limited. Example: 100
volumeType	No	String	EVS disk type. <b>SSD</b> , <b>GPSSD</b> , <b>SAS</b> , <b>ESSD</b> , and <b>SATA</b> are supported. If this parameter is left blank, there is no restriction on the disk type.
metadataEncrypted	No	String	Disk encryption identifier. If this parameter is set to <b>0</b> , it specifies that the disk is not encrypted. If this parameter is set to <b>1</b> , it specifies that the disk is encrypted. If this parameter is left blank, there is no restriction on the disk encryption identifier.

Parameter	Mandatory	Type	Description
metadataCmkid	No	String	CMK ID of an encrypted disk. The value is a string of 36 bytes. If this parameter is left blank, there is no restriction on the disk key ID.
count	No	String	Number of disks to be selected. If this parameter is left blank, all disks of this type are selected.

**Table 4-405** StorageGroups

Parameter	Mandatory	Type	Description
name	Yes	String	Name of a virtual storage group, which must be unique. <b>NOTE</b> <ul style="list-style-type: none"> <li>If <b>cceManaged</b> is set to <b>true</b>, the name must be set to <b>vgpaas</b>.</li> <li>If the data disk is used as a temporary storage volume, the name must be <b>vg-everest-localvolume-ephemeral</b>.</li> <li>If the data disk is used as a persistent storage volume, the name must be <b>vg-everest-localvolume-persistent</b>.</li> </ul>
cceManaged	No	Boolean	Storage space for Kubernetes and runtime components. Only one group can be set to <b>true</b> . If this parameter is left blank, the default value <b>false</b> is used.
selectorNames	Yes	Array of strings	Corresponds to <b>name</b> in <b>storageSelectors</b> . One group can select multiple selectors, but one selector can be selected by only one group.
virtualSpaces	Yes	Array of <b>VirtualSpace</b> objects	Detailed management of space configuration in a group.

**Table 4-406** VirtualSpace

Parameter	Mandatory	Type	Description
name	Yes	String	<p>Name of a virtualSpace. Options: <b>share</b>, <b>kubernetes</b>, <b>runtime</b>, and <b>user</b></p> <ul style="list-style-type: none"> <li>• <b>share</b>: shared space configuration (with the runtime and Kubernetes partitions cancelled). <b>lvmConfig</b> needs to be configured.</li> <li>• <b>Kubernetes</b>: Kubernetes space configuration. <b>lvmConfig</b> needs to be configured.</li> <li>• <b>runtime</b>: runtime space configuration. <b>runtimeConfig</b> needs to be configured.</li> <li>• <b>user</b>: user space configuration. <b>lvmConfig</b> needs to be configured.</li> </ul>
size	Yes	String	<p>Size of a virtualSpace. The value must be an integer in percentage. Example: 90%.</p> <p><b>NOTE</b> The value ranges from 10 to 90. The sum of percentages of all virtualSpaces in a group cannot exceed 100%.</p>
lvmConfig	No	<b>LVMConfig</b> object	<p>LVM configuration managements, applicable to the <i>share</i>*, <b>kubernetes</b>, and <b>user</b> spaces. Note that one virtual space supports only one config.</p>
runtimeConfig	No	<b>RuntimeConfig</b> object	<p>runtime configurations, applicable to the <b>runtime</b> space. Note that one virtual space supports only one config.</p>

**Table 4-407** LVMConfig

Parameter	Mandatory	Type	Description
lvType	Yes	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.
path	No	String	Disk mount path. This parameter takes effect only in user configuration. The value is an absolute path. Digits, letters, periods (.), hyphens (-), and underscores (_) are allowed.  <b>NOTE</b> Do not leave the path empty or set it to a key OS path such as the root directory. Key OS paths that cannot be specified: /, /home, /home/, /bin, /bin/, /lib, /lib/, /root, /root/, /boot, /boot/, /dev, /dev/, /etc, /etc/, /lost+found, /lost+found/, /mnt, /mnt/, /proc, /proc/, /sbin, /sbin/, /srv, /srv/, /tmp, /tmp/, /var, /var/, /media, /media/, /opt, /opt/, /selinux, /selinux/, /sys, /sys/, /usr, /usr/, /opt/cloud/, /mnt/paas/, /home/paas/, /var/paas/, /var/lib/, /var/script/

**Table 4-408** RuntimeConfig

Parameter	Mandatory	Type	Description
lvType	Yes	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.

**Table 4-409** ReinstallRuntimeConfig

Parameter	Mandatory	Type	Description
dockerBaseSize	No	Integer	<p>Available disk space of a single container on a node, in GB. This parameter has been discarded. Use <b>containerBaseSize</b>.</p> <p>If this parameter is left blank or is set to <b>0</b>, the default value is used. In Device Mapper mode, the default value is <b>10</b>. In OverlayFS mode, the available space of a single container is not limited by default, and the <b>dockerBaseSize</b> setting takes effect only on nodes running EulerOS in clusters of the new versions.</p> <p>For details about how to allocate the space for the container runtime, see <a href="#">Data Disk Space Allocation</a>.</p> <p>When Device Mapper is used, it is recommended that <b>dockerBaseSize</b> be set to a value less than or equal to 80 GB. If the value is too large, the container runtime may fail to be started due to long initialization. If there are special requirements for the container disk space, you can mount an external or local storage device.</p>

Parameter	Mandatory	Type	Description
containerBaseSize	No	Integer	<p>Available disk space of a single container on a node, in GiB.</p> <p>If this parameter is left blank or is set to <b>0</b>, the default value is used. In OverlayFS mode, the available space of a single container is not limited by default. In Device Mapper mode, the default value is <b>10</b>, and the <b>containerBaseSize</b> setting takes effect only on nodes running EulerOS in clusters v1.23.14-r0, v1.25.9-r0, v1.27.6-r0, v1.28.4-r0 or later.</p> <p>For details about how to allocate the space for the container runtime, see <a href="#">Data Disk Space Allocation</a>.</p> <p>When Device Mapper is used, it is recommended that <b>containerBaseSize</b> be set to a value less than or equal to 80 GB. If the value is too large, the container runtime may fail to be started due to long initialization. If there are special requirements for the container disk space, you can mount an external or local storage device. In new versions, the Device Mapper mode is used only by BMSs in the same resource pool and is being discarded.</p> <p><b>NOTE</b> This parameter cannot be updated when a node pool is being updated.</p>

Parameter	Mandatory	Type	Description
runtime	No	<b>Runtime</b> object	Container runtime. Default value: <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>

**Table 4-410** Runtime

Parameter	Mandatory	Type	Description
name	No	String	Container runtime. Default value: <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>

**Table 4-411** ReinstallK8sOptionsConfig

Parameter	Mandatory	Type	Description
labels	No	Map<String,String>	<p>Defined in key-value pairs. A maximum of 20 key-value pairs are allowed.</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain containing a maximum of 253 characters (such as <b>example.com/my-key</b>) can be prefixed to a key.</li> <li>• <b>Value:</b> The value can be left blank or contain 1 to 63 characters that start and end with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed.</li> </ul> <p>Example:</p> <pre>"k8sTags": {   "key": "value" }</pre>



Parameter	Mandatory	Type	Description
taints	No	Array of <b>Taint</b> objects	<p>You can add taints to created nodes to configure anti-affinity. A maximum of 20 taints can be added. Each taint contains the following parameters:</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain can be prefixed to a key.</li> <li>• <b>Value:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed.</li> <li>• <b>Effect:</b> Available options are <b>NoSchedule</b>, <b>PreferNoSchedule</b>, and <b>NoExecute</b>.</li> </ul> <p>Example:</p> <pre>"taints": [{   "key": "status",   "value": "unavailable",   "effect": "NoSchedule" }, {   "key": "looks",   "value": "bad",   "effect": "NoSchedule" }]</pre>
maxPods	No	Integer	Maximum number of pods that can be created on a node, including the default system pods. Value range: 16 to 256. This limit prevents the node from being overloaded of pods.

**Table 4-412** Taint

Parameter	Mandatory	Type	Description
key	Yes	String	Key.

Parameter	Mandatory	Type	Description
value	No	String	Value.
effect	Yes	String	Effect.

**Table 4-413** NodeLifecycleConfig

Parameter	Mandatory	Type	Description
preInstall	No	String	Pre-installation scripts. The total number of characters after transcoding cannot exceed 10240. <b>NOTE</b> The input value must be encoded using Base64. The method is as follows: <pre>echo -n "Content to be encoded"   base64</pre>
postInstall	No	String	Post-installation scripts. The total number of characters after transcoding cannot exceed 10240. <b>NOTE</b> The input value must be encoded using Base64. The method is as follows: <pre>echo -n "Content to be encoded"   base64</pre>

**Table 4-414** ReinstallExtendParam

Parameter	Mandatory	Type	Description
alpha.cce/ NodeImageID	No	String	(Discarded) ID of the user image to run the target OS. Specifying this parameter is equivalent to specifying <b>imageID</b> in <b>ReinstallVolumeSpec</b> . The original value will be overwritten.

**Table 4-415** HostnameConfig

Parameter	Mandatory	Type	Description
type	Yes	String	<p>Configuration type of the Kubernetes node name. The default value is <b>privatelp</b>.</p> <ul style="list-style-type: none"> <li>• <b>privatelp</b>: The Kubernetes node is named after its IP address.</li> <li>• <b>cceNodeName</b>: The Kubernetes node is named after the CCE node.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• For a node which is configured using <b>cceNodeName</b>, the name is the same as the Kubernetes node name and the ECS name. The node name cannot be changed. If the ECS name is changed on the ECS console, the node name will retain unchanged after ECS synchronization.</li> <li>• For a node which is configured using <b>cceNodeName</b>, to avoid a conflict between Kubernetes nodes, the system automatically adds a suffix to each node name. The suffix is in the format of <b>A hyphen (-) Five random characters</b>. The value of the random characters is a lowercase letter or a digit ranging from 0 to 9.</li> </ul>

## Response Parameters

Status code: 200

**Table 4-416** Response body parameters

Parameter	Type	Description
jobid	String	Job ID returned after the job is delivered. The job ID can be used to query the job execution status.

## Example Requests

- Reset the nodes running EulerOS 2.5 in the default node pool.

```
POST /api/v3/projects/{project_id}/clusters/{cluster_id}/nodes/reset
```

```
{
  "kind": "List",
  "apiVersion": "v3",
  "nodeList": [ {
    "nodeID": "yyyyyyyy-yyyy-yyyy-yyyyyyyyyyyy",
    "spec": {
      "name": "my-ecs-0001",
      "os": "EulerOS 2.5",
      "login": {
        "sshKey": "KeyPair-001"
      }
    }
  }
]
```

- Reset a node in a node pool (invalid spec).

```
POST /api/v3/projects/{project_id}/clusters/{cluster_id}/nodes/reset
```

```
{
  "kind": "List",
  "apiVersion": "v3",
  "nodeList": [ {
    "nodeID": "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
    "spec": {
      "name": "my-ecs-0001",
      "os": "EulerOS 2.5",
      "login": {
        "sshKey": "KeyPair-001"
      }
    }
  }
]
```

## Example Responses

**Status code: 200**

The job for resetting a node in a specified cluster is successfully delivered.

```
{
  "jobid": "2ec9b78d-9368-46f3-8f29-d1a95622a568"
}
```

## Status Codes

Status Code	Description
200	The job for resetting a node in a specified cluster is successfully delivered.

## Error Codes

See [Error Codes](#).

## 4.3.12 Removing a Node

### Function

This API is used to remove a node from a specified cluster.

 **NOTE**

The URL for cluster management is in the format of `https://Endpoint/uri`. In the URL, **uri** indicates the resource path, that is, the path for API access.

### URI

PUT `/api/v3/projects/{project_id}/clusters/{cluster_id}/nodes/operation/remove`

**Table 4-417** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

### Request Parameters

**Table 4-418** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

**Table 4-419** Request body parameters

Parameter	Mandatory	Type	Description
apiVersion	No	String	API version. The value is fixed at <b>v3</b> .
kind	No	String	API type. The value is fixed at <b>RemoveNodesTask</b> .
spec	Yes	<b>RemoveNodeSpec</b> object	Configuration information.
status	No	<b>TaskStatus</b> object	Job status.

**Table 4-420** RemoveNodesSpec

Parameter	Mandatory	Type	Description
login	Yes	<b>Login</b> object	Node login mode. Either the key pair or password must be used for login. <b>NOTE</b> This parameter cannot be modified when a node pool is being updated.
nodes	Yes	Array of <b>NodeItem</b> objects	List of nodes to be removed. A maximum of 200 nodes can be removed at a time.

**Table 4-421** Login

Parameter	Mandatory	Type	Description
sshKey	No	String	Name of the key pair used for login.
userPassword	No	<b>UserPassword</b> object	Password used for node login.

**Table 4-422** UserPassword

Parameter	Mandatory	Type	Description
username	No	String	Login account. The default value is <b>root</b> .

Parameter	Mandatory	Type	Description
password	Yes	String	<p>Login password. If a username and password are used to create a node, this field is shielded in the response body. A password must meet the following complexity requirements:</p> <ul style="list-style-type: none"> <li>• Contains 8 to 26 characters.</li> <li>• Contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters !@\$%^_-=+[{ }];,./?</li> <li>• Cannot contain the username or the username spelled backwards. The <b>password</b> field must be salted during node creation. For details, see <a href="#">Adding a Salt in the password Field When Creating a Node</a>.</li> </ul>

**Table 4-423** NodeItem

Parameter	Mandatory	Type	Description
uid	Yes	String	Node ID.

**Table 4-424** TaskStatus

Parameter	Mandatory	Type	Description
jobID	No	String	Job ID, which is used by the caller to query the job progress.

## Response Parameters

Status code: 200

**Table 4-425** Response body parameters

Parameter	Type	Description
apiVersion	String	API version. The value is fixed at <b>v3</b> .
kind	String	API type. The value is fixed at <b>RemoveNodesTask</b> .
spec	<b>RemoveNodeSpec</b> object	Configuration information.
status	<b>TaskStatus</b> object	Job status.

**Table 4-426** RemoveNodesSpec

Parameter	Type	Description
login	<b>Login</b> object	Node login mode. Either the key pair or password must be used for login. <b>NOTE</b> This parameter cannot be modified when a node pool is being updated.
nodes	Array of <b>NodeItem</b> objects	List of nodes to be removed. A maximum of 200 nodes can be removed at a time.

**Table 4-427** Login

Parameter	Type	Description
sshKey	String	Name of the key pair used for login.
userPassword	<b>UserPassword</b> object	Password used for node login.

**Table 4-428** UserPassword

Parameter	Type	Description
username	String	Login account. The default value is <b>root</b> .



Parameter	Type	Description
password	String	<p>Login password. If a username and password are used to create a node, this field is shielded in the response body. A password must meet the following complexity requirements:</p> <ul style="list-style-type: none"> <li>• Contains 8 to 26 characters.</li> <li>• Contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters !@\$%^-_=+[{ }];,./?</li> <li>• Cannot contain the username or the username spelled backwards. The <b>password</b> field must be salted during node creation. For details, see <a href="#">Adding a Salt in the password Field When Creating a Node</a>.</li> </ul>

**Table 4-429** NodeItem

Parameter	Type	Description
uid	String	Node ID.

**Table 4-430** TaskStatus

Parameter	Type	Description
jobID	String	Job ID, which is used by the caller to query the job progress.

## Example Requests

### Removing a Node

```
PUT /api/v3/projects/{project_id}/clusters/{cluster_id}/nodes/operation/remove
```

```
{
  "spec" : {
    "login" : {
      "sshKey" : "KeyPair-001"
    },
    "nodes" : [ {
      "uid" : "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx"
    }, {
      "uid" : "yyyyyyyy-yyyy-yyyy-yyyy-yyyyyyyyyyyy"
    } ]
  }
}
```

## Example Responses

### Status code: 200

The job for removing a node in a specified cluster is successfully delivered.

```
{
  "spec": {
    "login": {
      "sshKey": "KeyPair-001"
    },
    "nodes": [ {
      "uid": "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx"
    }, {
      "uid": "yyyyyyyy-yyyy-yyyy-yyyy-yyyyyyyyyy"
    } ]
  },
  "status": {
    "jobID": "2ec9b78d-9368-46f3-8f29-d1a95622a568"
  }
}
```

## Status Codes

Status Code	Description
200	The job for removing a node in a specified cluster is successfully delivered.

## Error Codes

See [Error Codes](#).

## 4.3.13 Migrating a Node

### Function

This API is used to migrate a node from a specified cluster to another cluster. The node can be migrated only between different clusters in the same VPC and project, and the cluster types before and after the migration must be the same.

#### NOTE

The URL for cluster management is in the format of **https://Endpoint/uri**, where **uri** indicates the resource path for API access.

### Constraints

- Data can be migrated only between clusters in the same VPC and project.
- Data cannot be migrated between CCE Turbo clusters and CCE standard clusters.
- Data cannot be migrated between DEC and non-DEC clusters.
- Data cannot be migrated between clusters with IPv6 enabled and clusters without IPv6 enabled.

- Data can be migrated only to the default node pool of the target cluster.

## URI

PUT /api/v3/projects/{project\_id}/clusters/{cluster\_id}/nodes/operation/migrateto/{target\_cluster\_id}

**Table 4-431** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
target_cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

## Request Parameters

**Table 4-432** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

**Table 4-433** Request body parameters

Parameter	Mandatory	Type	Description
apiVersion	No	String	API version. The value is fixed at <b>v3</b> .

Parameter	Mandatory	Type	Description
kind	No	String	API type. The value is fixed at <b>MigrateNodesTask</b> .
spec	Yes	<b>MigrateNodeSpec</b> object	Configuration data.
status	No	<b>TaskStatus</b> object	Job status.

**Table 4-434** MigrateNodesSpec

Parameter	Mandatory	Type	Description
os	Yes	String	Operating system type, which must be accurate to the version number. When <b>alpha.cce/NodeImageID</b> is specified, the value of <b>os</b> must be the same as the OS of the custom image.
extendParam	No	<b>MigrateNodeExtendParam</b> object	Extended parameters for node migration.
login	Yes	<b>Login</b> object	Node login mode. Either the key pair or password must be used for login. <b>NOTE</b> This parameter cannot be modified when a node pool is being updated.
runtime	No	<b>Runtime</b> object	Container runtime
nodes	Yes	Array of <b>NodeItem</b> objects	List of nodes to be migrated. A maximum of 200 nodes can be migrated at a time.

**Table 4-435** MigrateNodeExtendParam

Parameter	Mandatory	Type	Description
maxPods	No	Integer	Maximum number of pods that can be created on a node, including the default system pods. Value range: 16 to 256. This limit prevents the node from being overloaded of pods.

Parameter	Mandatory	Type	Description
DockerLVMConfigOverride	No	String	<p>Docker data disk configurations.</p> <p>The disk type of the nodes to be migrated must be the same as that specified during node creation (that is, the value of <code>diskType</code> in <code>DockerLVMConfigOverride</code> must be the same as that specified during node creation). Ensure that the disk types of the nodes selected for a single API call are the same.</p> <p>The following is the default configuration:</p> <pre>"DockerLVMConfigOverride": "dockerThinpool=vgpaas/90%VG;kubernetesLV=vgpaas/10%VG;diskType=evs;lvType=linear"</pre> <p>The configuration contains the following fields:</p> <ul style="list-style-type: none"> <li>• <code>userLV</code> (optional): size of the user space, for example, <code>vgpaas/20%VG</code>.</li> <li>• (Optional) <code>userPath</code>: mount path of the user space, for example, <code>/home/wqt-test</code>.</li> <li>• <code>diskType</code>: specifies the disk type. Currently, only <code>evs</code>, <code>hdd</code>, and <code>ssd</code> are supported.</li> <li>• <code>lvType</code>: type of the logical volume. Currently, the value can be <code>linear</code> or <code>striped</code>. For example, <code>striped</code>.</li> <li>• <code>dockerThinpool</code>: Docker disk size, for example, <code>vgpaas/60%VG</code>.</li> <li>• <code>kubernetesLV</code>: Kubelet space size, for example, <code>vgpaas/20%VG</code>.</li> </ul>

Parameter	Mandatory	Type	Description
alpha.cce/ preInstall	No	String	Pre-installation scripts. The total number of characters after transcoding cannot exceed 10240. <b>NOTE</b> The input value must be encoded using Base64. The method is as follows: <pre>echo -n "Content to be encoded"   base64</pre>
alpha.cce/ postInstall	No	String	Post-installation scripts. The total number of characters after transcoding cannot exceed 10240. <b>NOTE</b> The input value must be encoded using Base64. The method is as follows: <pre>echo -n "Content to be encoded"   base64</pre>
alpha.cce/ NodeImageID	No	String	ID of the user image to run the target OS. When <b>alpha.cce/NodeImageID</b> is specified, the value of <b>os</b> must be the same as the OS of the custom image.

**Table 4-436** Login

Parameter	Mandatory	Type	Description
sshKey	No	String	Name of the key pair used for login.
userPassword	No	<b>UserPassword</b> object	Password used for node login.

**Table 4-437** UserPassword

Parameter	Mandatory	Type	Description
username	No	String	Login account. The default value is <b>root</b> .

Parameter	Mandatory	Type	Description
password	Yes	String	<p>Login password. If a username and password are used to create a node, this field is shielded in the response body. A password must meet the following complexity requirements:</p> <ul style="list-style-type: none"> <li>• Contains 8 to 26 characters.</li> <li>• Contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters !@\$%^_-=+[{ }];,./?</li> <li>• Cannot contain the username or the username spelled backwards. The <b>password</b> field must be salted during node creation. For details, see <a href="#">Adding a Salt in the password Field When Creating a Node</a>.</li> </ul>

**Table 4-438** Runtime

Parameter	Mandatory	Type	Description
name	No	String	<p>Container runtime. Default value:</p> <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>



**Table 4-439** NodeItem

Parameter	Mandatory	Type	Description
uid	Yes	String	Node ID.

**Table 4-440** TaskStatus

Parameter	Mandatory	Type	Description
jobID	No	String	Job ID, which is used by the caller to query the job progress.

## Response Parameters

Status code: 200

**Table 4-441** Response body parameters

Parameter	Type	Description
apiVersion	String	API version. The value is fixed at <b>v3</b> .
kind	String	API type. The value is fixed at <b>MigrateNodesTask</b> .
spec	<b>MigrateNodeSpec</b> object	Configuration data.
status	<b>TaskStatus</b> object	Job status.

**Table 4-442** MigrateNodesSpec

Parameter	Type	Description
os	String	Operating system type, which must be accurate to the version number. When <b>alpha.cce/NodeImageID</b> is specified, the value of <b>os</b> must be the same as the OS of the custom image.
extendParam	<b>MigrateNodeExtendParam</b> object	Extended parameters for node migration.

Parameter	Type	Description
login	Login object	Node login mode. Either the key pair or password must be used for login. <b>NOTE</b> This parameter cannot be modified when a node pool is being updated.
runtime	Runtime object	Container runtime
nodes	Array of NodeItem objects	List of nodes to be migrated. A maximum of 200 nodes can be migrated at a time.

**Table 4-443** MigrateNodeExtendParam

Parameter	Type	Description
maxPods	Integer	Maximum number of pods that can be created on a node, including the default system pods. Value range: 16 to 256. This limit prevents the node from being overloaded of pods.
DockerLVMConfigOverride	String	<p>Docker data disk configurations.</p> <p>The disk type of the nodes to be migrated must be the same as that specified during node creation (that is, the value of diskType in DockerLVMConfigOverride must be the same as that specified during node creation). Ensure that the disk types of the nodes selected for a single API call are the same.</p> <p>The following is the default configuration: "DockerLVMConfigOverride":"dockerThinpool=vgpaas/90%VG;kubernetesLV=vgpaas/10%VG;diskType=evs;lvType=linear"</p> <p>The configuration contains the following fields:</p> <ul style="list-style-type: none"> <li>• userLV (optional): size of the user space, for example, vgpaas/20%VG.</li> <li>• (Optional) userPath: mount path of the user space, for example, /home/wqt-test.</li> <li>• diskType: specifies the disk type. Currently, only evs, hdd, and ssd are supported.</li> <li>• lvType: type of the logical volume. Currently, the value can be linear or striped. For example, striped.</li> <li>• dockerThinpool: Docker disk size, for example, vgpaas/60%VG.</li> <li>• kubernetesLV: Kubelet space size, for example, vgpaas/20%VG.</li> </ul>

Parameter	Type	Description
alpha.cce/ preInstall	String	Pre-installation scripts. The total number of characters after transcoding cannot exceed 10240. <b>NOTE</b> The input value must be encoded using Base64. The method is as follows: <code>echo -n "Content to be encoded"   base64</code>
alpha.cce/ postInstall	String	Post-installation scripts. The total number of characters after transcoding cannot exceed 10240. <b>NOTE</b> The input value must be encoded using Base64. The method is as follows: <code>echo -n "Content to be encoded"   base64</code>
alpha.cce/ NodeImageID	String	ID of the user image to run the target OS. When <b>alpha.cce/NodeImageID</b> is specified, the value of <b>os</b> must be the same as the OS of the custom image.

**Table 4-444** Login

Parameter	Type	Description
sshKey	String	Name of the key pair used for login.
userPassword	<b>UserPassword</b> object	Password used for node login.

**Table 4-445** UserPassword

Parameter	Type	Description
username	String	Login account. The default value is <b>root</b> .

Parameter	Type	Description
password	String	<p>Login password. If a username and password are used to create a node, this field is shielded in the response body. A password must meet the following complexity requirements:</p> <ul style="list-style-type: none"> <li>• Contains 8 to 26 characters.</li> <li>• Contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters !@\$%^-_=+[{ }];:./?</li> <li>• Cannot contain the username or the username spelled backwards. The <b>password</b> field must be salted during node creation. For details, see <a href="#">Adding a Salt in the password Field When Creating a Node</a>.</li> </ul>

**Table 4-446** Runtime

Parameter	Type	Description
name	String	<p>Container runtime. Default value:</p> <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>

**Table 4-447** NodeItem

Parameter	Type	Description
uid	String	Node ID.

**Table 4-448** TaskStatus

Parameter	Type	Description
jobID	String	Job ID, which is used by the caller to query the job progress.

## Example Requests

- Migrate a node running EulerOS 2.5 to another cluster.

```
PUT /api/v3/projects/{project_id}/clusters/{cluster_id}/nodes/operation/migrateto/{target_cluster_id}
```

```
{
  "spec" : {
    "os" : "EulerOS 2.5",
    "login" : {
      "sshKey" : "KeyPair-001"
    },
    "nodes" : [ {
      "uid" : "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx"
    }, {
      "uid" : "yyyyyyyy-yyyy-yyyy-yyyy-yyyyyyyyyyyy"
    } ]
  }
}
```

- Migrate a node running EulerOS 2.5 and created using a specified image ID to another cluster.

```
PUT /api/v3/projects/{project_id}/clusters/{cluster_id}/nodes/operation/migrateto/{target_cluster_id}
```

```
{
  "spec" : {
    "os" : "EulerOS 2.5",
    "extendParam" : {
      "alpha.cce/NodeImageID" : "cc697ad7-9563-11e8-8ea7-0255ac106311"
    },
    "login" : {
      "sshKey" : "KeyPair-001"
    },
    "nodes" : [ {
      "uid" : "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx"
    }, {
      "uid" : "yyyyyyyy-yyyy-yyyy-yyyy-yyyyyyyyyyyy"
    } ]
  }
}
```

## Example Responses

**Status code: 200**

The job for migrating a node from a specified cluster to another cluster is successfully delivered.

```
{
  "spec" : {
    "os" : "EulerOS 2.5",
    "login" : {
      "sshKey" : "KeyPair-001"
    },
    "runtime" : {
      "name" : "docker"
    },
    "nodes" : [ {
      "uid" : "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx"
    }, {
      "uid" : "yyyyyyyy-yyyy-yyyy-yyyy-yyyyyyyyyyyy"
    } ]
  },
  "status" : {
    "jobID" : "2ec9b78d-9368-46f3-8f29-d1a95622a568"
  }
}
```

## Status Codes

Status Code	Description
200	The job for migrating a node from a specified cluster to another cluster is successfully delivered.

## Error Codes

See [Error Codes](#).

# 4.4 Node Pool Management

## 4.4.1 Creating a Node Pool

### Function

This API is used to create a node pool in a specified cluster. This API can be called only when the cluster is in the available, scaling-out, or scaling-in state.

When creating a node pool in a CCE Turbo cluster of v1.21 or earlier, you can bind security groups to the node pool. A maximum of five security groups can be bound to a node pool.

After the security groups of a node pool are updated, the update takes effect only for newly created pods. You are advised to evict the original pods on the node.

#### NOTE

If there is no cluster, [create one](#). The URL for cluster management is in the format of **https://Endpoint/uri**, where **uri** specifies the resource path for API access.

## URI

POST /api/v3/projects/{project\_id}/clusters/{cluster\_id}/nodepools

**Table 4-449** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

## Request Parameters

**Table 4-450** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

**Table 4-451** Request body parameters

Parameter	Mandatory	Type	Description
kind	Yes	String	API type. The value is fixed at <b>NodePool</b> .
apiVersion	Yes	String	API version. The value is fixed at <b>v3</b> .
metadata	Yes	<a href="#">NodePoolMetadata</a> object	Metadata information of the node pool
spec	Yes	<a href="#">NodePoolSpec</a> object	Node pool specifications. <b>NOTE</b> <ul style="list-style-type: none"> <li>This field is optional for updating a node pool.</li> </ul>
status	No	<a href="#">NodePoolStatus</a> object	Node pool status

**Table 4-452** NodePoolMetadata

Parameter	Mandatory	Type	Description
name	Yes	String	Node pool name. <b>NOTE</b> Naming rules: Enter 1 to 50 characters, starting with a lowercase letter and not ending with a hyphen (-). Only lowercase letters, digits, and hyphens (-) are allowed. <ul style="list-style-type: none"> <li>You cannot create node pools named DefaultPool.</li> </ul>
uid	No	String	UID of the node pool. The value is automatically generated after the object is updated. A user-defined value will not take effect.
annotations	No	Map<String,String>	Annotations of a node pool in key-value pairs. This parameter is used only for query and cannot be input through a request. Entered data of this parameter is invalid.
updateTimestamp	No	String	Update time.
creationTimestamp	No	String	Creation time.

**Table 4-453** NodePoolSpec

Parameter	Mandatory	Type	Description
type	No	String	Node pool type. If this parameter is left blank, the value <b>vm</b> is used by default. <ul style="list-style-type: none"> <li><b>vm</b>: ECS</li> <li><b>ElasticBMS</b>: C6 general computing-plus BMS. An example flavor is <b>c6.22xlarge.2.physical</b>.</li> <li><b>bms</b>: BMS</li> </ul>
nodeTemplate	Yes	<b>NodeSpec</b> object	Detailed parameters of the node pool template.



Parameter	Mandatory	Type	Description
initialNodeCount	No	Integer	Initial number of nodes for the node pool. When queried, the value is the number of target nodes in the node pool.
autoscaling	No	<a href="#">NodePoolNodeAutoscaling</a> object	Auto scaling parameters
nodeManagement	No	<a href="#">NodeManagement</a> object	Node management configuration
podSecurityGroups	No	Array of <a href="#">SecurityID</a> objects	Security group configuration. This parameter is available only for CCE Turbo clusters.
extensionScaleGroups	No	Array of <a href="#">ExtensionScaleGroup</a> objects	Configuration list of the extended scaling groups in a node pool. For details, see the definition of <b>ExtensionScaleGroup</b> .
customSecurityGroups	No	Array of strings	Custom security group settings for a node pool. New nodes scaled out in a node pool can be bound to a specified security group. <ul style="list-style-type: none"> <li>• Specifying no security group ID will add the new nodes to the default security group of the worker nodes.</li> <li>• Specifying a valid security group ID will put new nodes in that security group.</li> <li>• When specifying a security group, do not modify the rules of the port on which CCE running depends.</li> </ul>

**Table 4-454** NodeSpec

Parameter	Mandatory	Type	Description
flavor	Yes	String	Node specifications. For details about the node specifications supported by CCE, see <a href="#">Node Flavor Description</a> .
az	Yes	String	AZ where the node to be created is located. You need to specify the AZ name. A node created using APIs does not support random AZs. <b>NOTE</b> <ul style="list-style-type: none"> <li>When you are creating a node pool and configuring a scaling group, this parameter cannot be set to <b>random</b>.</li> </ul>
os	No	String	Node OS. For details about the supported OSs, see <a href="#">Node OS</a> . <b>NOTE</b> <ul style="list-style-type: none"> <li>The system automatically selects the supported OS based on the cluster version. If the current cluster version does not support the OS, an error will be reported.</li> <li>If <b>alpha.cce/NodeImageID</b> in <b>extendParam</b> is specified during node creation, you do not need to configure this parameter.</li> <li>This parameter is mandatory when creating a node pool.</li> <li>If you are using shared disk space when creating a node, this parameter is mandatory. Specifically, the disk initialization configuration management parameter <b>storage</b> is used and the <b>name</b> field of <b>virtualSpaces</b> in <b>StorageGroups</b> is set to <b>share</b>.</li> </ul>
login	Yes	<a href="#">Login</a> object	Node login mode. Either the key pair or password must be used for login. <b>NOTE</b> This parameter cannot be modified when a node pool is being updated.

Parameter	Mandatory	Type	Description
rootVolume	Yes	<a href="#">Volume</a> object	Information about disks on the node
dataVolumes	Yes	Array of <a href="#">Volume</a> objects	Data disk parameters of a node. A VM can have up to 16 disks attached, while a BMS can only have up to 10 disks attached. The maximum number of disks that can be attached to either a VM or BMS is limited by their respective maximum disk limits. (You can add more data disks to each of your nodes through the CCE console and APIs.) The data disk used by the container runtime and kubelet cannot be detached, or the node will become unavailable. For DeC nodes, the parameter details are the same as those for <b>rootVolume</b> .
storage	No	<a href="#">Storage</a> object	Disk initialization management parameter. This parameter is complex to configure. For details, see <a href="#">Attaching Disks to a Node</a> . <b>NOTE</b> <ul style="list-style-type: none"> <li>If a node has both local and EVS disks attached, make sure to not keep the default value for this parameter, or it may result in unexpected disk partitions.</li> </ul> <b>NOTE</b> <ul style="list-style-type: none"> <li>If you want to change the value range of a data disk to 20 to 32768, do not use the default value.</li> </ul> <b>NOTE</b> <ul style="list-style-type: none"> <li>If you want to use the shared disk space (with the runtime and Kubernetes partitions cancelled), do not retain the default value of this parameter. For details about the shared disk space, see <a href="#">Data Disk Space Allocation</a>.</li> </ul>

Parameter	Mandatory	Type	Description
publicIP	No	<b>NodePublicIP</b> object	EIP of a node. <b>NOTE</b> This parameter is not supported when you add a node to a node pool.
nodeNicSpec	No	<b>NodeNicSpec</b> object	NIC of the node
count	No	Integer	Number of nodes to be created in a batch. The value must be a positive integer greater than or equal to 1 and less than or equal to the defined limit. This parameter can be left blank when it is used for a node pool.
billingMode	No	Integer	Node billing mode. <ul style="list-style-type: none"> <li>• <b>0</b>: pay-per-use</li> </ul>

Parameter	Mandatory	Type	Description
taints	No	Array of <b>Taint</b> objects	<p>You can add taints to created nodes to configure anti-affinity. A maximum of 20 taints can be added. Each taint contains the following parameters:</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain can be prefixed to a key.</li> <li>• <b>Value:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed.</li> <li>• <b>Effect:</b> Available options are <b>NoSchedule</b>, <b>PreferNoSchedule</b>, and <b>NoExecute</b>. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</li> </ul> <p>Example:</p> <pre>"taints": [{   "key": "status",   "value": "unavailable",   "effect": "NoSchedule" }, {   "key": "looks",   "value": "bad",   "effect": "NoSchedule" }]</pre>

Parameter	Mandatory	Type	Description
k8sTags	No	Map<String,String>	<p>Defined in key-value pairs. A maximum of 20 key-value pairs are allowed.</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain containing a maximum of 253 characters (such as <b>example.com/my-key</b>) can be prefixed to a key.</li> <li>• <b>Value:</b> The value can be left blank or contain 1 to 63 characters that start and end with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</li> </ul> <p>Example:</p> <pre>"k8sTags": {   "key": "value" }</pre>
ecsGroupId	No	String	<p>ECS group ID. If this parameter is configured, nodes will be created in the specified ECS group.</p> <p><b>NOTE</b> This configuration does not take effect when a node pool is created. To ensure that all nodes in a node pool are in the same ECS group, configure <b>nodeManagement</b> in the node pool.</p>

Parameter	Mandatory	Type	Description
dedicatedHostId	No	String	<p>ID of the DeH to which the node is scheduled.</p> <p><b>NOTE</b> This parameter is not supported when you add a node during node pool creation.</p>
userTags	No	Array of <b>UserTag</b> objects	<p>Cloud server tag (resource tag). The key of a tag must be unique. The maximum number of custom tags supported by CCE varies depending on regions and cannot exceed 8. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</p> <p><b>NOTE</b> A tag key can have a maximum of 36 characters. It can only contain letters, digits, Unicode characters, and special characters (-).</p>
runtime	No	<b>Runtime</b> object	<p>Container runtime. Default value:</p> <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>

Parameter	Mandatory	Type	Description
initializedConditions	No	Array of strings	<p>Custom initialization flag, which is left blank by default. Before CCE nodes are initialized, <b>node.cloudprovider.kubernetes.io/uninitialized</b> is added to the nodes to prevent pods from being scheduled to these nodes. When creating a node, you can configure the <b>initializedConditions</b> parameter to control the taint removal time. By default, the timeout period is not configured.</p> <p>Example:</p> <ol style="list-style-type: none"> <li>1. Create a node and add parameter <code>"initializedConditions": ["CCEInitial", "CustomedInitial"]</code>.</li> <li>2. After custom initialization is complete, CCE will call a Kubernetes API (for example, <b>PATCH /v1/nodes/{node_ip}/status</b>) to update the node <b>conditions</b> by adding two labels of types <b>CCEInitial</b> and <b>CustomedInitial</b> and setting the status to <b>True</b>, as shown in the following:</li> </ol> <pre>status:   conditions:   - type: CCEInitial     status: 'True'   - type: CustomedInitial     status: 'True'</pre> <ol style="list-style-type: none"> <li>1. CCE polls <b>status.Conditions</b> of nodes to check whether there are conditions of types <b>CCEInitial</b> and <b>CustomedInitial</b>. If such conditions exist and the status is <b>True</b>, the node initialization is complete and the initialization taint is removed.</li> </ol>



Parameter	Mandatory	Type	Description
			<p>2. <b>initializedConditions</b> allows you to configure a timeout period during node creation. For example, "initializedConditions": ["CCEInitial:15m", "CustomedInitial:15m"] indicates that the timeout period is 15 minutes. After the timeout period is reached, the initialization conditions will be automatically ignored and the initialization taint will be removed when CCE polls the node.</p> <ul style="list-style-type: none"> <li>• The timeout interval in <b>initializedConditions</b> ranges from 1 to 99 seconds.</li> <li>• Use only letters and digits. Max. characters: 20.</li> <li>• Max. flags: 2.</li> <li>• The unit of the timeout period is minute (m).</li> </ul>
extendParam	No	<b>NodeExtendParam</b> object	Extended parameters for creating a node.
hostnameConfig	No	<b>HostnameConfig</b> object	Kubernetes node name configuration parameter, which is supported by clusters of v1.23.6-r0 to v1.25 or clusters of v1.25.2-r0 or later versions.

Parameter	Mandatory	Type	Description
serverEnterpriseProjectID	No	String	Enterprise project ID of a server. CCE does not support features provided by EPS. This field is used only for synchronizing enterprise project IDs of servers. To create a node or node pool, you can specify an existing enterprise project for it. If this field is left blank, it inherits the enterprise project attributes of the cluster that runs the node or node pool. When updating a node pool, any configuration modifications you made will apply only to new nodes. You need to modify the enterprise projects of existing nodes on the EPS console. If no value is specified during the update, this field will not be updated. If this parameter is left blank, the cluster enterprise project is returned.

**Table 4-455** Login

Parameter	Mandatory	Type	Description
sshKey	No	String	Name of the key pair used for login.
userPassword	No	UserPassword object	Password used for node login.

**Table 4-456** UserPassword

Parameter	Mandatory	Type	Description
username	No	String	Login account. The default value is <b>root</b> .

Parameter	Mandatory	Type	Description
password	Yes	String	<p>Login password. If a username and password are used to create a node, this field is shielded in the response body. A password must meet the following complexity requirements:</p> <ul style="list-style-type: none"> <li>• Contains 8 to 26 characters.</li> <li>• Contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters !@\$%^_-=+[{ }];,./?</li> <li>• Cannot contain the username or the username spelled backwards. The <b>password</b> field must be salted during node creation. For details, see <a href="#">Adding a Salt in the password Field When Creating a Node</a>.</li> </ul>

**Table 4-457** Volume

Parameter	Mandatory	Type	Description
size	Yes	Integer	<p>Disk size, in GiB.</p> <ul style="list-style-type: none"> <li>• System disk: 40 to 1024</li> <li>• Value range for data disks: 100 to 32768</li> </ul>
volumetype	Yes	String	<p>Disk type. For details about possible values, see the description of the <b>root_volume</b> parameter in the API used to create an ECS.</p> <ul style="list-style-type: none"> <li>• <b>SAS</b>: high I/O SAS disk</li> <li>• <b>SSD</b>: ultra-high I/O SSD disk</li> <li>• <b>SATA</b>: common I/O SATA disk SATA disks have been removed from EVS. You can find them attached only on existing nodes.</li> </ul>

Parameter	Mandatory	Type	Description
extendParam	No	Map<String,Object>	Extended disk parameters, defined in <b>extendparam</b> in the API used to create an ECS.
cluster_id	No	String	ID of the storage pool used by the ECS system disk. This field is used only for DeC clusters, which functions as <b>dssPoolID</b> , that is, the ID of the DSS storage pool.
cluster_type	No	String	Storage class of the cloud server system disk. The value is always <b>dss</b> . This field is used only for DeC clusters.
hw:passthrough	No	Boolean	<ul style="list-style-type: none"> <li>Pay attention to this field if your ECS is SDI-compliant. If the value of this field is <b>true</b>, the created disk is of SCSI type.</li> <li>If the node pool type is <b>ElasticBMS</b>, this field must be set to <b>true</b>.</li> <li>If a node specification involves both local and EVS disks, configure the disk initialization parameters. For details, see <a href="#">Attaching Disks to a Node</a>.</li> </ul>
metadata	No	<b>VolumeMetadata</b> object	EVS disk encryption information. This field is mandatory only when you need to encrypt the system disk or data disks of the node to be created.

**Table 4-458** VolumeMetadata

Parameter	Mandatory	Type	Description
<code>__system__encrypted</code>	No	String	Whether the EVS disk is encrypted. The value <b>0</b> indicates that the EVS disk is not encrypted, and the value <b>1</b> indicates that the EVS disk is encrypted.  If this parameter is not specified, EVS disks will not be encrypted by default.
<code>__system__cmkid</code>	No	String	CMK ID, which indicates encryption in <b>metadata</b> . This field is used with <b>__system__encrypted</b> .

**Table 4-459** Storage

Parameter	Mandatory	Type	Description
<code>storageSelectors</code>	Yes	Array of <b>StorageSelectors</b> objects	Disk selection. Matched disks are managed according to <b>matchLabels</b> and <b>storageType</b> .
<code>storageGroups</code>	Yes	Array of <b>StorageGroups</b> objects	A storage group consists of multiple storage devices. It is used to divide storage space.

**Table 4-460** StorageSelectors

Parameter	Mandatory	Type	Description
<code>name</code>	Yes	String	Selector name, used as the index of <b>selectorNames</b> in <b>storageGroup</b> . Therefore, the name of each selector must be unique.
<code>storageType</code>	Yes	String	Storage type. Only <b>evs</b> (EVS disks) and <b>local</b> (local disks) are supported. The local storage does not support disk selection. All local disks will form a VG, so only one storageSelector of the local type is allowed.

Parameter	Mandatory	Type	Description
matchLabels	No	<a href="#">matchLabels</a> object	Matching field of an EVS disk. The value can be <b>size</b> , <b>volumeType</b> , <b>metadataEncrypted</b> , <b>metadataCmkid</b> , or <b>count</b> in <b>DataVolume</b> .

**Table 4-461** matchLabels

Parameter	Mandatory	Type	Description
size	No	String	Matched disk size. If this parameter is left unspecified, the disk size is not limited. Example: 100
volumeType	No	String	EVS disk type. <b>SSD</b> , <b>GPSSD</b> , <b>SAS</b> , <b>ESSD</b> , and <b>SATA</b> are supported. If this parameter is left blank, there is no restriction on the disk type.
metadataEncrypted	No	String	Disk encryption identifier. If this parameter is set to <b>0</b> , it specifies that the disk is not encrypted. If this parameter is set to <b>1</b> , it specifies that the disk is encrypted. If this parameter is left blank, there is no restriction on the disk encryption identifier.
metadataCmkid	No	String	CMK ID of an encrypted disk. The value is a string of 36 bytes. If this parameter is left blank, there is no restriction on the disk key ID.
count	No	String	Number of disks to be selected. If this parameter is left blank, all disks of this type are selected.

**Table 4-462** StorageGroups

Parameter	Mandatory	Type	Description
name	Yes	String	Name of a virtual storage group, which must be unique. <b>NOTE</b> <ul style="list-style-type: none"> <li>If <b>cceManaged</b> is set to <b>true</b>, the name must be set to <b>vgpaas</b>.</li> <li>If the data disk is used as a temporary storage volume, the name must be <b>vg-everest-localvolume-ephemeral</b>.</li> <li>If the data disk is used as a persistent storage volume, the name must be <b>vg-everest-localvolume-persistent</b>.</li> </ul>
cceManaged	No	Boolean	Storage space for Kubernetes and runtime components. Only one group can be set to <b>true</b> . If this parameter is left blank, the default value <b>false</b> is used.
selectorNames	Yes	Array of strings	Corresponds to <b>name</b> in <b>storageSelectors</b> . One group can select multiple selectors, but one selector can be selected by only one group.
virtualSpaces	Yes	Array of <b>VirtualSpace</b> objects	Detailed management of space configuration in a group.

**Table 4-463** VirtualSpace

Parameter	Mandatory	Type	Description
name	Yes	String	<p>Name of a virtualSpace. Options: <b>share</b>, <b>kubernetes</b>, <b>runtime</b>, and <b>user</b></p> <ul style="list-style-type: none"> <li>• <b>share</b>: shared space configuration (with the runtime and Kubernetes partitions cancelled). <b>lvmConfig</b> needs to be configured.</li> <li>• <b>Kubernetes</b>: Kubernetes space configuration. <b>lvmConfig</b> needs to be configured.</li> <li>• <b>runtime</b>: runtime space configuration. <b>runtimeConfig</b> needs to be configured.</li> <li>• <b>user</b>: user space configuration. <b>lvmConfig</b> needs to be configured.</li> </ul>
size	Yes	String	<p>Size of a virtualSpace. The value must be an integer in percentage. Example: 90%.</p> <p><b>NOTE</b> The value ranges from 10 to 90. The sum of percentages of all virtualSpaces in a group cannot exceed 100%.</p>
lvmConfig	No	<b>LVMConfig</b> object	<p>LVM configuration managements, applicable to the <i>share</i>*, <b>kubernetes</b>, and <b>user</b> spaces. Note that one virtual space supports only one config.</p>
runtimeConfig	No	<b>RuntimeConfig</b> object	<p>runtime configurations, applicable to the <b>runtime</b> space. Note that one virtual space supports only one config.</p>



**Table 4-464** LVMConfig

Parameter	Mandatory	Type	Description
lvType	Yes	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.
path	No	String	Disk mount path. This parameter takes effect only in user configuration. The value is an absolute path. Digits, letters, periods (.), hyphens (-), and underscores (_) are allowed.  <b>NOTE</b> Do not leave the path empty or set it to a key OS path such as the root directory. Key OS paths that cannot be specified: /, /home, /home/, /bin, /bin/, /lib, /lib/, /root, /root/, /boot, /boot/, /dev, /dev/, /etc, /etc/, /lost+found, /lost+found/, /mnt, /mnt/, /proc, /proc/, /sbin, /sbin/, /srv, /srv/, /tmp, /tmp/, /var, /var/, /media, /media/, /opt, /opt/, /selinux, /selinux/, /sys, /sys/, /usr, /usr/, /opt/cloud/, /mnt/paas/, /home/paas/, /var/paas/, /var/lib/, /var/script/

**Table 4-465** RuntimeConfig

Parameter	Mandatory	Type	Description
lvType	Yes	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.

**Table 4-466** NodePublicIP

Parameter	Mandatory	Type	Description
ids	No	Array of strings	IDs of existing EIPs. The quantity cannot be greater than the number of nodes to be created. <b>NOTE</b> If <b>ids</b> has been set, you do not need to set <b>count</b> and <b>eip</b> .
count	No	Integer	Number of EIPs to be dynamically created. <b>NOTE</b> <b>count</b> and <b>eip</b> must be set at the same time.
eip	No	<b>NodeEIPSpec</b> object	EIP configuration. <b>NOTE</b> This parameter is not supported when creating a node pool.

**Table 4-467** NodeEIPSpec

Parameter	Mandatory	Type	Description
iptype	Yes	String	EIP type, specified in <b>publicip.type</b> in the API for assigning an EIP.
bandwidth	No	<b>NodeBandwidth</b> object	Bandwidth parameters of the EIP

**Table 4-468** NodeBandwidth

Parameter	Mandatory	Type	Description
chargemode	No	String	<p>Bandwidth billing mode.</p> <ul style="list-style-type: none"> <li>If this field is not specified, the billing is based on bandwidth.</li> <li>If the field is null, the billing is based on bandwidth.</li> <li>If the field value is <b>traffic</b>, the billing is based on traffic.</li> <li>If the value is out of the preceding options, the cloud server will fail to be created.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>Billed by bandwidth: The billing will be based on the data transmission rate (in Mbps) of public networks. This billing mode is recommended if your bandwidth usage is higher than 10%.</li> <li>Billed by traffic: The billing is based on the total amount of data (in GB) transmitted over the public network. This mode is available only when you are creating a pay-per-use node. This billing mode is recommended if your bandwidth usage is lower than 10%.</li> </ul>
size	No	Integer	Bandwidth size, specified in <b>bandwidth.size</b> in the API for assigning an EIP.
sharetype	No	String	Bandwidth sharing type. Value options: <b>PER</b> (exclusive bandwidth)

**Table 4-469** NodeNicSpec

Parameter	Mandatory	Type	Description
primaryNic	No	<a href="#">NicSpec</a> object	Description of the primary NIC.

Parameter	Mandatory	Type	Description
extNics	No	Array of <b>NicSpec</b> objects	Extension NIC <b>NOTE</b> This parameter is not supported when you add a node to a node pool.

**Table 4-470** NicSpec

Parameter	Mandatory	Type	Description
subnetId	No	String	Network ID of the subnet to which a NIC belongs. If <b>subnetId</b> is not specified when a primary NIC is creating, the cluster subnet will be used. If <b>subnetList</b> is also configured for a node pool, the <b>subnetList</b> field is used for adding subnets to the node pool. When creating an extension NIC, you must specify <b>subnetId</b> .
fixedIps	No	Array of strings	The IP address of the primary ENI is specified using <b>fixedIps</b> . The number of IP addresses cannot be greater than the number of created nodes. Either <b>fixedIps</b> or <b>ipBlock</b> can be specified. <b>fixedIps</b> cannot be specified for extension ENIs.
ipBlock	No	String	CIDR format of the primary NIC IP range. The IP address of the created node falls in this range. <b>fixedIps</b> and <b>ipBlock</b> cannot be specified at the same time.
subnetList	No	Array of strings	Network ID list of the subnet where the NIC resides. A maximum of 20 subnets can be configured for a node pool.

**Table 4-471** Taint

Parameter	Mandatory	Type	Description
key	Yes	String	Key.
value	No	String	Value.
effect	Yes	String	Effect.

**Table 4-472** UserTag

Parameter	Mandatory	Type	Description
key	No	String	Key of the cloud server label. The value cannot start with <b>CCE-</b> or <b>__type_baremetal</b> .
value	No	String	Value of the cloud server label.

**Table 4-473** Runtime

Parameter	Mandatory	Type	Description
name	No	String	Container runtime. Default value: <ul style="list-style-type: none"> <li>Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>

**Table 4-474** NodeExtendParam

Parameter	Mandatory	Type	Description
ecs:performancetype	No	String	ECS flavor types. This field is returned in the response.

Parameter	Mandatory	Type	Description
maxPods	No	Integer	<p>Maximum number of pods that can be created on a node, including the default system pods. Value range: 16 to 256.</p> <p>This limit prevents the node from being overloaded of pods.</p> <p>The number of pods that can be created on a node is determined by multiple parameters. For details, see <a href="#">Maximum Number of Pods That Can Be Created on a Node</a>.</p>
DockerLVMConfigOverride	No	String	<p>Docker data disk configuration item. (This parameter has been discarded. Use the <b>storage</b> field instead.)</p> <p>Example default configuration:  <code>"DockerLVMConfigOverride":"dockerThinpool=vgpaas/90%VG;kubernetesLV=vgpaas/10%VG;diskType=evs;lvType=linear"</code></p> <p>By default, if no VD disk is available, an error occurs because the data disk fails to be found. Set <b>diskType</b> based on the actual drive letter type. The following fields are included:</p> <ul style="list-style-type: none"> <li>• <b>userLV</b> (optional): size of the user space, for example, <b>vgpaas/20%VG</b>.</li> <li>• <b>userPath</b> (optional): mount path of the user space, for example, <b>/home/wqt-test</b>.</li> <li>• <b>diskType</b>: disk type. Currently, only <b>evs</b>, <b>hdd</b>, and <b>ssd</b> are supported.</li> <li>• <b>lvType</b>: type of a logic volume. The value can be <b>linear</b> or <b>striped</b>.</li> <li>• <b>dockerThinpool</b>: Docker space size, for example, <b>vgpaas/60%VG</b>.</li> <li>• <b>kubernetesLV</b>: kubelet space size, for example, <b>vgpaas/20%VG</b>.</li> </ul>

Parameter	Mandatory	Type	Description
dockerBaseSize	No	Integer	<p>Available disk space of a single container on a node, in GB. This parameter has been discarded. Use <b>containerBaseSize</b>.</p> <p>If this parameter is left blank or is set to <b>0</b>, the default value is used. In Device Mapper mode, the default value is <b>10</b>. In OverlayFS mode, the available space of a single container is not limited by default, and the <b>dockerBaseSize</b> setting takes effect only on nodes running EulerOS in clusters of the new versions.</p> <p>For details about how to allocate the space for the container runtime, see <a href="#">Data Disk Space Allocation</a>.</p> <p>When Device Mapper is used, it is recommended that <b>dockerBaseSize</b> be set to a value less than or equal to 80 GB. If the value is too large, the container runtime may fail to be started due to long initialization. If there are special requirements for the container disk space, you can mount an external or local storage device.</p>

Parameter	Mandatory	Type	Description
containerBaseSize	No	Integer	<p>Available disk space of a single container on a node, in GiB.</p> <p>If this parameter is left blank or is set to <b>0</b>, the default value is used. In OverlayFS mode, the available space of a single container is not limited by default. In Device Mapper mode, the default value is <b>10</b>, and the <b>containerBaseSize</b> setting takes effect only on nodes running EulerOS in clusters v1.23.14-r0, v1.25.9-r0, v1.27.6-r0, v1.28.4-r0 or later.</p> <p>For details about how to allocate the space for the container runtime, see <a href="#">Data Disk Space Allocation</a>.</p> <p>When Device Mapper is used, it is recommended that <b>containerBaseSize</b> be set to a value less than or equal to 80 GB. If the value is too large, the container runtime may fail to be started due to long initialization. If there are special requirements for the container disk space, you can mount an external or local storage device. In new versions, the Device Mapper mode is used only by BMSs in the same resource pool and is being discarded.</p> <p><b>NOTE</b> This parameter cannot be updated when a node pool is being updated.</p>
publicKey	No	String	Public key of a node.



Parameter	Mandatory	Type	Description
alpha.cce/ preInstall	No	String	Pre-installation scripts. The total number of characters after transcoding cannot exceed 10240. <b>NOTE</b> The input value must be encoded using Base64. The method is as follows: <pre>echo -n "Content to be encoded"   base64</pre>
alpha.cce/ postInstall	No	String	Post-installation scripts. The total number of characters after transcoding cannot exceed 10240. <b>NOTE</b> The input value must be encoded using Base64. The method is as follows: <pre>echo -n "Content to be encoded"   base64</pre>
alpha.cce/ NodeImageID	No	String	This parameter is required when a custom image is used to create a BMS node.
chargingMode	No	Integer	Billing mode of a node. This parameter has been deprecated. Use the <b>billingMode</b> parameter in <b>NodeSpec</b> .
agency_name	No	String	Name of an agency An agency is created by a tenant administrator on Identity and Access Management (IAM) to provide temporary credentials for CCE nodes to access cloud servers. This parameter is returned only when it is transferred during node creation.
kubeReservedMem	No	Integer	Reserved node memory, which is reserved for Kubernetes components.
systemReservedMem	No	Integer	Reserved node memory, which is reserved for system components.

Parameter	Mandatory	Type	Description
init-node-password	No	String	Node password. If this parameter is used as a response parameter, asterisks (*) will be displayed.

**Table 4-475** HostnameConfig

Parameter	Mandatory	Type	Description
type	Yes	String	<p>Configuration type of the Kubernetes node name. The default value is <b>privatelp</b>.</p> <ul style="list-style-type: none"> <li>• <b>privatelp</b>: The Kubernetes node is named after its IP address.</li> <li>• <b>cceNodeName</b>: The Kubernetes node is named after the CCE node.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• For a node which is configured using <b>cceNodeName</b>, the name is the same as the Kubernetes node name and the ECS name. The node name cannot be changed. If the ECS name is changed on the ECS console, the node name will retain unchanged after ECS synchronization.</li> <li>• For a node which is configured using <b>cceNodeName</b>, to avoid a conflict between Kubernetes nodes, the system automatically adds a suffix to each node name. The suffix is in the format of <b>A hyphen (-) Five random characters</b>. The value of the random characters is a lowercase letter or a digit ranging from 0 to 9.</li> </ul>

**Table 4-476** NodePoolNodeAutoscaling

Parameter	Mandatory	Type	Description
enable	No	Boolean	Whether to enable auto scaling.
minNodeCount	No	Integer	Minimum number of nodes allowed if auto scaling is enabled. The value cannot be greater than the maximum number of nodes allowed by the cluster specifications.
maxNodeCount	No	Integer	Maximum number of nodes allowed if auto scaling is enabled. This value must be greater than or equal to the value of <b>minNodeCount</b> and cannot exceed the maximum number of nodes in the cluster specifications.
scaleDownCooldownTime	No	Integer	Interval between two scaling operations, in minutes. During this period, nodes added after a scale-up will not be deleted.
priority	No	Integer	Weight of a node pool. A node pool with a higher weight has a higher priority during scaling.

**Table 4-477** NodeManagement

Parameter	Mandatory	Type	Description
serverGroupReference	No	String	Cloud server group ID. If this field is specified, all nodes in the node pool will be created in this group. The group ID can be specified only when you create the node pool and cannot be modified. When you specify a cloud server group, the number of nodes in the node pool cannot exceed the group quota.

**Table 4-478** SecurityID

Parameter	Mandatory	Type	Description
id	No	String	Security group ID

**Table 4-479** ExtensionScaleGroup

Parameter	Mandatory	Type	Description
metadata	No	<a href="#">ExtensionScaleGroupMetadata</a> object	Basic information about an extended scaling group
spec	No	<a href="#">ExtensionScaleGroupSpec</a> object	Configurations of an extended scaling group, which carry different configurations from those of the default scaling group

**Table 4-480** ExtensionScaleGroupMetadata

Parameter	Mandatory	Type	Description
uid	No	String	UUID of an extended scaling group, which is automatically generated by the system.
name	No	String	Name of an extended scaling group. The value cannot be <b>default</b> and can contain a maximum of 55 characters. Only digits, lowercase letters, and hyphens (-) are allowed.

**Table 4-481** ExtensionScaleGroupSpec

Parameter	Mandatory	Type	Description
flavor	No	String	Node flavor
az	No	String	AZ of a node. If this parameter is not specified or left blank, the default scaling group configurations take effect.
capacityReservationSpecification	No	<a href="#">CapacityReservationSpecification</a> object	Capacity reservation configurations of an extended scaling group

Parameter	Mandatory	Type	Description
autoscaling	No	<a href="#">ScaleGroupAutoscaling</a> object	Auto scaling configurations of an extended scaling group

**Table 4-482** CapacityReservationSpecification

Parameter	Mandatory	Type	Description
id	No	String	Private pool ID. The parameter value can be ignored when <b>preference</b> is set to <b>none</b> .
preference	No	String	Capacity of a private storage pool. If the value is <b>none</b> , the capacity reservation is not specified. If the value is <b>targeted</b> , the capacity reservation is specified. In this case, the ID cannot be left blank.

**Table 4-483** ScaleGroupAutoscaling

Parameter	Mandatory	Type	Description
enable	No	Boolean	Whether to enable auto scaling for a scaling group. This function is disabled by default.
extensionPriority	No	Integer	Priority of a scaling group. If no value is specified for this parameter, it defaults to <b>0</b> . A higher value indicates a greater priority.
minNodeCount	No	Integer	Minimum number of nodes in a scaling group during auto scaling. The value must be greater than <b>0</b> .

Parameter	Mandatory	Type	Description
maxNodeCount	No	Integer	Maximum number of nodes that can be retained in a scaling group during auto scaling. The value must be greater than or equal to that of <b>minNodeCount</b> , and can neither be greater than the maximum number of nodes allowed by the cluster nor the maximum number of nodes in the node pool.

**Table 4-484** NodePoolStatus

Parameter	Mandatory	Type	Description
currentNode	No	Integer	Total number of nodes in the current node pool (excluding the nodes that are being deleted)
creatingNode	No	Integer	Number of nodes in the creation process in the node pool
deletingNode	No	Integer	Number of nodes being deleted in the current node pool.

Parameter	Mandatory	Type	Description
phase	No	String	<p>Node pool status.</p> <ul style="list-style-type: none"> <li>• Null: available (the number of current nodes in the node pool has reached the expected value, and no node scaling is being performed.)</li> <li>• Synchronizing: scaling in progress (the number of current nodes in the node pool does not reach the expected value and no node scaling is being performed.)</li> <li>• Synchronized: pending scaling (the number of current nodes in the node pool does not reach the expected value, or node scaling is being performed.)</li> <li>• SoldOut: The node pool cannot be scaled out. (This field is used in multiple scenarios, for example, resources in the node pool have been sold out or the resource quota is insufficient.)</li> </ul> <p><b>NOTE</b> This node pool status has been discarded and is reserved only for compatibility. You are not advised to use it. The replacements are as follows:</p> <ul style="list-style-type: none"> <li>• Node pool scaling status: You can obtain accurate status of the current node pool based on parameters such as <b>currentNode</b>, <b>creatingNode</b>, or <b>deletingNode</b>.</li> <li>• Node pool scale-out: You can use <b>conditions</b> to obtain the detailed status of a node pool. <b>Scalable</b> can replace <b>SoldOut</b>.</li> <li>• <b>Deleting</b>: The object is being deleted.</li> <li>• <b>Error</b>: An error occurs.</li> </ul>

Parameter	Mandatory	Type	Description
jobId	No	String	ID of a job executed on the node pool. This field is available only when the node pool is being deleted.
conditions	No	Array of <b>NodePoolCondition</b> objects	Node pool status details. For details, see the definition of <b>Condition</b> .
scaleGroupStatuses	No	Array of <b>ScaleGroupStatus</b> objects	Detailed status of a scaling group. For details, see the definition of <b>ScaleGroupStatus</b> .

**Table 4-485** ScaleGroupStatus

Parameter	Mandatory	Type	Description
name	No	String	Scaling group name
uid	No	String	UUID of a scaling group
creationTimestamp	No	String	The time when a scaling group was created
updateTimestamp	No	String	The time when a scaling group was updated



Parameter	Mandatory	Type	Description
phase	No	String	<p>Status of a scaling group.</p> <ul style="list-style-type: none"> <li>• <b>Null:</b> The scaling group is available. (The number of nodes in the current scaling group has reached the expected value, and no node scaling is being performed.)</li> <li>• <b>Synchronizing:</b> The scaling group is being scaled. (The number of nodes in the current scaling group has not reached the expected value, and no node scaling is being performed.)</li> <li>• <b>Synchronized:</b> The scaling is pending. (The number of nodes in the current scaling group has not reached the expected value, or node scaling is being performed.)</li> <li>• <b>SoldOut:</b> The scaling group cannot be scaled out. (This field is used in multiple scenarios, for example, resources in the node pool have been sold out or the resource quota is insufficient.)</li> </ul> <p><b>NOTE</b> This scaling group status has been deprecated. It is reserved only for compatibility. Do not use it anymore. Use the following instead:</p> <ul style="list-style-type: none"> <li>• Scaling status of a scaling group: You can obtain the status of the current scaling group using parameters such as <b>desiredNodeCount</b>, <b>existingNodeCount</b>, or <b>upcomingNodeCount</b>.</li> <li>• Scaling group scale-out: You can use <b>conditions</b> to obtain the detailed status of a scaling group. <b>Scalable</b> can replace <b>SoldOut</b>.</li> <li>• <b>Deleting:</b> The scaling group being deleted.</li> </ul>

Parameter	Mandatory	Type	Description
			<ul style="list-style-type: none"> <li><b>Error:</b> An error occurred in the scaling group.</li> </ul>
desiredNodeCount	No	Integer	Expected number of nodes in a scaling group
unpaidScaleNodeCount	No	Integer	Number of unpaid nodes in an order
existingNodeCount	No	<a href="#">existingNodeCount</a> object	Statistics on existing nodes in a scaling group
upcomingNodeCount	No	<a href="#">upcomingNodeCount</a> object	Statistics on the nodes to be created in a scaling group
scaleDownDisabledNodeCount	No	Integer	Number of nodes that cannot be deleted in a scaling group
conditions	No	Array of <a href="#">NodePoolCondition</a> objects	Scaling group status details. For details, see the definition of <b>Condition</b> .

**Table 4-486** existingNodeCount

Parameter	Mandatory	Type	Description
postPaid	No	Integer	Number of pay-per-use nodes
total	No	Integer	Total number of pay-per-use nodes

**Table 4-487** upcomingNodeCount

Parameter	Mandatory	Type	Description
postPaid	No	Integer	Number of pay-per-use nodes
total	No	Integer	Total number of pay-per-use nodes

**Table 4-488** NodePoolCondition

Parameter	Mandatory	Type	Description
type	No	String	Condition type. The options are as follows: <ul style="list-style-type: none"> <li>• <b>Scalable</b>: whether a node pool can be scaled. If the status is <b>False</b>, node pool scaling will not be triggered again.</li> <li>• <b>QuotaInsufficient</b>: Quotas on which node pool scaling depends are insufficient, affecting the node pool scaling status.</li> <li>• <b>ResourceInsufficient</b>: Resources on which node pool scaling depends are insufficient, affecting the node pool scaling status.</li> <li>• <b>UnexpectedError</b>: The node pool fails to be scaled out due to unexpected reasons, affecting the node pool scaling status.</li> <li>• <b>Error</b>: A node pool error occurs. A common trigger is deletion failure.</li> </ul>
status	No	String	Current status of <b>Condition</b> . The options are as follows: <ul style="list-style-type: none"> <li>• "True"</li> <li>• "False"</li> </ul>
lastProbeTime	No	String	Time when the status was last checked
lastTransitTime	No	String	Time when the status was last changed
reason	No	String	Reason why the status was last changed
message	No	String	Detailed condition description

## Response Parameters

Status code: 201

**Table 4-489** Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at <b>NodePool</b> .
apiVersion	String	API version. The value is fixed at <b>v3</b> .
metadata	<b>NodePoolMetadata</b> object	Metadata of the node pool.
spec	<b>NodePoolSpec</b> object	Node pool specifications. <b>NOTE</b> <ul style="list-style-type: none"> <li>This field is optional for updating a node pool.</li> </ul>
status	<b>CreateNodePoolStatus</b> object	Node pool status.

**Table 4-490** NodePoolMetadata

Parameter	Type	Description
name	String	Node pool name. <b>NOTE</b> Naming rules: Enter 1 to 50 characters, starting with a lowercase letter and not ending with a hyphen (-). Only lowercase letters, digits, and hyphens (-) are allowed. <ul style="list-style-type: none"> <li>You cannot create node pools named DefaultPool.</li> </ul>
uid	String	UID of the node pool. The value is automatically generated after the object is updated. A user-defined value will not take effect.
annotations	Map<String,String>	Annotations of a node pool in key-value pairs. This parameter is used only for query and cannot be input through a request. Entered data of this parameter is invalid.
updateTimestamp	String	Update time.
creationTimestamp	String	Creation time.

**Table 4-491** NodePoolSpec

Parameter	Type	Description
type	String	Node pool type. If this parameter is left blank, the value <b>vm</b> is used by default. <ul style="list-style-type: none"> <li>• <b>vm</b>: ECS</li> <li>• <b>ElasticBMS</b>: C6 general computing-plus BMS. An example flavor is <b>c6.22xlarge.2.physical</b>.</li> <li>• <b>bms</b>: BMS</li> </ul>
nodeTemplate	<b>NodeSpec</b> object	Detailed parameters of the node pool template.
initialNodeCount	Integer	Initial number of nodes for the node pool. When queried, the value is the number of target nodes in the node pool.
autoscaling	<b>NodePoolNodeAutoscaling</b> object	Auto scaling parameters
nodeManagement	<b>NodeManagement</b> object	Node management configuration
podSecurityGroups	Array of <b>SecurityID</b> objects	Security group configuration. This parameter is available only for CCE Turbo clusters.
extensionScaleGroups	Array of <b>ExtensionScaleGroup</b> objects	Configuration list of the extended scaling groups in a node pool. For details, see the definition of <b>ExtensionScaleGroup</b> .
customSecurityGroups	Array of strings	Custom security group settings for a node pool. New nodes scaled out in a node pool can be bound to a specified security group. <ul style="list-style-type: none"> <li>• Specifying no security group ID will add the new nodes to the default security group of the worker nodes.</li> <li>• Specifying a valid security group ID will put new nodes in that security group.</li> <li>• When specifying a security group, do not modify the rules of the port on which CCE running depends.</li> </ul>

**Table 4-492** NodeSpec

Parameter	Type	Description
flavor	String	Node specifications. For details about the node specifications supported by CCE, see <a href="#">Node Flavor Description</a> .
az	String	AZ where the node to be created is located. You need to specify the AZ name. A node created using APIs does not support random AZs. <b>NOTE</b> <ul style="list-style-type: none"> <li>When you are creating a node pool and configuring a scaling group, this parameter cannot be set to <b>random</b>.</li> </ul>
os	String	Node OS. For details about the supported OSs, see <a href="#">Node OS</a> . <b>NOTE</b> <ul style="list-style-type: none"> <li>The system automatically selects the supported OS based on the cluster version. If the current cluster version does not support the OS, an error will be reported.</li> <li>If <b>alpha.cce/NodeImageID</b> in <b>extendParam</b> is specified during node creation, you do not need to configure this parameter.</li> <li>This parameter is mandatory when creating a node pool.</li> <li>If you are using shared disk space when creating a node, this parameter is mandatory. Specifically, the disk initialization configuration management parameter <b>storage</b> is used and the <b>name</b> field of <b>virtualSpaces</b> in <b>StorageGroups</b> is set to <b>share</b>.</li> </ul>
login	<a href="#">Login</a> object	Node login mode. Either the key pair or password must be used for login. <b>NOTE</b> This parameter cannot be modified when a node pool is being updated.
rootVolume	<a href="#">Volume</a> object	Information about disks on the node

Parameter	Type	Description
dataVolumes	Array of <a href="#">Volume</a> objects	Data disk parameters of a node. A VM can have up to 16 disks attached, while a BMS can only have up to 10 disks attached. The maximum number of disks that can be attached to either a VM or BMS is limited by their respective maximum disk limits. (You can add more data disks to each of your nodes through the CCE console and APIs.) The data disk used by the container runtime and kubelet cannot be detached, or the node will become unavailable. For DeC nodes, the parameter details are the same as those for <b>rootVolume</b> .
storage	<a href="#">Storage</a> object	Disk initialization management parameter. This parameter is complex to configure. For details, see <a href="#">Attaching Disks to a Node</a> . <b>NOTE</b> <ul style="list-style-type: none"> <li>If a node has both local and EVS disks attached, make sure to not keep the default value for this parameter, or it may result in unexpected disk partitions.</li> </ul> <b>NOTE</b> <ul style="list-style-type: none"> <li>If you want to change the value range of a data disk to 20 to 32768, do not use the default value.</li> </ul> <b>NOTE</b> <ul style="list-style-type: none"> <li>If you want to use the shared disk space (with the runtime and Kubernetes partitions cancelled), do not retain the default value of this parameter. For details about the shared disk space, see <a href="#">Data Disk Space Allocation</a>.</li> </ul>
publicIP	<a href="#">NodePublicIP</a> object	EIP of a node. <b>NOTE</b> This parameter is not supported when you add a node to a node pool.
nodeNicSpec	<a href="#">NodeNicSpec</a> object	NIC of the node
count	Integer	Number of nodes to be created in a batch. The value must be a positive integer greater than or equal to 1 and less than or equal to the defined limit. This parameter can be left blank when it is used for a node pool.
billingMode	Integer	Node billing mode. <ul style="list-style-type: none"> <li><b>0</b>: pay-per-use</li> </ul>

Parameter	Type	Description
taints	Array of <b>Taint</b> objects	<p>You can add taints to created nodes to configure anti-affinity. A maximum of 20 taints can be added. Each taint contains the following parameters:</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain can be prefixed to a key.</li> <li>• <b>Value:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed.</li> <li>• <b>Effect:</b> Available options are <b>NoSchedule</b>, <b>PreferNoSchedule</b>, and <b>NoExecute</b>. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</li> </ul> <p>Example:</p> <pre data-bbox="815 1108 1428 1339">"taints": [{   "key": "status",   "value": "unavailable",   "effect": "NoSchedule" }, {   "key": "looks",   "value": "bad",   "effect": "NoSchedule" }]</pre>



Parameter	Type	Description
k8sTags	Map<String,String>	<p>Defined in key-value pairs. A maximum of 20 key-value pairs are allowed.</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain containing a maximum of 253 characters (such as <b>example.com/my-key</b>) can be prefixed to a key.</li> <li>• <b>Value:</b> The value can be left blank or contain 1 to 63 characters that start and end with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</li> </ul> <p>Example:</p> <pre>"k8sTags": {   "key": "value" }</pre>
ecsGroupId	String	<p>ECS group ID. If this parameter is configured, nodes will be created in the specified ECS group.</p> <p><b>NOTE</b> This configuration does not take effect when a node pool is created. To ensure that all nodes in a node pool are in the same ECS group, configure <b>nodeManagement</b> in the node pool.</p>
dedicatedHostId	String	<p>ID of the DeH to which the node is scheduled.</p> <p><b>NOTE</b> This parameter is not supported when you add a node during node pool creation.</p>

Parameter	Type	Description
userTags	Array of <a href="#">UserTag</a> objects	<p>Cloud server tag (resource tag). The key of a tag must be unique. The maximum number of custom tags supported by CCE varies depending on regions and cannot exceed 8. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</p> <p><b>NOTE</b> A tag key can have a maximum of 36 characters. It can only contain letters, digits, Unicode characters, and special characters (-_).</p>
runtime	<a href="#">Runtime</a> object	<p>Container runtime. Default value:</p> <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>

Parameter	Type	Description
initializedConditions	Array of strings	<p>Custom initialization flag, which is left blank by default.</p> <p>Before CCE nodes are initialized, <b>node.cloudprovider.kubernetes.io/uninitialized</b> is added to the nodes to prevent pods from being scheduled to these nodes. When creating a node, you can configure the <b>initializedConditions</b> parameter to control the taint removal time. By default, the timeout period is not configured.</p> <p>Example:</p> <ol style="list-style-type: none"> <li>1. Create a node and add parameter "initializedConditions": ["CCEInitial", "CustomedInitial"].</li> <li>2. After custom initialization is complete, CCE will call a Kubernetes API (for example, <b>PATCH /v1/nodes/{node_ip}/status</b>) to update the node <b>conditions</b> by adding two labels of types <b>CCEInitial</b> and <b>CustomedInitial</b> and setting the status to <b>True</b>, as shown in the following:</li> </ol> <pre data-bbox="815 1055 1430 1205"> status:   conditions:   - type: CCEInitial     status: 'True'   - type: CustomedInitial     status: 'True'                     </pre> <ol style="list-style-type: none"> <li>1. CCE polls <b>status.Conditions</b> of nodes to check whether there are conditions of types <b>CCEInitial</b> and <b>CustomedInitial</b>. If such conditions exist and the status is <b>True</b>, the node initialization is complete and the initialization taint is removed.</li> <li>2. <b>initializedConditions</b> allows you to configure a timeout period during node creation. For example, "initializedConditions": ["CCEInitial:15m", "CustomedInitial:15m"] indicates that the timeout period is 15 minutes. After the timeout period is reached, the initialization conditions will be automatically ignored and the initialization taint will be removed when CCE polls the node. <ul style="list-style-type: none"> <li>• The timeout interval in <b>initializedConditions</b> ranges from 1 to 99 seconds.</li> <li>• Use only letters and digits. Max. characters: 20.</li> <li>• Max. flags: 2.</li> </ul> </li> </ol>

Parameter	Type	Description
		<ul style="list-style-type: none"> <li>The unit of the timeout period is minute (m).</li> </ul>
extendParam	<a href="#">NodeExtendParam</a> object	Extended parameters for creating a node.
hostnameConfig	<a href="#">HostnameConfig</a> object	Kubernetes node name configuration parameter, which is supported by clusters of v1.23.6-r0 to v1.25 or clusters of v1.25.2-r0 or later versions.
serverEnterpriseProjectID	String	Enterprise project ID of a server. CCE does not support features provided by EPS. This field is used only for synchronizing enterprise project IDs of servers. To create a node or node pool, you can specify an existing enterprise project for it. If this field is left blank, it inherits the enterprise project attributes of the cluster that runs the node or node pool. When updating a node pool, any configuration modifications you made will apply only to new nodes. You need to modify the enterprise projects of existing nodes on the EPS console. If no value is specified during the update, this field will not be updated. If this parameter is left blank, the cluster enterprise project is returned.

**Table 4-493** Login

Parameter	Type	Description
sshKey	String	Name of the key pair used for login.
userPassword	<a href="#">UserPassword</a> object	Password used for node login.

**Table 4-494** UserPassword

Parameter	Type	Description
username	String	Login account. The default value is <b>root</b> .

Parameter	Type	Description
password	String	<p>Login password. If a username and password are used to create a node, this field is shielded in the response body. A password must meet the following complexity requirements:</p> <ul style="list-style-type: none"> <li>• Contains 8 to 26 characters.</li> <li>• Contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters !@\$%^-_=+[{ }];,./?</li> <li>• Cannot contain the username or the username spelled backwards. The <b>password</b> field must be salted during node creation. For details, see <a href="#">Adding a Salt in the password Field When Creating a Node</a>.</li> </ul>

Table 4-495 Volume

Parameter	Type	Description
size	Integer	<p>Disk size, in GiB.</p> <ul style="list-style-type: none"> <li>• System disk: 40 to 1024</li> <li>• Value range for data disks: 100 to 32768</li> </ul>
volumetype	String	<p>Disk type. For details about possible values, see the description of the <b>root_volume</b> parameter in the API used to create an ECS.</p> <ul style="list-style-type: none"> <li>• <b>SAS</b>: high I/O SAS disk</li> <li>• <b>SSD</b>: ultra-high I/O SSD disk</li> <li>• <b>SATA</b>: common I/O SATA disk SATA disks have been removed from EVS. You can find them attached only on existing nodes.</li> </ul>
extendParam	Map<String, Object>	<p>Extended disk parameters, defined in <b>extendparam</b> in the API used to create an ECS.</p>
cluster_id	String	<p>ID of the storage pool used by the ECS system disk. This field is used only for DeC clusters, which functions as <b>dssPoolID</b>, that is, the ID of the DSS storage pool.</p>
cluster_type	String	<p>Storage class of the cloud server system disk. The value is always <b>dss</b>. This field is used only for DeC clusters.</p>

Parameter	Type	Description
hw:passthrough	Boolean	<ul style="list-style-type: none"> <li>Pay attention to this field if your ECS is SDI-compliant. If the value of this field is <b>true</b>, the created disk is of SCSI type.</li> <li>If the node pool type is <b>ElasticBMS</b>, this field must be set to <b>true</b>.</li> <li>If a node specification involves both local and EVS disks, configure the disk initialization parameters. For details, see <a href="#">Attaching Disks to a Node</a>.</li> </ul>
metadata	<b>VolumeMetadata</b> object	EVS disk encryption information. This field is mandatory only when you need to encrypt the system disk or data disks of the node to be created.

**Table 4-496** VolumeMetadata

Parameter	Type	Description
__system__encrypted	String	Whether the EVS disk is encrypted. The value <b>0</b> indicates that the EVS disk is not encrypted, and the value <b>1</b> indicates that the EVS disk is encrypted. If this parameter is not specified, EVS disks will not be encrypted by default.
__system__cmkid	String	CMK ID, which indicates encryption in <b>metadata</b> . This field is used with <b>__system__encrypted</b> .

**Table 4-497** Storage

Parameter	Type	Description
storageSelectors	Array of <b>StorageSelectors</b> objects	Disk selection. Matched disks are managed according to <b>matchLabels</b> and <b>storageType</b> .
storageGroups	Array of <b>StorageGroups</b> objects	A storage group consists of multiple storage devices. It is used to divide storage space.

**Table 4-498** StorageSelectors

Parameter	Type	Description
name	String	Selector name, used as the index of <b>selectorNames</b> in <b>storageGroup</b> . Therefore, the name of each selector must be unique.
storageType	String	Storage type. Only <b>evs</b> (EVS disks) and <b>local</b> (local disks) are supported. The local storage does not support disk selection. All local disks will form a VG, so only one storageSelector of the local type is allowed.
matchLabels	<b>matchLabels</b> object	Matching field of an EVS disk. The value can be <b>size</b> , <b>volumeType</b> , <b>metadataEncrypted</b> , <b>metadataCmkid</b> , or <b>count</b> in <b>DataVolume</b> .

**Table 4-499** matchLabels

Parameter	Type	Description
size	String	Matched disk size. If this parameter is left unspecified, the disk size is not limited. Example: 100
volumeType	String	EVS disk type. <b>SSD</b> , <b>GPSSD</b> , <b>SAS</b> , <b>ESSD</b> , and <b>SATA</b> are supported. If this parameter is left blank, there is no restriction on the disk type.
metadataEncrypted	String	Disk encryption identifier. If this parameter is set to <b>0</b> , it specifies that the disk is not encrypted. If this parameter is set to <b>1</b> , it specifies that the disk is encrypted. If this parameter is left blank, there is no restriction on the disk encryption identifier.
metadataCmkid	String	CMK ID of an encrypted disk. The value is a string of 36 bytes. If this parameter is left blank, there is no restriction on the disk key ID.
count	String	Number of disks to be selected. If this parameter is left blank, all disks of this type are selected.

**Table 4-500** StorageGroups

Parameter	Type	Description
name	String	Name of a virtual storage group, which must be unique. <b>NOTE</b> <ul style="list-style-type: none"> <li>If <b>cceManaged</b> is set to <b>true</b>, the name must be set to <b>vgpaas</b>.</li> <li>If the data disk is used as a temporary storage volume, the name must be <b>vg-everest-localvolume-ephemeral</b>.</li> <li>If the data disk is used as a persistent storage volume, the name must be <b>vg-everest-localvolume-persistent</b>.</li> </ul>
cceManaged	Boolean	Storage space for Kubernetes and runtime components. Only one group can be set to <b>true</b> . If this parameter is left blank, the default value <b>false</b> is used.
selectorNames	Array of strings	Corresponds to <b>name</b> in <b>storageSelectors</b> . One group can select multiple selectors, but one selector can be selected by only one group.
virtualSpaces	Array of <a href="#">VirtualSpace</a> objects	Detailed management of space configuration in a group.

**Table 4-501** VirtualSpace

Parameter	Type	Description
name	String	Name of a virtualSpace. Options: <b>share</b> , <b>kubernetes</b> , <b>runtime</b> , and <b>user</b> <ul style="list-style-type: none"> <li><b>share</b>: shared space configuration (with the runtime and Kubernetes partitions cancelled). <b>lvmConfig</b> needs to be configured.</li> <li><b>Kubernetes</b>: Kubernetes space configuration. <b>lvmConfig</b> needs to be configured.</li> <li><b>runtime</b>: runtime space configuration. <b>runtimeConfig</b> needs to be configured.</li> <li><b>user</b>: user space configuration. <b>lvmConfig</b> needs to be configured.</li> </ul>



Parameter	Type	Description
size	String	Size of a virtualSpace. The value must be an integer in percentage. Example: 90%. <b>NOTE</b> The value ranges from 10 to 90. The sum of percentages of all virtualSpaces in a group cannot exceed 100%.
lvmConfig	<b>LVMConfig</b> object	LVM configuration managements, applicable to the <i>share*</i> , <b>kubernetes</b> , and <b>user</b> spaces. Note that one virtual space supports only one config.
runtimeConfig	<b>RuntimeConfig</b> object	runtime configurations, applicable to the <b>runtime</b> space. Note that one virtual space supports only one config.

**Table 4-502** LVMConfig

Parameter	Type	Description
lvType	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.
path	String	Disk mount path. This parameter takes effect only in user configuration. The value is an absolute path. Digits, letters, periods (.), hyphens (-), and underscores (_) are allowed. <b>NOTE</b> Do not leave the path empty or set it to a key OS path such as the root directory. Key OS paths that cannot be specified: /, /home, /home/, /bin, /bin/, /lib, /lib/, /root, /root/, /boot, /boot/, /dev, /dev/, /etc, /etc/, /lost+found, /lost+found/, /mnt, /mnt/, /proc, /proc/, /sbin, /sbin/, /srv, /srv/, /tmp, /tmp/, /var, /var/, /media, /media/, /opt, /opt/, /selinux, /selinux/, /sys, /sys/, /usr, /usr/, /opt/cloud/, /mnt/paas/, /home/paas/, /var/paas/, /var/lib/, /var/script/

**Table 4-503** RuntimeConfig

Parameter	Type	Description
lvType	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.

**Table 4-504** NodePublicIP

Parameter	Type	Description
ids	Array of strings	IDs of existing EIPs. The quantity cannot be greater than the number of nodes to be created. <b>NOTE</b> If <b>ids</b> has been set, you do not need to set <b>count</b> and <b>eip</b> .
count	Integer	Number of EIPs to be dynamically created. <b>NOTE</b> <b>count</b> and <b>eip</b> must be set at the same time.
eip	<b>NodeEIPSpec</b> object	EIP configuration. <b>NOTE</b> This parameter is not supported when creating a node pool.

**Table 4-505** NodeEIPSpec

Parameter	Type	Description
iptype	String	EIP type, specified in <b>publicip.type</b> in the API for assigning an EIP.
bandwidth	<b>NodeBandwidth</b> object	Bandwidth parameters of the EIP

**Table 4-506** NodeBandwidth

Parameter	Type	Description
chargemode	String	<p>Bandwidth billing mode.</p> <ul style="list-style-type: none"> <li>• If this field is not specified, the billing is based on bandwidth.</li> <li>• If the field is null, the billing is based on bandwidth.</li> <li>• If the field value is <b>traffic</b>, the billing is based on traffic.</li> <li>• If the value is out of the preceding options, the cloud server will fail to be created.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• Billed by bandwidth: The billing will be based on the data transmission rate (in Mbps) of public networks. This billing mode is recommended if your bandwidth usage is higher than 10%.</li> <li>• Billed by traffic: The billing is based on the total amount of data (in GB) transmitted over the public network. This mode is available only when you are creating a pay-per-use node. This billing mode is recommended if your bandwidth usage is lower than 10%.</li> </ul>
size	Integer	Bandwidth size, specified in <b>bandwidth.size</b> in the API for assigning an EIP.
sharetype	String	Bandwidth sharing type. Value options: <b>PER</b> (exclusive bandwidth)

**Table 4-507** NodeNicSpec

Parameter	Type	Description
primaryNic	<b>NicSpec</b> object	Description of the primary NIC.
extNics	Array of <b>NicSpec</b> objects	<p>Extension NIC</p> <p><b>NOTE</b> This parameter is not supported when you add a node to a node pool.</p>

**Table 4-508 NicSpec**

Parameter	Type	Description
subnetId	String	Network ID of the subnet to which a NIC belongs. If <b>subnetId</b> is not specified when a primary NIC is creating, the cluster subnet will be used. If <b>subnetList</b> is also configured for a node pool, the <b>subnetList</b> field is used for adding subnets to the node pool. When creating an extension NIC, you must specify <b>subnetId</b> .
fixedIps	Array of strings	The IP address of the primary ENI is specified using <b>fixedIps</b> . The number of IP addresses cannot be greater than the number of created nodes. Either <b>fixedIps</b> or <b>ipBlock</b> can be specified. <b>fixedIps</b> cannot be specified for extension ENIs.
ipBlock	String	CIDR format of the primary NIC IP range. The IP address of the created node falls in this range. <b>fixedIps</b> and <b>ipBlock</b> cannot be specified at the same time.
subnetList	Array of strings	Network ID list of the subnet where the NIC resides. A maximum of 20 subnets can be configured for a node pool.

**Table 4-509 Taint**

Parameter	Type	Description
key	String	Key.
value	String	Value.
effect	String	Effect.

**Table 4-510 UserTag**

Parameter	Type	Description
key	String	Key of the cloud server label. The value cannot start with <b>CCE-</b> or <b>__type_baremetal</b> .
value	String	Value of the cloud server label.

**Table 4-511** Runtime

Parameter	Type	Description
name	String	<p>Container runtime. Default value:</p> <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>

**Table 4-512** NodeExtendParam

Parameter	Type	Description
ecs:performancetype	String	ECS flavor types. This field is returned in the response.
maxPods	Integer	<p>Maximum number of pods that can be created on a node, including the default system pods. Value range: 16 to 256.</p> <p>This limit prevents the node from being overloaded of pods.</p> <p>The number of pods that can be created on a node is determined by multiple parameters. For details, see <a href="#">Maximum Number of Pods That Can Be Created on a Node</a>.</p>

Parameter	Type	Description
DockerLVMConfigOverride	String	<p>Docker data disk configuration item. (This parameter has been discarded. Use the <b>storage</b> field instead.) Example default configuration: "DockerLVMConfigOverride":"dockerThinpool=vgpaas/90%VG;kubernetesLV=vgpaas/10%VG;diskType=evs;lvType=linear"</p> <p>By default, if no VD disk is available, an error occurs because the data disk fails to be found. Set <b>diskType</b> based on the actual drive letter type. The following fields are included:</p> <ul style="list-style-type: none"> <li>• <b>userLV</b> (optional): size of the user space, for example, <b>vgpaas/20%VG</b>.</li> <li>• <b>userPath</b> (optional): mount path of the user space, for example, <b>/home/wqt-test</b>.</li> <li>• <b>diskType</b>: disk type. Currently, only <b>evs</b>, <b>hdd</b>, and <b>ssd</b> are supported.</li> <li>• <b>lvType</b>: type of a logic volume. The value can be <b>linear</b> or <b>striped</b>.</li> <li>• <b>dockerThinpool</b>: Docker space size, for example, <b>vgpaas/60%VG</b>.</li> <li>• <b>kubernetesLV</b>: kubelet space size, for example, <b>vgpaas/20%VG</b>.</li> </ul>
dockerBaseSize	Integer	<p>Available disk space of a single container on a node, in GB. This parameter has been discarded. Use <b>containerBaseSize</b>.</p> <p>If this parameter is left blank or is set to <b>0</b>, the default value is used. In Device Mapper mode, the default value is <b>10</b>. In OverlayFS mode, the available space of a single container is not limited by default, and the <b>dockerBaseSize</b> setting takes effect only on nodes running EulerOS in clusters of the new versions.</p> <p>For details about how to allocate the space for the container runtime, see <a href="#">Data Disk Space Allocation</a>.</p> <p>When Device Mapper is used, it is recommended that <b>dockerBaseSize</b> be set to a value less than or equal to 80 GB. If the value is too large, the container runtime may fail to be started due to long initialization. If there are special requirements for the container disk space, you can mount an external or local storage device.</p>

Parameter	Type	Description
containerBaseSize	Integer	<p>Available disk space of a single container on a node, in GiB.</p> <p>If this parameter is left blank or is set to <b>0</b>, the default value is used. In OverlayFS mode, the available space of a single container is not limited by default. In Device Mapper mode, the default value is <b>10</b>, and the <b>containerBaseSize</b> setting takes effect only on nodes running EulerOS in clusters v1.23.14-r0, v1.25.9-r0, v1.27.6-r0, v1.28.4-r0 or later.</p> <p>For details about how to allocate the space for the container runtime, see <a href="#">Data Disk Space Allocation</a>.</p> <p>When Device Mapper is used, it is recommended that <b>containerBaseSize</b> be set to a value less than or equal to 80 GB. If the value is too large, the container runtime may fail to be started due to long initialization. If there are special requirements for the container disk space, you can mount an external or local storage device. In new versions, the Device Mapper mode is used only by BMSs in the same resource pool and is being discarded.</p> <p><b>NOTE</b> This parameter cannot be updated when a node pool is being updated.</p>
publicKey	String	Public key of a node.
alpha.cce/preInstall	String	<p>Pre-installation scripts. The total number of characters after transcoding cannot exceed 10240.</p> <p><b>NOTE</b> The input value must be encoded using Base64. The method is as follows:</p> <pre>echo -n "Content to be encoded"   base64</pre>
alpha.cce/postInstall	String	<p>Post-installation scripts. The total number of characters after transcoding cannot exceed 10240.</p> <p><b>NOTE</b> The input value must be encoded using Base64. The method is as follows:</p> <pre>echo -n "Content to be encoded"   base64</pre>
alpha.cce/NodeImageID	String	This parameter is required when a custom image is used to create a BMS node.

Parameter	Type	Description
chargingMode	Integer	Billing mode of a node. This parameter has been deprecated. Use the <b>billingMode</b> parameter in <b>NodeSpec</b> .
agency_name	String	Name of an agency An agency is created by a tenant administrator on Identity and Access Management (IAM) to provide temporary credentials for CCE nodes to access cloud servers. This parameter is returned only when it is transferred during node creation.
kubeReservedMem	Integer	Reserved node memory, which is reserved for Kubernetes components.
systemReservedMem	Integer	Reserved node memory, which is reserved for system components.
init-node-password	String	Node password. If this parameter is used as a response parameter, asterisks (*) will be displayed.

**Table 4-513** HostnameConfig

Parameter	Type	Description
type	String	Configuration type of the Kubernetes node name. The default value is <b>privatelp</b> . <ul style="list-style-type: none"> <li>• <b>privatelp</b>: The Kubernetes node is named after its IP address.</li> <li>• <b>cceNodeName</b>: The Kubernetes node is named after the CCE node.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• For a node which is configured using <b>cceNodeName</b>, the name is the same as the Kubernetes node name and the ECS name. The node name cannot be changed. If the ECS name is changed on the ECS console, the node name will retain unchanged after ECS synchronization.</li> <li>• For a node which is configured using <b>cceNodeName</b>, to avoid a conflict between Kubernetes nodes, the system automatically adds a suffix to each node name. The suffix is in the format of <b>A hyphen (-) Five random characters</b>. The value of the random characters is a lowercase letter or a digit ranging from 0 to 9.</li> </ul>



**Table 4-514** NodePoolNodeAutoscaling

Parameter	Type	Description
enable	Boolean	Whether to enable auto scaling.
minNodeCount	Integer	Minimum number of nodes allowed if auto scaling is enabled. The value cannot be greater than the maximum number of nodes allowed by the cluster specifications.
maxNodeCount	Integer	Maximum number of nodes allowed if auto scaling is enabled. This value must be greater than or equal to the value of <b>minNodeCount</b> and cannot exceed the maximum number of nodes in the cluster specifications.
scaleDownCooldownTime	Integer	Interval between two scaling operations, in minutes. During this period, nodes added after a scale-up will not be deleted.
priority	Integer	Weight of a node pool. A node pool with a higher weight has a higher priority during scaling.

**Table 4-515** NodeManagement

Parameter	Type	Description
serverGroupReference	String	Cloud server group ID. If this field is specified, all nodes in the node pool will be created in this group. The group ID can be specified only when you create the node pool and cannot be modified. When you specify a cloud server group, the number of nodes in the node pool cannot exceed the group quota.

**Table 4-516** SecurityID

Parameter	Type	Description
id	String	Security group ID

**Table 4-517** ExtensionScaleGroup

Parameter	Type	Description
metadata	<a href="#">ExtensionScaleGroupMetadata</a> object	Basic information about an extended scaling group

Parameter	Type	Description
spec	<a href="#">ExtensionScaleGroupSpec</a> object	Configurations of an extended scaling group, which carry different configurations from those of the default scaling group

**Table 4-518** ExtensionScaleGroupMetadata

Parameter	Type	Description
uid	String	UUID of an extended scaling group, which is automatically generated by the system.
name	String	Name of an extended scaling group. The value cannot be <b>default</b> and can contain a maximum of 55 characters. Only digits, lowercase letters, and hyphens (-) are allowed.

**Table 4-519** ExtensionScaleGroupSpec

Parameter	Type	Description
flavor	String	Node flavor
az	String	AZ of a node. If this parameter is not specified or left blank, the default scaling group configurations take effect.
capacityReservationSpecification	<a href="#">CapacityReservationSpecification</a> object	Capacity reservation configurations of an extended scaling group
autoscaling	<a href="#">ScaleGroupAutoscaling</a> object	Auto scaling configurations of an extended scaling group

**Table 4-520** CapacityReservationSpecification

Parameter	Type	Description
id	String	Private pool ID. The parameter value can be ignored when <b>preference</b> is set to <b>none</b> .
preference	String	Capacity of a private storage pool. If the value is <b>none</b> , the capacity reservation is not specified. If the value is <b>targeted</b> , the capacity reservation is specified. In this case, the ID cannot be left blank.

**Table 4-521** ScaleGroupAutoscaling

Parameter	Type	Description
enable	Boolean	Whether to enable auto scaling for a scaling group. This function is disabled by default.
extensionPriority	Integer	Priority of a scaling group. If no value is specified for this parameter, it defaults to <b>0</b> . A higher value indicates a greater priority.
minNodeCount	Integer	Minimum number of nodes in a scaling group during auto scaling. The value must be greater than <b>0</b> .
maxNodeCount	Integer	Maximum number of nodes that can be retained in a scaling group during auto scaling. The value must be greater than or equal to that of <b>minNodeCount</b> , and can neither be greater than the maximum number of nodes allowed by the cluster nor the maximum number of nodes in the node pool.

**Table 4-522** CreateNodePoolStatus

Parameter	Type	Description
currentNode	Integer	Total number of nodes in the current node pool (excluding the nodes that are being deleted).
creatingNode	Integer	Number of nodes that are being created in the node pool.
deletingNode	Integer	Number of nodes that are being deleted in the current node pool.

Parameter	Type	Description
phase	String	<p>Node pool status. Options:</p> <ul style="list-style-type: none"> <li>• <b>Null:</b> The node pool is available. (The number of nodes in the current node pool has reached the expected value, and no node scaling is being performed.)</li> <li>• <b>Synchronizing:</b> The node pool is being scaled. (The number of nodes in the current node pool has not reached the expected value, and no node scaling is being performed.)</li> <li>• <b>Synchronized:</b> The node pool scaling is pending. (The number of nodes in the current node pool has not reached the expected value, or node scaling is being performed.)</li> <li>• <b>SoldOut:</b> The node pool cannot be scaled out. (This field is used in multiple scenarios, for example, resources in the node pool have been sold out or the resource quota is insufficient.)</li> </ul> <p><b>NOTE</b> This node pool status has been deprecated. It is reserved only for compatibility. Do not use it anymore. Use the following instead:</p> <ul style="list-style-type: none"> <li>• Node pool scaling status: You can obtain the status of the current node pool using parameters such as <b>currentNode</b>, <b>creatingNode</b>, or <b>deletingNode</b>.</li> <li>• Node pool scale-out: You can use <b>conditions</b> to obtain the detailed status of a node pool. <b>Scalable</b> can replace <b>SoldOut</b>.</li> <li>• <b>Deleting:</b> The node pool is being deleted.</li> <li>• <b>Error:</b> An error occurred in the node pool.</li> </ul>
conditions	Array of <b>NodePoolCondition</b> objects	Node pool status details. For details, see the definition of <b>Condition</b> .
scaleGroupStatuses	Array of <b>ScaleGroupStatus</b> objects	Detailed status of a scaling group. For details, see the definition of <b>ScaleGroupStatus</b> .

**Table 4-523** ScaleGroupStatus

Parameter	Type	Description
name	String	Scaling group name

Parameter	Type	Description
uid	String	UUID of a scaling group
creationTimes tamp	String	The time when a scaling group was created
updateTimest amp	String	The time when a scaling group was updated
phase	String	<p>Status of a scaling group.</p> <ul style="list-style-type: none"> <li>• <b>Null:</b> The scaling group is available. (The number of nodes in the current scaling group has reached the expected value, and no node scaling is being performed.)</li> <li>• <b>Synchronizing:</b> The scaling group is being scaled. (The number of nodes in the current scaling group has not reached the expected value, and no node scaling is being performed.)</li> <li>• <b>Synchronized:</b> The scaling is pending. (The number of nodes in the current scaling group has not reached the expected value, or node scaling is being performed.)</li> <li>• <b>SoldOut:</b> The scaling group cannot be scaled out. (This field is used in multiple scenarios, for example, resources in the node pool have been sold out or the resource quota is insufficient.)</li> </ul> <p><b>NOTE</b> This scaling group status has been deprecated. It is reserved only for compatibility. Do not use it anymore. Use the following instead:</p> <ul style="list-style-type: none"> <li>• Scaling status of a scaling group: You can obtain the status of the current scaling group using parameters such as <b>desiredNodeCount</b>, <b>existingNodeCount</b>, or <b>upcomingNodeCount</b>.</li> <li>• Scaling group scale-out: You can use <b>conditions</b> to obtain the detailed status of a scaling group. <b>Scalable</b> can replace <b>SoldOut</b>.</li> <li>• <b>Deleting:</b> The scaling group being deleted.</li> <li>• <b>Error:</b> An error occurred in the scaling group.</li> </ul>
desiredNodeC ount	Integer	Expected number of nodes in a scaling group
unpaidScaleN odeCount	Integer	Number of unpaid nodes in an order
existingNodeC ount	<b>existingNode Count</b> object	Statistics on existing nodes in a scaling group

Parameter	Type	Description
upcomingNodeCount	<a href="#">upcomingNodeCount</a> object	Statistics on the nodes to be created in a scaling group
scaleDownDisabledNodeCount	Integer	Number of nodes that cannot be deleted in a scaling group
conditions	Array of <a href="#">NodePoolCondition</a> objects	Scaling group status details. For details, see the definition of <b>Condition</b> .

**Table 4-524** existingNodeCount

Parameter	Type	Description
postPaid	Integer	Number of pay-per-use nodes
total	Integer	Total number of pay-per-use nodes

**Table 4-525** upcomingNodeCount

Parameter	Type	Description
postPaid	Integer	Number of pay-per-use nodes
total	Integer	Total number of pay-per-use nodes

**Table 4-526** NodePoolCondition

Parameter	Type	Description
type	String	Condition type. The options are as follows: <ul style="list-style-type: none"> <li>● <b>Scalable</b>: whether a node pool can be scaled. If the status is <b>False</b>, node pool scaling will not be triggered again.</li> <li>● <b>QuotaInsufficient</b>: Quotas on which node pool scaling depends are insufficient, affecting the node pool scaling status.</li> <li>● <b>ResourceInsufficient</b>: Resources on which node pool scaling depends are insufficient, affecting the node pool scaling status.</li> <li>● <b>UnexpectedError</b>: The node pool fails to be scaled out due to unexpected reasons, affecting the node pool scaling status.</li> <li>● <b>Error</b>: A node pool error occurs. A common trigger is deletion failure.</li> </ul>
status	String	Current status of <b>Condition</b> . The options are as follows: <ul style="list-style-type: none"> <li>● "True"</li> <li>● "False"</li> </ul>
lastProbeTime	String	Time when the status was last checked
lastTransitTime	String	Time when the status was last changed
reason	String	Reason why the status was last changed
message	String	Detailed condition description

## Example Requests

Create a node pool with the following configurations: The number of nodes is 0. The node pool flavor is 2 vCPUs and 4 GiB of memory. The node OS is EulerOS 2.5. Docker containers are deployed. The sizes of the system disk and data disk of the node are 40 GB and 100 GB, respectively. Both the disks are of the high I/O type.

```
{
  "kind": "NodePool",
  "apiVersion": "v3",
  "metadata": {
    "name": "lc-it-nodepool-79796"
  },
  "spec": {
    "initialNodeCount": 0,
    "type": "vm",
    "autoscaling": {
      "enable": false,
      "minNodeCount": 0,
      "maxNodeCount": 1,
      "scaleDownCooldownTime": 0,

```

```

    "priority" : 0
  },
  "nodeManagement" : {
    "serverGroupReference" : ""
  },
  "nodeTemplate" : {
    "flavor" : "s6.large.2",
    "az" : "*****",
    "os" : "EulerOS 2.5",
    "login" : {
      "sshKey" : "KeyPair-001"
    },
    "rootVolume" : {
      "volumetype" : "SAS",
      "size" : 40
    },
    "dataVolumes" : [ {
      "volumetype" : "SAS",
      "size" : 100,
      "extendParam" : {
        "useType" : "docker"
      }
    } ],
    "billingMode" : 0,
    "extendParam" : {
      "alpha.cce/preInstall" : "",
      "alpha.cce/postInstall" : "",
      "alpha.cce/NodeImageID" : "",
      "maxPods" : 110
    },
    "nodeNicSpec" : {
      "primaryNic" : {
        "subnetId" : "7e767d10-7548-4df5-ad72-aeac1d08bd8a"
      }
    }
  },
  "podSecurityGroups" : [ {
    "id" : ""
  } ]
}

```

## Example Responses

### Status code: 201

The job for creating a node pool in a specified cluster is successfully delivered.

```

{
  "kind" : "NodePool",
  "apiVersion" : "v3",
  "metadata" : {
    "name" : "lc-it-nodepool-79796",
    "uid" : "99addaa2-69eb-11ea-a592-0255ac1001bb"
  },
  "spec" : {
    "type" : "vm",
    "nodeTemplate" : {
      "flavor" : "s6.large.2",
      "az" : "*****",
      "os" : "EulerOS 2.5",
      "login" : {
        "sshKey" : "KeyPair-001"
      },
      "rootVolume" : {
        "volumetype" : "SAS",
        "size" : 40
      },
      "dataVolumes" : [ {

```



```

"volumetype" : "SAS",
"size" : 100,
"extendParam" : {
  "useType" : "docker"
}
}],
"publicIP" : {
  "eip" : {
    "bandwidth" : { }
  }
},
"nodeNicSpec" : {
  "primaryNic" : {
    "subnetId" : "7e767d10-7548-4df5-ad72-aeac1d08bd8a"
  }
},
"billingMode" : 0,
"extendParam" : {
  "alpha.cce/NodeImageID" : "",
  "alpha.cce/postInstall" : "",
  "alpha.cce/preInstall" : "",
  "maxPods" : 110
},
"k8sTags" : {
  "cce.cloud.com/cce-nodepool" : "lc-it-nodepool-79796"
}
},
"autoscaling" : {
  "maxNodeCount" : 1
},
"nodeManagement" : { }
},
"status" : {
  "phase" : ""
}
}

```

## Status Codes

Status Code	Description
201	The job for creating a node pool in a specified cluster is successfully delivered.

## Error Codes

See [Error Codes](#).

### 4.4.2 Reading a Specified Node Pool

#### Function

This API is used to obtain details about a specified node pool.

#### NOTE

The URL for cluster management is in the format of `https://Endpoint/uri`. In the URL, **uri** indicates the resource path, that is, the path for API access.

## URI

GET /api/v3/projects/{project\_id}/clusters/{cluster\_id}/nodepools/{nodepool\_id}

**Table 4-527** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
nodepool_id	Yes	String	Node pool ID.

## Request Parameters

**Table 4-528** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

## Response Parameters

**Status code: 200**

**Table 4-529** Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at <b>NodePool</b> .
apiVersion	String	API version. The value is fixed at <b>v3</b> .
metadata	<a href="#">NodePoolMetadata</a> object	Metadata of the node pool.

Parameter	Type	Description
spec	<a href="#">NodePoolSpec</a> object	Node pool specifications. <b>NOTE</b> <ul style="list-style-type: none"> <li>This field is optional for updating a node pool.</li> </ul>
status	<a href="#">NodePoolStatus</a> object	Node pool status.

**Table 4-530** NodePoolMetadata

Parameter	Type	Description
name	String	Node pool name. <b>NOTE</b> Naming rules: Enter 1 to 50 characters, starting with a lowercase letter and not ending with a hyphen (-). Only lowercase letters, digits, and hyphens (-) are allowed. <ul style="list-style-type: none"> <li>You cannot create node pools named DefaultPool.</li> </ul>
uid	String	UID of the node pool. The value is automatically generated after the object is updated. A user-defined value will not take effect.
annotations	Map<String,String>	Annotations of a node pool in key-value pairs. This parameter is used only for query and cannot be input through a request. Entered data of this parameter is invalid.
updateTimestamp	String	Update time.
creationTimestamp	String	Creation time.

**Table 4-531** NodePoolSpec

Parameter	Type	Description
type	String	Node pool type. If this parameter is left blank, the value <b>vm</b> is used by default. <ul style="list-style-type: none"> <li><b>vm</b>: ECS</li> <li><b>ElasticBMS</b>: C6 general computing-plus BMS. An example flavor is <b>c6.22xlarge.2.physical</b>.</li> <li><b>bms</b>: BMS</li> </ul>

Parameter	Type	Description
nodeTemplate	<a href="#">NodeSpec</a> object	Detailed parameters of the node pool template.
initialNodeCount	Integer	Initial number of nodes for the node pool. When queried, the value is the number of target nodes in the node pool.
autoscaling	<a href="#">NodePoolNodeAutoscaling</a> object	Auto scaling parameters
nodeManagement	<a href="#">NodeManagement</a> object	Node management configuration
podSecurityGroups	Array of <a href="#">SecurityID</a> objects	Security group configuration. This parameter is available only for CCE Turbo clusters.
extensionScaleGroups	Array of <a href="#">ExtensionScaleGroup</a> objects	Configuration list of the extended scaling groups in a node pool. For details, see the definition of <b>ExtensionScaleGroup</b> .
customSecurityGroups	Array of strings	Custom security group settings for a node pool. New nodes scaled out in a node pool can be bound to a specified security group. <ul style="list-style-type: none"> <li>• Specifying no security group ID will add the new nodes to the default security group of the worker nodes.</li> <li>• Specifying a valid security group ID will put new nodes in that security group.</li> <li>• When specifying a security group, do not modify the rules of the port on which CCE running depends.</li> </ul>

**Table 4-532** NodeSpec

Parameter	Type	Description
flavor	String	Node specifications. For details about the node specifications supported by CCE, see <a href="#">Node Flavor Description</a> .

Parameter	Type	Description
az	String	AZ where the node to be created is located. You need to specify the AZ name. A node created using APIs does not support random AZs. <b>NOTE</b> <ul style="list-style-type: none"> <li>When you are creating a node pool and configuring a scaling group, this parameter cannot be set to <b>random</b>.</li> </ul>
os	String	Node OS. For details about the supported OSs, see <a href="#">Node OS</a> . <b>NOTE</b> <ul style="list-style-type: none"> <li>The system automatically selects the supported OS based on the cluster version. If the current cluster version does not support the OS, an error will be reported.</li> <li>If <b>alpha.cce/NodeImageID</b> in <b>extendParam</b> is specified during node creation, you do not need to configure this parameter.</li> <li>This parameter is mandatory when creating a node pool.</li> <li>If you are using shared disk space when creating a node, this parameter is mandatory. Specifically, the disk initialization configuration management parameter <b>storage</b> is used and the <b>name</b> field of <b>virtualSpaces</b> in <b>StorageGroups</b> is set to <b>share</b>.</li> </ul>
login	<a href="#">Login</a> object	Node login mode. Either the key pair or password must be used for login. <b>NOTE</b> This parameter cannot be modified when a node pool is being updated.
rootVolume	<a href="#">Volume</a> object	Information about disks on the node
dataVolumes	Array of <a href="#">Volume</a> objects	Data disk parameters of a node. A VM can have up to 16 disks attached, while a BMS can only have up to 10 disks attached. The maximum number of disks that can be attached to either a VM or BMS is limited by their respective maximum disk limits. (You can add more data disks to each of your nodes through the CCE console and APIs.) The data disk used by the container runtime and kubelet cannot be detached, or the node will become unavailable. For DeC nodes, the parameter details are the same as those for <b>rootVolume</b> .

Parameter	Type	Description
storage	<a href="#">Storage</a> object	<p>Disk initialization management parameter. This parameter is complex to configure. For details, see <a href="#">Attaching Disks to a Node</a>.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If a node has both local and EVS disks attached, make sure to not keep the default value for this parameter, or it may result in unexpected disk partitions.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If you want to change the value range of a data disk to 20 to 32768, do not use the default value.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If you want to use the shared disk space (with the runtime and Kubernetes partitions cancelled), do not retain the default value of this parameter. For details about the shared disk space, see <a href="#">Data Disk Space Allocation</a>.</li> </ul>
publicIP	<a href="#">NodePublicIP</a> object	<p>EIP of a node.</p> <p><b>NOTE</b></p> <p>This parameter is not supported when you add a node to a node pool.</p>
nodeNicSpec	<a href="#">NodeNicSpec</a> object	NIC of the node
count	Integer	Number of nodes to be created in a batch. The value must be a positive integer greater than or equal to 1 and less than or equal to the defined limit. This parameter can be left blank when it is used for a node pool.
billingMode	Integer	<p>Node billing mode.</p> <ul style="list-style-type: none"> <li><b>0</b>: pay-per-use</li> </ul>

Parameter	Type	Description
taints	Array of <b>Taint</b> objects	<p>You can add taints to created nodes to configure anti-affinity. A maximum of 20 taints can be added. Each taint contains the following parameters:</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain can be prefixed to a key.</li> <li>• <b>Value:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed.</li> <li>• <b>Effect:</b> Available options are <b>NoSchedule</b>, <b>PreferNoSchedule</b>, and <b>NoExecute</b>. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</li> </ul> <p>Example:</p> <pre data-bbox="815 1108 1428 1339">"taints": [{   "key": "status",   "value": "unavailable",   "effect": "NoSchedule" }, {   "key": "looks",   "value": "bad",   "effect": "NoSchedule" }]</pre>

Parameter	Type	Description
k8sTags	Map<String,String>	<p>Defined in key-value pairs. A maximum of 20 key-value pairs are allowed.</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain containing a maximum of 253 characters (such as <b>example.com/my-key</b>) can be prefixed to a key.</li> <li>• <b>Value:</b> The value can be left blank or contain 1 to 63 characters that start and end with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</li> </ul> <p>Example:</p> <pre>"k8sTags": {   "key": "value" }</pre>
ecsGroupId	String	<p>ECS group ID. If this parameter is configured, nodes will be created in the specified ECS group.</p> <p><b>NOTE</b> This configuration does not take effect when a node pool is created. To ensure that all nodes in a node pool are in the same ECS group, configure <b>nodeManagement</b> in the node pool.</p>
dedicatedHostId	String	<p>ID of the DeH to which the node is scheduled.</p> <p><b>NOTE</b> This parameter is not supported when you add a node during node pool creation.</p>



Parameter	Type	Description
userTags	Array of <a href="#">UserTag</a> objects	<p>Cloud server tag (resource tag). The key of a tag must be unique. The maximum number of custom tags supported by CCE varies depending on regions and cannot exceed 8. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</p> <p><b>NOTE</b> A tag key can have a maximum of 36 characters. It can only contain letters, digits, Unicode characters, and special characters (-_).</p>
runtime	<a href="#">Runtime</a> object	<p>Container runtime. Default value:</p> <ul style="list-style-type: none"><li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li><li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li></ul>

Parameter	Type	Description
initializedConditions	Array of strings	<p>Custom initialization flag, which is left blank by default.</p> <p>Before CCE nodes are initialized, <b>node.cloudprovider.kubernetes.io/uninitialized</b> is added to the nodes to prevent pods from being scheduled to these nodes. When creating a node, you can configure the <b>initializedConditions</b> parameter to control the taint removal time. By default, the timeout period is not configured.</p> <p>Example:</p> <ol style="list-style-type: none"> <li>1. Create a node and add parameter "initializedConditions": ["CCEInitial", "CustomedInitial"].</li> <li>2. After custom initialization is complete, CCE will call a Kubernetes API (for example, <b>PATCH /v1/nodes/{node_ip}/status</b>) to update the node <b>conditions</b> by adding two labels of types <b>CCEInitial</b> and <b>CustomedInitial</b> and setting the status to <b>True</b>, as shown in the following:</li> </ol> <pre>status:   conditions:   - type: CCEInitial     status: 'True'   - type: CustomedInitial     status: 'True'</pre> <ol style="list-style-type: none"> <li>1. CCE polls <b>status.Conditions</b> of nodes to check whether there are conditions of types <b>CCEInitial</b> and <b>CustomedInitial</b>. If such conditions exist and the status is <b>True</b>, the node initialization is complete and the initialization taint is removed.</li> <li>2. <b>initializedConditions</b> allows you to configure a timeout period during node creation. For example, "initializedConditions": ["CCEInitial:15m", "CustomedInitial:15m"] indicates that the timeout period is 15 minutes. After the timeout period is reached, the initialization conditions will be automatically ignored and the initialization taint will be removed when CCE polls the node. <ul style="list-style-type: none"> <li>• The timeout interval in <b>initializedConditions</b> ranges from 1 to 99 seconds.</li> <li>• Use only letters and digits. Max. characters: 20.</li> <li>• Max. flags: 2.</li> </ul> </li> </ol>

Parameter	Type	Description
		<ul style="list-style-type: none"> <li>The unit of the timeout period is minute (m).</li> </ul>
extendParam	<b>NodeExtendParam</b> object	Extended parameters for creating a node.
hostnameConfig	<b>HostnameConfig</b> object	Kubernetes node name configuration parameter, which is supported by clusters of v1.23.6-r0 to v1.25 or clusters of v1.25.2-r0 or later versions.
serverEnterpriseProjectID	String	Enterprise project ID of a server. CCE does not support features provided by EPS. This field is used only for synchronizing enterprise project IDs of servers. To create a node or node pool, you can specify an existing enterprise project for it. If this field is left blank, it inherits the enterprise project attributes of the cluster that runs the node or node pool. When updating a node pool, any configuration modifications you made will apply only to new nodes. You need to modify the enterprise projects of existing nodes on the EPS console. If no value is specified during the update, this field will not be updated. If this parameter is left blank, the cluster enterprise project is returned.

**Table 4-533** Login

Parameter	Type	Description
sshKey	String	Name of the key pair used for login.
userPassword	<b>UserPassword</b> object	Password used for node login.

**Table 4-534** UserPassword

Parameter	Type	Description
username	String	Login account. The default value is <b>root</b> .

Parameter	Type	Description
password	String	<p>Login password. If a username and password are used to create a node, this field is shielded in the response body. A password must meet the following complexity requirements:</p> <ul style="list-style-type: none"> <li>• Contains 8 to 26 characters.</li> <li>• Contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters !@\$%^-_=+[{ }];,./?</li> <li>• Cannot contain the username or the username spelled backwards. The <b>password</b> field must be salted during node creation. For details, see <a href="#">Adding a Salt in the password Field When Creating a Node</a>.</li> </ul>

Table 4-535 Volume

Parameter	Type	Description
size	Integer	<p>Disk size, in GiB.</p> <ul style="list-style-type: none"> <li>• System disk: 40 to 1024</li> <li>• Value range for data disks: 100 to 32768</li> </ul>
volumetype	String	<p>Disk type. For details about possible values, see the description of the <b>root_volume</b> parameter in the API used to create an ECS.</p> <ul style="list-style-type: none"> <li>• <b>SAS</b>: high I/O SAS disk</li> <li>• <b>SSD</b>: ultra-high I/O SSD disk</li> <li>• <b>SATA</b>: common I/O SATA disk SATA disks have been removed from EVS. You can find them attached only on existing nodes.</li> </ul>
extendParam	Map<String, Object>	<p>Extended disk parameters, defined in <b>extendparam</b> in the API used to create an ECS.</p>
cluster_id	String	<p>ID of the storage pool used by the ECS system disk. This field is used only for DeC clusters, which functions as <b>dssPoolID</b>, that is, the ID of the DSS storage pool.</p>
cluster_type	String	<p>Storage class of the cloud server system disk. The value is always <b>dss</b>. This field is used only for DeC clusters.</p>

Parameter	Type	Description
hw:passthrough	Boolean	<ul style="list-style-type: none"> <li>Pay attention to this field if your ECS is SDI-compliant. If the value of this field is <b>true</b>, the created disk is of SCSI type.</li> <li>If the node pool type is <b>ElasticBMS</b>, this field must be set to <b>true</b>.</li> <li>If a node specification involves both local and EVS disks, configure the disk initialization parameters. For details, see <a href="#">Attaching Disks to a Node</a>.</li> </ul>
metadata	<b>VolumeMetadata</b> object	EVS disk encryption information. This field is mandatory only when you need to encrypt the system disk or data disks of the node to be created.

**Table 4-536** VolumeMetadata

Parameter	Type	Description
__system__encrypted	String	Whether the EVS disk is encrypted. The value <b>0</b> indicates that the EVS disk is not encrypted, and the value <b>1</b> indicates that the EVS disk is encrypted. If this parameter is not specified, EVS disks will not be encrypted by default.
__system__cmkid	String	CMK ID, which indicates encryption in <b>metadata</b> . This field is used with <b>__system__encrypted</b> .

**Table 4-537** Storage

Parameter	Type	Description
storageSelectors	Array of <b>StorageSelectors</b> objects	Disk selection. Matched disks are managed according to <b>matchLabels</b> and <b>storageType</b> .
storageGroups	Array of <b>StorageGroups</b> objects	A storage group consists of multiple storage devices. It is used to divide storage space.

**Table 4-538** StorageSelectors

Parameter	Type	Description
name	String	Selector name, used as the index of <b>selectorNames</b> in <b>storageGroup</b> . Therefore, the name of each selector must be unique.
storageType	String	Storage type. Only <b>evs</b> (EVS disks) and <b>local</b> (local disks) are supported. The local storage does not support disk selection. All local disks will form a VG, so only one storageSelector of the local type is allowed.
matchLabels	<b>matchLabels</b> object	Matching field of an EVS disk. The value can be <b>size</b> , <b>volumeType</b> , <b>metadataEncrypted</b> , <b>metadataCmkid</b> , or <b>count</b> in <b>DataVolume</b> .

**Table 4-539** matchLabels

Parameter	Type	Description
size	String	Matched disk size. If this parameter is left unspecified, the disk size is not limited. Example: 100
volumeType	String	EVS disk type. <b>SSD</b> , <b>GPSSD</b> , <b>SAS</b> , <b>ESSD</b> , and <b>SATA</b> are supported. If this parameter is left blank, there is no restriction on the disk type.
metadataEncrypted	String	Disk encryption identifier. If this parameter is set to <b>0</b> , it specifies that the disk is not encrypted. If this parameter is set to <b>1</b> , it specifies that the disk is encrypted. If this parameter is left blank, there is no restriction on the disk encryption identifier.
metadataCmkid	String	CMK ID of an encrypted disk. The value is a string of 36 bytes. If this parameter is left blank, there is no restriction on the disk key ID.
count	String	Number of disks to be selected. If this parameter is left blank, all disks of this type are selected.

**Table 4-540** StorageGroups

Parameter	Type	Description
name	String	Name of a virtual storage group, which must be unique. <b>NOTE</b> <ul style="list-style-type: none"> <li>If <b>cceManaged</b> is set to <b>true</b>, the name must be set to <b>vgpaas</b>.</li> <li>If the data disk is used as a temporary storage volume, the name must be <b>vg-everest-localvolume-ephemeral</b>.</li> <li>If the data disk is used as a persistent storage volume, the name must be <b>vg-everest-localvolume-persistent</b>.</li> </ul>
cceManaged	Boolean	Storage space for Kubernetes and runtime components. Only one group can be set to <b>true</b> . If this parameter is left blank, the default value <b>false</b> is used.
selectorNames	Array of strings	Corresponds to <b>name</b> in <b>storageSelectors</b> . One group can select multiple selectors, but one selector can be selected by only one group.
virtualSpaces	Array of <a href="#">VirtualSpace</a> objects	Detailed management of space configuration in a group.

**Table 4-541** VirtualSpace

Parameter	Type	Description
name	String	Name of a virtualSpace. Options: <b>share</b> , <b>kubernetes</b> , <b>runtime</b> , and <b>user</b> <ul style="list-style-type: none"> <li><b>share</b>: shared space configuration (with the runtime and Kubernetes partitions cancelled). <b>lvmConfig</b> needs to be configured.</li> <li><b>Kubernetes</b>: Kubernetes space configuration. <b>lvmConfig</b> needs to be configured.</li> <li><b>runtime</b>: runtime space configuration. <b>runtimeConfig</b> needs to be configured.</li> <li><b>user</b>: user space configuration. <b>lvmConfig</b> needs to be configured.</li> </ul>

Parameter	Type	Description
size	String	Size of a virtualSpace. The value must be an integer in percentage. Example: 90%. <b>NOTE</b> The value ranges from 10 to 90. The sum of percentages of all virtualSpaces in a group cannot exceed 100%.
lvmConfig	<b>LVMConfig</b> object	LVM configuration managements, applicable to the <i>share*</i> , <b>kubernetes</b> , and <b>user</b> spaces. Note that one virtual space supports only one config.
runtimeConfig	<b>RuntimeConfig</b> object	runtime configurations, applicable to the <b>runtime</b> space. Note that one virtual space supports only one config.

**Table 4-542** LVMConfig

Parameter	Type	Description
lvType	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.
path	String	Disk mount path. This parameter takes effect only in user configuration. The value is an absolute path. Digits, letters, periods (.), hyphens (-), and underscores (_) are allowed. <b>NOTE</b> Do not leave the path empty or set it to a key OS path such as the root directory. Key OS paths that cannot be specified: /, /home, /home/, /bin, /bin/, /lib, /lib/, /root, /root/, /boot, /boot/, /dev, /dev/, /etc, /etc/, /lost+found, /lost+found/, /mnt, /mnt/, /proc, /proc/, /sbin, /sbin/, /srv, /srv/, /tmp, /tmp/, /var, /var/, /media, /media/, /opt, /opt/, /selinux, /selinux/, /sys, /sys/, /usr, /usr/, /opt/cloud/, /mnt/paas/, /home/paas/, /var/paas/, /var/lib/, /var/script/

**Table 4-543** RuntimeConfig

Parameter	Type	Description
lvType	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.



**Table 4-544** NodePublicIP

Parameter	Type	Description
ids	Array of strings	IDs of existing EIPs. The quantity cannot be greater than the number of nodes to be created. <b>NOTE</b> If <b>ids</b> has been set, you do not need to set <b>count</b> and <b>eip</b> .
count	Integer	Number of EIPs to be dynamically created. <b>NOTE</b> <b>count</b> and <b>eip</b> must be set at the same time.
eip	<b>NodeEIPSpec</b> object	EIP configuration. <b>NOTE</b> This parameter is not supported when creating a node pool.

**Table 4-545** NodeEIPSpec

Parameter	Type	Description
iptype	String	EIP type, specified in <b>publicip.type</b> in the API for assigning an EIP.
bandwidth	<b>NodeBandwidth</b> object	Bandwidth parameters of the EIP

**Table 4-546** NodeBandwidth

Parameter	Type	Description
chargemode	String	<p>Bandwidth billing mode.</p> <ul style="list-style-type: none"> <li>• If this field is not specified, the billing is based on bandwidth.</li> <li>• If the field is null, the billing is based on bandwidth.</li> <li>• If the field value is <b>traffic</b>, the billing is based on traffic.</li> <li>• If the value is out of the preceding options, the cloud server will fail to be created.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• Billed by bandwidth: The billing will be based on the data transmission rate (in Mbps) of public networks. This billing mode is recommended if your bandwidth usage is higher than 10%.</li> <li>• Billed by traffic: The billing is based on the total amount of data (in GB) transmitted over the public network. This mode is available only when you are creating a pay-per-use node. This billing mode is recommended if your bandwidth usage is lower than 10%.</li> </ul>
size	Integer	Bandwidth size, specified in <b>bandwidth.size</b> in the API for assigning an EIP.
sharetype	String	Bandwidth sharing type. Value options: <b>PER</b> (exclusive bandwidth)

**Table 4-547** NodeNicSpec

Parameter	Type	Description
primaryNic	<b>NicSpec</b> object	Description of the primary NIC.
extNics	Array of <b>NicSpec</b> objects	<p>Extension NIC</p> <p><b>NOTE</b> This parameter is not supported when you add a node to a node pool.</p>

**Table 4-548** NicSpec

Parameter	Type	Description
subnetId	String	Network ID of the subnet to which a NIC belongs. If <b>subnetId</b> is not specified when a primary NIC is creating, the cluster subnet will be used. If <b>subnetList</b> is also configured for a node pool, the <b>subnetList</b> field is used for adding subnets to the node pool. When creating an extension NIC, you must specify <b>subnetId</b> .
fixedIps	Array of strings	The IP address of the primary ENI is specified using <b>fixedIps</b> . The number of IP addresses cannot be greater than the number of created nodes. Either <b>fixedIps</b> or <b>ipBlock</b> can be specified. <b>fixedIps</b> cannot be specified for extension ENIs.
ipBlock	String	CIDR format of the primary NIC IP range. The IP address of the created node falls in this range. <b>fixedIps</b> and <b>ipBlock</b> cannot be specified at the same time.
subnetList	Array of strings	Network ID list of the subnet where the NIC resides. A maximum of 20 subnets can be configured for a node pool.

**Table 4-549** Taint

Parameter	Type	Description
key	String	Key.
value	String	Value.
effect	String	Effect.

**Table 4-550** UserTag

Parameter	Type	Description
key	String	Key of the cloud server label. The value cannot start with <b>CCE-</b> or <b>__type_baremetal</b> .
value	String	Value of the cloud server label.

**Table 4-551** Runtime

Parameter	Type	Description
name	String	<p>Container runtime. Default value:</p> <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>

**Table 4-552** NodeExtendParam

Parameter	Type	Description
ecs:performancetype	String	ECS flavor types. This field is returned in the response.
maxPods	Integer	<p>Maximum number of pods that can be created on a node, including the default system pods. Value range: 16 to 256.</p> <p>This limit prevents the node from being overloaded of pods.</p> <p>The number of pods that can be created on a node is determined by multiple parameters. For details, see <a href="#">Maximum Number of Pods That Can Be Created on a Node</a>.</p>

Parameter	Type	Description
DockerLVMConfigOverride	String	<p>Docker data disk configuration item. (This parameter has been discarded. Use the <b>storage</b> field instead.) Example default configuration: "DockerLVMConfigOverride":"dockerThinpool=vgpaas/90%VG;kubernetesLV=vgpaas/10%VG;diskType=evs;lvType=linear"</p> <p>By default, if no VD disk is available, an error occurs because the data disk fails to be found. Set <b>diskType</b> based on the actual drive letter type. The following fields are included:</p> <ul style="list-style-type: none"> <li>• <b>userLV</b> (optional): size of the user space, for example, <b>vgpaas/20%VG</b>.</li> <li>• <b>userPath</b> (optional): mount path of the user space, for example, <b>/home/wqt-test</b>.</li> <li>• <b>diskType</b>: disk type. Currently, only <b>evs</b>, <b>hdd</b>, and <b>ssd</b> are supported.</li> <li>• <b>lvType</b>: type of a logic volume. The value can be <b>linear</b> or <b>striped</b>.</li> <li>• <b>dockerThinpool</b>: Docker space size, for example, <b>vgpaas/60%VG</b>.</li> <li>• <b>kubernetesLV</b>: kubelet space size, for example, <b>vgpaas/20%VG</b>.</li> </ul>
dockerBaseSize	Integer	<p>Available disk space of a single container on a node, in GB. This parameter has been discarded. Use <b>containerBaseSize</b>.</p> <p>If this parameter is left blank or is set to <b>0</b>, the default value is used. In Device Mapper mode, the default value is <b>10</b>. In OverlayFS mode, the available space of a single container is not limited by default, and the <b>dockerBaseSize</b> setting takes effect only on nodes running EulerOS in clusters of the new versions.</p> <p>For details about how to allocate the space for the container runtime, see <a href="#">Data Disk Space Allocation</a>.</p> <p>When Device Mapper is used, it is recommended that <b>dockerBaseSize</b> be set to a value less than or equal to 80 GB. If the value is too large, the container runtime may fail to be started due to long initialization. If there are special requirements for the container disk space, you can mount an external or local storage device.</p>

Parameter	Type	Description
containerBaseSize	Integer	<p>Available disk space of a single container on a node, in GiB.</p> <p>If this parameter is left blank or is set to <b>0</b>, the default value is used. In OverlayFS mode, the available space of a single container is not limited by default. In Device Mapper mode, the default value is <b>10</b>, and the <b>containerBaseSize</b> setting takes effect only on nodes running EulerOS in clusters v1.23.14-r0, v1.25.9-r0, v1.27.6-r0, v1.28.4-r0 or later.</p> <p>For details about how to allocate the space for the container runtime, see <a href="#">Data Disk Space Allocation</a>.</p> <p>When Device Mapper is used, it is recommended that <b>containerBaseSize</b> be set to a value less than or equal to 80 GB. If the value is too large, the container runtime may fail to be started due to long initialization. If there are special requirements for the container disk space, you can mount an external or local storage device. In new versions, the Device Mapper mode is used only by BMSs in the same resource pool and is being discarded.</p> <p><b>NOTE</b> This parameter cannot be updated when a node pool is being updated.</p>
publicKey	String	Public key of a node.
alpha.cce/preInstall	String	<p>Pre-installation scripts. The total number of characters after transcoding cannot exceed 10240.</p> <p><b>NOTE</b> The input value must be encoded using Base64. The method is as follows:</p> <pre>echo -n "Content to be encoded"   base64</pre>
alpha.cce/postInstall	String	<p>Post-installation scripts. The total number of characters after transcoding cannot exceed 10240.</p> <p><b>NOTE</b> The input value must be encoded using Base64. The method is as follows:</p> <pre>echo -n "Content to be encoded"   base64</pre>
alpha.cce/NodeImageID	String	This parameter is required when a custom image is used to create a BMS node.

Parameter	Type	Description
chargingMode	Integer	Billing mode of a node. This parameter has been deprecated. Use the <b>billingMode</b> parameter in <b>NodeSpec</b> .
agency_name	String	Name of an agency An agency is created by a tenant administrator on Identity and Access Management (IAM) to provide temporary credentials for CCE nodes to access cloud servers. This parameter is returned only when it is transferred during node creation.
kubeReservedMem	Integer	Reserved node memory, which is reserved for Kubernetes components.
systemReservedMem	Integer	Reserved node memory, which is reserved for system components.
init-node-password	String	Node password. If this parameter is used as a response parameter, asterisks (*) will be displayed.

**Table 4-553** HostnameConfig

Parameter	Type	Description
type	String	Configuration type of the Kubernetes node name. The default value is <b>privatelp</b> . <ul style="list-style-type: none"> <li>• <b>privatelp</b>: The Kubernetes node is named after its IP address.</li> <li>• <b>cceNodeName</b>: The Kubernetes node is named after the CCE node.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• For a node which is configured using <b>cceNodeName</b>, the name is the same as the Kubernetes node name and the ECS name. The node name cannot be changed. If the ECS name is changed on the ECS console, the node name will retain unchanged after ECS synchronization.</li> <li>• For a node which is configured using <b>cceNodeName</b>, to avoid a conflict between Kubernetes nodes, the system automatically adds a suffix to each node name. The suffix is in the format of <b>A hyphen (-) Five random characters</b>. The value of the random characters is a lowercase letter or a digit ranging from 0 to 9.</li> </ul>

**Table 4-554** NodePoolNodeAutoscaling

Parameter	Type	Description
enable	Boolean	Whether to enable auto scaling.
minNodeCount	Integer	Minimum number of nodes allowed if auto scaling is enabled. The value cannot be greater than the maximum number of nodes allowed by the cluster specifications.
maxNodeCount	Integer	Maximum number of nodes allowed if auto scaling is enabled. This value must be greater than or equal to the value of <b>minNodeCount</b> and cannot exceed the maximum number of nodes in the cluster specifications.
scaleDownCooldownTime	Integer	Interval between two scaling operations, in minutes. During this period, nodes added after a scale-up will not be deleted.
priority	Integer	Weight of a node pool. A node pool with a higher weight has a higher priority during scaling.

**Table 4-555** NodeManagement

Parameter	Type	Description
serverGroupReference	String	Cloud server group ID. If this field is specified, all nodes in the node pool will be created in this group. The group ID can be specified only when you create the node pool and cannot be modified. When you specify a cloud server group, the number of nodes in the node pool cannot exceed the group quota.

**Table 4-556** SecurityID

Parameter	Type	Description
id	String	Security group ID

**Table 4-557** ExtensionScaleGroup

Parameter	Type	Description
metadata	<a href="#">ExtensionScaleGroupMetadata</a> object	Basic information about an extended scaling group



Parameter	Type	Description
spec	<a href="#">ExtensionScaleGroupSpec</a> object	Configurations of an extended scaling group, which carry different configurations from those of the default scaling group

**Table 4-558** ExtensionScaleGroupMetadata

Parameter	Type	Description
uid	String	UUID of an extended scaling group, which is automatically generated by the system.
name	String	Name of an extended scaling group. The value cannot be <b>default</b> and can contain a maximum of 55 characters. Only digits, lowercase letters, and hyphens (-) are allowed.

**Table 4-559** ExtensionScaleGroupSpec

Parameter	Type	Description
flavor	String	Node flavor
az	String	AZ of a node. If this parameter is not specified or left blank, the default scaling group configurations take effect.
capacityReservationSpecification	<a href="#">CapacityReservationSpecification</a> object	Capacity reservation configurations of an extended scaling group
autoscaling	<a href="#">ScaleGroupAutoscaling</a> object	Auto scaling configurations of an extended scaling group

**Table 4-560** CapacityReservationSpecification

Parameter	Type	Description
id	String	Private pool ID. The parameter value can be ignored when <b>preference</b> is set to <b>none</b> .
preference	String	Capacity of a private storage pool. If the value is <b>none</b> , the capacity reservation is not specified. If the value is <b>targeted</b> , the capacity reservation is specified. In this case, the ID cannot be left blank.

**Table 4-561** ScaleGroupAutoscaling

Parameter	Type	Description
enable	Boolean	Whether to enable auto scaling for a scaling group. This function is disabled by default.
extensionPriority	Integer	Priority of a scaling group. If no value is specified for this parameter, it defaults to <b>0</b> . A higher value indicates a greater priority.
minNodeCount	Integer	Minimum number of nodes in a scaling group during auto scaling. The value must be greater than <b>0</b> .
maxNodeCount	Integer	Maximum number of nodes that can be retained in a scaling group during auto scaling. The value must be greater than or equal to that of <b>minNodeCount</b> , and can neither be greater than the maximum number of nodes allowed by the cluster nor the maximum number of nodes in the node pool.

**Table 4-562** NodePoolStatus

Parameter	Type	Description
currentNode	Integer	Total number of nodes in the current node pool (excluding the nodes that are being deleted)
creatingNode	Integer	Number of nodes in the creation process in the node pool
deletingNode	Integer	Number of nodes being deleted in the current node pool.

Parameter	Type	Description
phase	String	<p>Node pool status.</p> <ul style="list-style-type: none"> <li>• Null: available (the number of current nodes in the node pool has reached the expected value, and no node scaling is being performed.)</li> <li>• Synchronizing: scaling in progress (the number of current nodes in the node pool does not reach the expected value and no node scaling is being performed.)</li> <li>• Synchronized: pending scaling (the number of current nodes in the node pool does not reach the expected value, or node scaling is being performed.)</li> <li>• SoldOut: The node pool cannot be scaled out. (This field is used in multiple scenarios, for example, resources in the node pool have been sold out or the resource quota is insufficient.)</li> </ul> <p><b>NOTE</b> This node pool status has been discarded and is reserved only for compatibility. You are not advised to use it. The replacements are as follows:</p> <ul style="list-style-type: none"> <li>• Node pool scaling status: You can obtain accurate status of the current node pool based on parameters such as <b>currentNode</b>, <b>creatingNode</b>, or <b>deletingNode</b>.</li> <li>• Node pool scale-out: You can use <b>conditions</b> to obtain the detailed status of a node pool. <b>Scalable</b> can replace <b>SoldOut</b>.</li> <li>• <b>Deleting</b>: The object is being deleted.</li> <li>• <b>Error</b>: An error occurs.</li> </ul>
jobId	String	ID of a job executed on the node pool. This field is available only when the node pool is being deleted.
conditions	Array of <a href="#">NodePoolCondition</a> objects	Node pool status details. For details, see the definition of <b>Condition</b> .
scaleGroupStatuses	Array of <a href="#">ScaleGroupStatus</a> objects	Detailed status of a scaling group. For details, see the definition of <b>ScaleGroupStatus</b> .

**Table 4-563** ScaleGroupStatus

Parameter	Type	Description
name	String	Scaling group name
uid	String	UUID of a scaling group
creationTimes tamp	String	The time when a scaling group was created
updateTimest amp	String	The time when a scaling group was updated
phase	String	<p>Status of a scaling group.</p> <ul style="list-style-type: none"> <li>• <b>Null:</b> The scaling group is available. (The number of nodes in the current scaling group has reached the expected value, and no node scaling is being performed.)</li> <li>• <b>Synchronizing:</b> The scaling group is being scaled. (The number of nodes in the current scaling group has not reached the expected value, and no node scaling is being performed.)</li> <li>• <b>Synchronized:</b> The scaling is pending. (The number of nodes in the current scaling group has not reached the expected value, or node scaling is being performed.)</li> <li>• <b>SoldOut:</b> The scaling group cannot be scaled out. (This field is used in multiple scenarios, for example, resources in the node pool have been sold out or the resource quota is insufficient.)</li> </ul> <p><b>NOTE</b> This scaling group status has been deprecated. It is reserved only for compatibility. Do not use it anymore. Use the following instead:</p> <ul style="list-style-type: none"> <li>• Scaling status of a scaling group: You can obtain the status of the current scaling group using parameters such as <b>desiredNodeCount</b>, <b>existingNodeCount</b>, or <b>upcomingNodeCount</b>.</li> <li>• Scaling group scale-out: You can use <b>conditions</b> to obtain the detailed status of a scaling group. <b>Scalable</b> can replace <b>SoldOut</b>.</li> <li>• <b>Deleting:</b> The scaling group being deleted.</li> <li>• <b>Error:</b> An error occurred in the scaling group.</li> </ul>
desiredNodeC ount	Integer	Expected number of nodes in a scaling group
unpaidScaleN odeCount	Integer	Number of unpaid nodes in an order

Parameter	Type	Description
existingNodeCount	<a href="#">existingNodeCount</a> object	Statistics on existing nodes in a scaling group
upcomingNodeCount	<a href="#">upcomingNodeCount</a> object	Statistics on the nodes to be created in a scaling group
scaleDownDisabledNodeCount	Integer	Number of nodes that cannot be deleted in a scaling group
conditions	Array of <a href="#">NodePoolCondition</a> objects	Scaling group status details. For details, see the definition of <b>Condition</b> .

**Table 4-564** existingNodeCount

Parameter	Type	Description
postPaid	Integer	Number of pay-per-use nodes
total	Integer	Total number of pay-per-use nodes

**Table 4-565** upcomingNodeCount

Parameter	Type	Description
postPaid	Integer	Number of pay-per-use nodes
total	Integer	Total number of pay-per-use nodes

**Table 4-566** NodePoolCondition

Parameter	Type	Description
type	String	Condition type. The options are as follows: <ul style="list-style-type: none"> <li>• <b>Scalable</b>: whether a node pool can be scaled. If the status is <b>False</b>, node pool scaling will not be triggered again.</li> <li>• <b>QuotaInsufficient</b>: Quotas on which node pool scaling depends are insufficient, affecting the node pool scaling status.</li> <li>• <b>ResourceInsufficient</b>: Resources on which node pool scaling depends are insufficient, affecting the node pool scaling status.</li> <li>• <b>UnexpectedError</b>: The node pool fails to be scaled out due to unexpected reasons, affecting the node pool scaling status.</li> <li>• <b>Error</b>: A node pool error occurs. A common trigger is deletion failure.</li> </ul>
status	String	Current status of <b>Condition</b> . The options are as follows: <ul style="list-style-type: none"> <li>• "True"</li> <li>• "False"</li> </ul>
lastProbeTime	String	Time when the status was last checked
lastTransitTime	String	Time when the status was last changed
reason	String	Reason why the status was last changed
message	String	Detailed condition description

## Example Requests

None

## Example Responses

**Status code: 200**

Details about the specified node pool are obtained successfully.

```
{
  "kind": "NodePool",
  "apiVersion": "v3",
  "metadata": {
    "name": "lc-it-nodepool-79796",
    "uid": "99addaa2-69eb-11ea-a592-0255ac1001bb"
  },
  "spec": {
    "type": "vm",
    "nodeTemplate": {
```

```

"flavor" : "s6.large.2",
"az" : "*****",
"os" : "EulerOS 2.5",
"login" : {
  "sshKey" : "KeyPair-001"
},
"rootVolume" : {
  "volumetype" : "SAS",
  "size" : 40
},
"dataVolumes" : [ {
  "volumetype" : "SAS",
  "size" : 100,
  "extendParam" : {
    "useType" : "docker"
  }
} ],
"publicIP" : {
  "eip" : {
    "bandwidth" : { }
  }
},
"nodeNicSpec" : {
  "primaryNic" : {
    "subnetId" : "7e767d10-7548-4df5-ad72-aeac1d08bd8a"
  }
},
"billingMode" : 0,
"extendParam" : {
  "maxPods" : 110
},
"k8sTags" : {
  "cce.cloud.com/cce-nodepool" : "lc-it-nodepool-79796"
},
"autoscaling" : { },
"nodeManagement" : { }
},
"status" : {
  "phase" : "Deleting",
  "jobId" : "3281fa02-69ee-11ea-a592-0255ac1001bb"
}
}

```

## Status Codes

Status Code	Description
200	Details about the specified node pool are obtained successfully.

## Error Codes

See [Error Codes](#).

### 4.4.3 Listing All Node Pools in a Specified Cluster

#### Function

This API is used to obtain information about all node pools in a specified cluster.

 NOTE

- The URL for cluster management is in the format of **https://Endpoint/uri**. In the URL, **uri** indicates the resource path, that is, the path for API access.
- A node pool is a group of nodes within a cluster that all have the same configuration.

## URI

GET /api/v3/projects/{project\_id}/clusters/{cluster\_id}/nodepools

**Table 4-567** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

**Table 4-568** Query Parameters

Parameter	Mandatory	Type	Description
showDefaultNodePool	No	String	Whether to display the default node pool. By default, the default node pool is not displayed. If this parameter is set to <b>true</b> , the default node pool is displayed.

## Request Parameters

**Table 4-569** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).



Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

## Response Parameters

Status code: 200

**Table 4-570** Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed to List.
apiVersion	String	API version. The value is fixed to v3.
items	Array of <a href="#">NodePoolResp</a> objects	/

**Table 4-571** NodePoolResp

Parameter	Type	Description
kind	String	API type. The value is fixed at <b>NodePool</b> .
apiVersion	String	API version. The value is fixed at <b>v3</b> .
metadata	<a href="#">NodePoolMetadata</a> object	Metadata of the node pool.
spec	<a href="#">NodePoolSpec</a> object	Node pool specifications. <b>NOTE</b> <ul style="list-style-type: none"> <li>This field is optional for updating a node pool.</li> </ul>
status	<a href="#">NodePoolStatus</a> object	Node pool status.

**Table 4-572** NodePoolMetadata

Parameter	Type	Description
name	String	Node pool name. <b>NOTE</b> Naming rules: Enter 1 to 50 characters, starting with a lowercase letter and not ending with a hyphen (-). Only lowercase letters, digits, and hyphens (-) are allowed. <ul style="list-style-type: none"> <li>You cannot create node pools named DefaultPool.</li> </ul>
uid	String	UID of the node pool. The value is automatically generated after the object is updated. A user-defined value will not take effect.
annotations	Map<String,String>	Annotations of a node pool in key-value pairs. This parameter is used only for query and cannot be input through a request. Entered data of this parameter is invalid.
updateTimestamp	String	Update time.
creationTimestamp	String	Creation time.

**Table 4-573** NodePoolSpec

Parameter	Type	Description
type	String	Node pool type. If this parameter is left blank, the value <b>vm</b> is used by default. <ul style="list-style-type: none"> <li><b>vm</b>: ECS</li> <li><b>ElasticBMS</b>: C6 general computing-plus BMS. An example flavor is <b>c6.22xlarge.2.physical</b>.</li> <li><b>bms</b>: BMS</li> </ul>
nodeTemplate	<a href="#">NodeSpec</a> object	Detailed parameters of the node pool template.
initialNodeCount	Integer	Initial number of nodes for the node pool. When queried, the value is the number of target nodes in the node pool.
autoscaling	<a href="#">NodePoolNodeAutoscaling</a> object	Auto scaling parameters

Parameter	Type	Description
nodeManagement	<a href="#">NodeManagement</a> object	Node management configuration
podSecurityGroups	Array of <a href="#">SecurityID</a> objects	Security group configuration. This parameter is available only for CCE Turbo clusters.
extensionScaleGroups	Array of <a href="#">ExtensionScaleGroup</a> objects	Configuration list of the extended scaling groups in a node pool. For details, see the definition of <b>ExtensionScaleGroup</b> .
customSecurityGroups	Array of strings	Custom security group settings for a node pool. New nodes scaled out in a node pool can be bound to a specified security group. <ul style="list-style-type: none"> <li>• Specifying no security group ID will add the new nodes to the default security group of the worker nodes.</li> <li>• Specifying a valid security group ID will put new nodes in that security group.</li> <li>• When specifying a security group, do not modify the rules of the port on which CCE running depends.</li> </ul>

**Table 4-574** NodeSpec

Parameter	Type	Description
flavor	String	Node specifications. For details about the node specifications supported by CCE, see <a href="#">Node Flavor Description</a> .
az	String	AZ where the node to be created is located. You need to specify the AZ name. A node created using APIs does not support random AZs. <b>NOTE</b> <ul style="list-style-type: none"> <li>• When you are creating a node pool and configuring a scaling group, this parameter cannot be set to <b>random</b>.</li> </ul>

Parameter	Type	Description
os	String	<p>Node OS. For details about the supported OSs, see <a href="#">Node OS</a>.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>The system automatically selects the supported OS based on the cluster version. If the current cluster version does not support the OS, an error will be reported.</li> <li>If <code>alpha.cce/NodeImageID</code> in <code>extendParam</code> is specified during node creation, you do not need to configure this parameter.</li> <li>This parameter is mandatory when creating a node pool.</li> <li>If you are using shared disk space when creating a node, this parameter is mandatory. Specifically, the disk initialization configuration management parameter <code>storage</code> is used and the <code>name</code> field of <code>virtualSpaces</code> in <code>StorageGroups</code> is set to <code>share</code>.</li> </ul>
login	<a href="#">Login</a> object	<p>Node login mode. Either the key pair or password must be used for login.</p> <p><b>NOTE</b> This parameter cannot be modified when a node pool is being updated.</p>
rootVolume	<a href="#">Volume</a> object	Information about disks on the node
dataVolumes	Array of <a href="#">Volume</a> objects	<p>Data disk parameters of a node. A VM can have up to 16 disks attached, while a BMS can only have up to 10 disks attached. The maximum number of disks that can be attached to either a VM or BMS is limited by their respective maximum disk limits. (You can add more data disks to each of your nodes through the CCE console and APIs.) The data disk used by the container runtime and kubelet cannot be detached, or the node will become unavailable. For DeC nodes, the parameter details are the same as those for <b>rootVolume</b>.</p>

Parameter	Type	Description
storage	<a href="#">Storage</a> object	<p>Disk initialization management parameter. This parameter is complex to configure. For details, see <a href="#">Attaching Disks to a Node</a>.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If a node has both local and EVS disks attached, make sure to not keep the default value for this parameter, or it may result in unexpected disk partitions.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If you want to change the value range of a data disk to 20 to 32768, do not use the default value.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If you want to use the shared disk space (with the runtime and Kubernetes partitions cancelled), do not retain the default value of this parameter. For details about the shared disk space, see <a href="#">Data Disk Space Allocation</a>.</li> </ul>
publicIP	<a href="#">NodePublicIP</a> object	<p>EIP of a node.</p> <p><b>NOTE</b></p> <p>This parameter is not supported when you add a node to a node pool.</p>
nodeNicSpec	<a href="#">NodeNicSpec</a> object	NIC of the node
count	Integer	Number of nodes to be created in a batch. The value must be a positive integer greater than or equal to 1 and less than or equal to the defined limit. This parameter can be left blank when it is used for a node pool.
billingMode	Integer	<p>Node billing mode.</p> <ul style="list-style-type: none"> <li><b>0</b>: pay-per-use</li> </ul>

Parameter	Type	Description
taints	Array of <b>Taint</b> objects	<p>You can add taints to created nodes to configure anti-affinity. A maximum of 20 taints can be added. Each taint contains the following parameters:</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain can be prefixed to a key.</li> <li>• <b>Value:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed.</li> <li>• <b>Effect:</b> Available options are <b>NoSchedule</b>, <b>PreferNoSchedule</b>, and <b>NoExecute</b>. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</li> </ul> <p>Example:</p> <pre data-bbox="815 1108 1430 1346">"taints": [{   "key": "status",   "value": "unavailable",   "effect": "NoSchedule" }, {   "key": "looks",   "value": "bad",   "effect": "NoSchedule" }]</pre>

Parameter	Type	Description
k8sTags	Map<String,String>	<p>Defined in key-value pairs. A maximum of 20 key-value pairs are allowed.</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain containing a maximum of 253 characters (such as <b>example.com/my-key</b>) can be prefixed to a key.</li> <li>• <b>Value:</b> The value can be left blank or contain 1 to 63 characters that start and end with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</li> </ul> <p>Example:</p> <pre>"k8sTags": {   "key": "value" }</pre>
ecsGroupId	String	<p>ECS group ID. If this parameter is configured, nodes will be created in the specified ECS group.</p> <p><b>NOTE</b> This configuration does not take effect when a node pool is created. To ensure that all nodes in a node pool are in the same ECS group, configure <b>nodeManagement</b> in the node pool.</p>
dedicatedHostId	String	<p>ID of the DeH to which the node is scheduled.</p> <p><b>NOTE</b> This parameter is not supported when you add a node during node pool creation.</p>

Parameter	Type	Description
userTags	Array of <a href="#">UserTag</a> objects	<p>Cloud server tag (resource tag). The key of a tag must be unique. The maximum number of custom tags supported by CCE varies depending on regions and cannot exceed 8. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</p> <p><b>NOTE</b> A tag key can have a maximum of 36 characters. It can only contain letters, digits, Unicode characters, and special characters (-_).</p>
runtime	<a href="#">Runtime</a> object	<p>Container runtime. Default value:</p> <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>



Parameter	Type	Description
initializedConditions	Array of strings	<p>Custom initialization flag, which is left blank by default.</p> <p>Before CCE nodes are initialized, <b>node.cloudprovider.kubernetes.io/uninitialized</b> is added to the nodes to prevent pods from being scheduled to these nodes. When creating a node, you can configure the <b>initializedConditions</b> parameter to control the taint removal time. By default, the timeout period is not configured.</p> <p>Example:</p> <ol style="list-style-type: none"> <li>1. Create a node and add parameter "initializedConditions": ["CCEInitial", "CustomedInitial"].</li> <li>2. After custom initialization is complete, CCE will call a Kubernetes API (for example, <b>PATCH /v1/nodes/{node_ip}/status</b>) to update the node <b>conditions</b> by adding two labels of types <b>CCEInitial</b> and <b>CustomedInitial</b> and setting the status to <b>True</b>, as shown in the following:</li> </ol> <pre data-bbox="815 1055 1430 1205"> status:   conditions:   - type: CCEInitial     status: 'True'   - type: CustomedInitial     status: 'True'                     </pre> <ol style="list-style-type: none"> <li>1. CCE polls <b>status.Conditions</b> of nodes to check whether there are conditions of types <b>CCEInitial</b> and <b>CustomedInitial</b>. If such conditions exist and the status is <b>True</b>, the node initialization is complete and the initialization taint is removed.</li> <li>2. <b>initializedConditions</b> allows you to configure a timeout period during node creation. For example, "initializedConditions": ["CCEInitial:15m", "CustomedInitial:15m"] indicates that the timeout period is 15 minutes. After the timeout period is reached, the initialization conditions will be automatically ignored and the initialization taint will be removed when CCE polls the node. <ul style="list-style-type: none"> <li>● The timeout interval in <b>initializedConditions</b> ranges from 1 to 99 seconds.</li> <li>● Use only letters and digits. Max. characters: 20.</li> <li>● Max. flags: 2.</li> </ul> </li> </ol>

Parameter	Type	Description
		<ul style="list-style-type: none"> <li>The unit of the timeout period is minute (m).</li> </ul>
extendParam	<a href="#">NodeExtendParam</a> object	Extended parameters for creating a node.
hostnameConfig	<a href="#">HostnameConfig</a> object	Kubernetes node name configuration parameter, which is supported by clusters of v1.23.6-r0 to v1.25 or clusters of v1.25.2-r0 or later versions.
serverEnterpriseProjectID	String	Enterprise project ID of a server. CCE does not support features provided by EPS. This field is used only for synchronizing enterprise project IDs of servers. To create a node or node pool, you can specify an existing enterprise project for it. If this field is left blank, it inherits the enterprise project attributes of the cluster that runs the node or node pool. When updating a node pool, any configuration modifications you made will apply only to new nodes. You need to modify the enterprise projects of existing nodes on the EPS console. If no value is specified during the update, this field will not be updated. If this parameter is left blank, the cluster enterprise project is returned.

**Table 4-575** Login

Parameter	Type	Description
sshKey	String	Name of the key pair used for login.
userPassword	<a href="#">UserPassword</a> object	Password used for node login.

**Table 4-576** UserPassword

Parameter	Type	Description
username	String	Login account. The default value is <b>root</b> .

Parameter	Type	Description
password	String	<p>Login password. If a username and password are used to create a node, this field is shielded in the response body. A password must meet the following complexity requirements:</p> <ul style="list-style-type: none"> <li>• Contains 8 to 26 characters.</li> <li>• Contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters !@\$%^-_=+[{ }];,./?</li> <li>• Cannot contain the username or the username spelled backwards. The <b>password</b> field must be salted during node creation. For details, see <a href="#">Adding a Salt in the password Field When Creating a Node</a>.</li> </ul>

Table 4-577 Volume

Parameter	Type	Description
size	Integer	<p>Disk size, in GiB.</p> <ul style="list-style-type: none"> <li>• System disk: 40 to 1024</li> <li>• Value range for data disks: 100 to 32768</li> </ul>
volumetype	String	<p>Disk type. For details about possible values, see the description of the <b>root_volume</b> parameter in the API used to create an ECS.</p> <ul style="list-style-type: none"> <li>• <b>SAS</b>: high I/O SAS disk</li> <li>• <b>SSD</b>: ultra-high I/O SSD disk</li> <li>• <b>SATA</b>: common I/O SATA disk SATA disks have been removed from EVS. You can find them attached only on existing nodes.</li> </ul>
extendParam	Map<String, Object>	<p>Extended disk parameters, defined in <b>extendparam</b> in the API used to create an ECS.</p>
cluster_id	String	<p>ID of the storage pool used by the ECS system disk. This field is used only for DeC clusters, which functions as <b>dssPoolID</b>, that is, the ID of the DSS storage pool.</p>
cluster_type	String	<p>Storage class of the cloud server system disk. The value is always <b>dss</b>. This field is used only for DeC clusters.</p>

Parameter	Type	Description
hw:passthrough	Boolean	<ul style="list-style-type: none"> <li>Pay attention to this field if your ECS is SDI-compliant. If the value of this field is <b>true</b>, the created disk is of SCSI type.</li> <li>If the node pool type is <b>ElasticBMS</b>, this field must be set to <b>true</b>.</li> <li>If a node specification involves both local and EVS disks, configure the disk initialization parameters. For details, see <a href="#">Attaching Disks to a Node</a>.</li> </ul>
metadata	<b>VolumeMetadata</b> object	EVS disk encryption information. This field is mandatory only when you need to encrypt the system disk or data disks of the node to be created.

**Table 4-578** VolumeMetadata

Parameter	Type	Description
__system__encrypted	String	Whether the EVS disk is encrypted. The value <b>0</b> indicates that the EVS disk is not encrypted, and the value <b>1</b> indicates that the EVS disk is encrypted. If this parameter is not specified, EVS disks will not be encrypted by default.
__system__cmkid	String	CMK ID, which indicates encryption in <b>metadata</b> . This field is used with <b>__system__encrypted</b> .

**Table 4-579** Storage

Parameter	Type	Description
storageSelectors	Array of <b>StorageSelectors</b> objects	Disk selection. Matched disks are managed according to <b>matchLabels</b> and <b>storageType</b> .
storageGroups	Array of <b>StorageGroups</b> objects	A storage group consists of multiple storage devices. It is used to divide storage space.

**Table 4-580** StorageSelectors

Parameter	Type	Description
name	String	Selector name, used as the index of <b>selectorNames</b> in <b>storageGroup</b> . Therefore, the name of each selector must be unique.
storageType	String	Storage type. Only <b>evs</b> (EVS disks) and <b>local</b> (local disks) are supported. The local storage does not support disk selection. All local disks will form a VG, so only one storageSelector of the local type is allowed.
matchLabels	<b>matchLabels</b> object	Matching field of an EVS disk. The value can be <b>size</b> , <b>volumeType</b> , <b>metadataEncrypted</b> , <b>metadataCmkid</b> , or <b>count</b> in <b>DataVolume</b> .

**Table 4-581** matchLabels

Parameter	Type	Description
size	String	Matched disk size. If this parameter is left unspecified, the disk size is not limited. Example: 100
volumeType	String	EVS disk type. <b>SSD</b> , <b>GPSSD</b> , <b>SAS</b> , <b>ESSD</b> , and <b>SATA</b> are supported. If this parameter is left blank, there is no restriction on the disk type.
metadataEncrypted	String	Disk encryption identifier. If this parameter is set to <b>0</b> , it specifies that the disk is not encrypted. If this parameter is set to <b>1</b> , it specifies that the disk is encrypted. If this parameter is left blank, there is no restriction on the disk encryption identifier.
metadataCmkid	String	CMK ID of an encrypted disk. The value is a string of 36 bytes. If this parameter is left blank, there is no restriction on the disk key ID.
count	String	Number of disks to be selected. If this parameter is left blank, all disks of this type are selected.

**Table 4-582** StorageGroups

Parameter	Type	Description
name	String	Name of a virtual storage group, which must be unique. <b>NOTE</b> <ul style="list-style-type: none"> <li>If <b>cceManaged</b> is set to <b>true</b>, the name must be set to <b>vgpaas</b>.</li> <li>If the data disk is used as a temporary storage volume, the name must be <b>vg-everest-localvolume-ephemeral</b>.</li> <li>If the data disk is used as a persistent storage volume, the name must be <b>vg-everest-localvolume-persistent</b>.</li> </ul>
cceManaged	Boolean	Storage space for Kubernetes and runtime components. Only one group can be set to <b>true</b> . If this parameter is left blank, the default value <b>false</b> is used.
selectorNames	Array of strings	Corresponds to <b>name</b> in <b>storageSelectors</b> . One group can select multiple selectors, but one selector can be selected by only one group.
virtualSpaces	Array of <a href="#">VirtualSpace</a> objects	Detailed management of space configuration in a group.

**Table 4-583** VirtualSpace

Parameter	Type	Description
name	String	Name of a virtualSpace. Options: <b>share</b> , <b>kubernetes</b> , <b>runtime</b> , and <b>user</b> <ul style="list-style-type: none"> <li><b>share</b>: shared space configuration (with the runtime and Kubernetes partitions cancelled). <b>lvmConfig</b> needs to be configured.</li> <li><b>Kubernetes</b>: Kubernetes space configuration. <b>lvmConfig</b> needs to be configured.</li> <li><b>runtime</b>: runtime space configuration. <b>runtimeConfig</b> needs to be configured.</li> <li><b>user</b>: user space configuration. <b>lvmConfig</b> needs to be configured.</li> </ul>

Parameter	Type	Description
size	String	Size of a virtualSpace. The value must be an integer in percentage. Example: 90%. <b>NOTE</b> The value ranges from 10 to 90. The sum of percentages of all virtualSpaces in a group cannot exceed 100%.
lvmConfig	<b>LVMConfig</b> object	LVM configuration managements, applicable to the <i>share*</i> , <b>kubernetes</b> , and <b>user</b> spaces. Note that one virtual space supports only one config.
runtimeConfig	<b>RuntimeConfig</b> object	runtime configurations, applicable to the <b>runtime</b> space. Note that one virtual space supports only one config.

**Table 4-584** LVMConfig

Parameter	Type	Description
lvType	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.
path	String	Disk mount path. This parameter takes effect only in user configuration. The value is an absolute path. Digits, letters, periods (.), hyphens (-), and underscores (_) are allowed. <b>NOTE</b> Do not leave the path empty or set it to a key OS path such as the root directory. Key OS paths that cannot be specified: /, /home, /home/, /bin, /bin/, /lib, /lib/, /root, /root/, /boot, /boot/, /dev, /dev/, /etc, /etc/, /lost+found, /lost+found/, /mnt, /mnt/, /proc, /proc/, /sbin, /sbin/, /srv, /srv/, /tmp, /tmp/, /var, /var/, /media, /media/, /opt, /opt/, /selinux, /selinux/, /sys, /sys/, /usr, /usr/, /opt/cloud/, /mnt/paas/, /home/paas/, /var/paas/, /var/lib/, /var/script/

**Table 4-585** RuntimeConfig

Parameter	Type	Description
lvType	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.

**Table 4-586** NodePublicIP

Parameter	Type	Description
ids	Array of strings	IDs of existing EIPs. The quantity cannot be greater than the number of nodes to be created. <b>NOTE</b> If <b>ids</b> has been set, you do not need to set <b>count</b> and <b>eip</b> .
count	Integer	Number of EIPs to be dynamically created. <b>NOTE</b> <b>count</b> and <b>eip</b> must be set at the same time.
eip	<b>NodeEIPSpec</b> object	EIP configuration. <b>NOTE</b> This parameter is not supported when creating a node pool.

**Table 4-587** NodeEIPSpec

Parameter	Type	Description
iptype	String	EIP type, specified in <b>publicip.type</b> in the API for assigning an EIP.
bandwidth	<b>NodeBandwidth</b> object	Bandwidth parameters of the EIP



**Table 4-588** NodeBandwidth

Parameter	Type	Description
chargemode	String	<p>Bandwidth billing mode.</p> <ul style="list-style-type: none"> <li>• If this field is not specified, the billing is based on bandwidth.</li> <li>• If the field is null, the billing is based on bandwidth.</li> <li>• If the field value is <b>traffic</b>, the billing is based on traffic.</li> <li>• If the value is out of the preceding options, the cloud server will fail to be created.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• Billed by bandwidth: The billing will be based on the data transmission rate (in Mbps) of public networks. This billing mode is recommended if your bandwidth usage is higher than 10%.</li> <li>• Billed by traffic: The billing is based on the total amount of data (in GB) transmitted over the public network. This mode is available only when you are creating a pay-per-use node. This billing mode is recommended if your bandwidth usage is lower than 10%.</li> </ul>
size	Integer	Bandwidth size, specified in <b>bandwidth.size</b> in the API for assigning an EIP.
sharetype	String	Bandwidth sharing type. Value options: <b>PER</b> (exclusive bandwidth)

**Table 4-589** NodeNicSpec

Parameter	Type	Description
primaryNic	<b>NicSpec</b> object	Description of the primary NIC.
extNics	Array of <b>NicSpec</b> objects	<p>Extension NIC</p> <p><b>NOTE</b> This parameter is not supported when you add a node to a node pool.</p>

**Table 4-590 NicSpec**

Parameter	Type	Description
subnetId	String	Network ID of the subnet to which a NIC belongs. If <b>subnetId</b> is not specified when a primary NIC is creating, the cluster subnet will be used. If <b>subnetList</b> is also configured for a node pool, the <b>subnetList</b> field is used for adding subnets to the node pool. When creating an extension NIC, you must specify <b>subnetId</b> .
fixedIps	Array of strings	The IP address of the primary ENI is specified using <b>fixedIps</b> . The number of IP addresses cannot be greater than the number of created nodes. Either <b>fixedIps</b> or <b>ipBlock</b> can be specified. <b>fixedIps</b> cannot be specified for extension ENIs.
ipBlock	String	CIDR format of the primary NIC IP range. The IP address of the created node falls in this range. <b>fixedIps</b> and <b>ipBlock</b> cannot be specified at the same time.
subnetList	Array of strings	Network ID list of the subnet where the NIC resides. A maximum of 20 subnets can be configured for a node pool.

**Table 4-591 Taint**

Parameter	Type	Description
key	String	Key.
value	String	Value.
effect	String	Effect.

**Table 4-592 UserTag**

Parameter	Type	Description
key	String	Key of the cloud server label. The value cannot start with <b>CCE-</b> or <b>__type_baremetal</b> .
value	String	Value of the cloud server label.

**Table 4-593** Runtime

Parameter	Type	Description
name	String	<p>Container runtime. Default value:</p> <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>

**Table 4-594** NodeExtendParam

Parameter	Type	Description
ecs:performancetype	String	ECS flavor types. This field is returned in the response.
maxPods	Integer	<p>Maximum number of pods that can be created on a node, including the default system pods. Value range: 16 to 256.</p> <p>This limit prevents the node from being overloaded of pods.</p> <p>The number of pods that can be created on a node is determined by multiple parameters. For details, see <a href="#">Maximum Number of Pods That Can Be Created on a Node</a>.</p>

Parameter	Type	Description
DockerLVMConfigOverride	String	<p>Docker data disk configuration item. (This parameter has been discarded. Use the <b>storage</b> field instead.) Example default configuration: "DockerLVMConfigOverride":"dockerThinpool=vgpaas/90%VG;kubernetesLV=vgpaas/10%VG;diskType=evs;lvType=linear"</p> <p>By default, if no VD disk is available, an error occurs because the data disk fails to be found. Set <b>diskType</b> based on the actual drive letter type. The following fields are included:</p> <ul style="list-style-type: none"> <li>• <b>userLV</b> (optional): size of the user space, for example, <b>vgpaas/20%VG</b>.</li> <li>• <b>userPath</b> (optional): mount path of the user space, for example, <b>/home/wqt-test</b>.</li> <li>• <b>diskType</b>: disk type. Currently, only <b>evs</b>, <b>hdd</b>, and <b>ssd</b> are supported.</li> <li>• <b>lvType</b>: type of a logic volume. The value can be <b>linear</b> or <b>striped</b>.</li> <li>• <b>dockerThinpool</b>: Docker space size, for example, <b>vgpaas/60%VG</b>.</li> <li>• <b>kubernetesLV</b>: kubelet space size, for example, <b>vgpaas/20%VG</b>.</li> </ul>
dockerBaseSize	Integer	<p>Available disk space of a single container on a node, in GB. This parameter has been discarded. Use <b>containerBaseSize</b>.</p> <p>If this parameter is left blank or is set to <b>0</b>, the default value is used. In Device Mapper mode, the default value is <b>10</b>. In OverlayFS mode, the available space of a single container is not limited by default, and the <b>dockerBaseSize</b> setting takes effect only on nodes running EulerOS in clusters of the new versions.</p> <p>For details about how to allocate the space for the container runtime, see <a href="#">Data Disk Space Allocation</a>.</p> <p>When Device Mapper is used, it is recommended that <b>dockerBaseSize</b> be set to a value less than or equal to 80 GB. If the value is too large, the container runtime may fail to be started due to long initialization. If there are special requirements for the container disk space, you can mount an external or local storage device.</p>

Parameter	Type	Description
containerBaseSize	Integer	<p>Available disk space of a single container on a node, in GiB.</p> <p>If this parameter is left blank or is set to <b>0</b>, the default value is used. In OverlayFS mode, the available space of a single container is not limited by default. In Device Mapper mode, the default value is <b>10</b>, and the <b>containerBaseSize</b> setting takes effect only on nodes running EulerOS in clusters v1.23.14-r0, v1.25.9-r0, v1.27.6-r0, v1.28.4-r0 or later.</p> <p>For details about how to allocate the space for the container runtime, see <a href="#">Data Disk Space Allocation</a>.</p> <p>When Device Mapper is used, it is recommended that <b>containerBaseSize</b> be set to a value less than or equal to 80 GB. If the value is too large, the container runtime may fail to be started due to long initialization. If there are special requirements for the container disk space, you can mount an external or local storage device. In new versions, the Device Mapper mode is used only by BMSs in the same resource pool and is being discarded.</p> <p><b>NOTE</b> This parameter cannot be updated when a node pool is being updated.</p>
publicKey	String	Public key of a node.
alpha.cce/preInstall	String	<p>Pre-installation scripts. The total number of characters after transcoding cannot exceed 10240.</p> <p><b>NOTE</b> The input value must be encoded using Base64. The method is as follows:</p> <pre>echo -n "Content to be encoded"   base64</pre>
alpha.cce/postInstall	String	<p>Post-installation scripts. The total number of characters after transcoding cannot exceed 10240.</p> <p><b>NOTE</b> The input value must be encoded using Base64. The method is as follows:</p> <pre>echo -n "Content to be encoded"   base64</pre>
alpha.cce/NodeImageID	String	This parameter is required when a custom image is used to create a BMS node.

Parameter	Type	Description
chargingMode	Integer	Billing mode of a node. This parameter has been deprecated. Use the <b>billingMode</b> parameter in <b>NodeSpec</b> .
agency_name	String	Name of an agency An agency is created by a tenant administrator on Identity and Access Management (IAM) to provide temporary credentials for CCE nodes to access cloud servers. This parameter is returned only when it is transferred during node creation.
kubeReservedMem	Integer	Reserved node memory, which is reserved for Kubernetes components.
systemReservedMem	Integer	Reserved node memory, which is reserved for system components.
init-node-password	String	Node password. If this parameter is used as a response parameter, asterisks (*) will be displayed.

**Table 4-595** HostnameConfig

Parameter	Type	Description
type	String	Configuration type of the Kubernetes node name. The default value is <b>privatelp</b> . <ul style="list-style-type: none"> <li>• <b>privatelp</b>: The Kubernetes node is named after its IP address.</li> <li>• <b>cceNodeName</b>: The Kubernetes node is named after the CCE node.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• For a node which is configured using <b>cceNodeName</b>, the name is the same as the Kubernetes node name and the ECS name. The node name cannot be changed. If the ECS name is changed on the ECS console, the node name will retain unchanged after ECS synchronization.</li> <li>• For a node which is configured using <b>cceNodeName</b>, to avoid a conflict between Kubernetes nodes, the system automatically adds a suffix to each node name. The suffix is in the format of <b>A hyphen (-) Five random characters</b>. The value of the random characters is a lowercase letter or a digit ranging from 0 to 9.</li> </ul>

**Table 4-596** NodePoolNodeAutoscaling

Parameter	Type	Description
enable	Boolean	Whether to enable auto scaling.
minNodeCount	Integer	Minimum number of nodes allowed if auto scaling is enabled. The value cannot be greater than the maximum number of nodes allowed by the cluster specifications.
maxNodeCount	Integer	Maximum number of nodes allowed if auto scaling is enabled. This value must be greater than or equal to the value of <b>minNodeCount</b> and cannot exceed the maximum number of nodes in the cluster specifications.
scaleDownCooldownTime	Integer	Interval between two scaling operations, in minutes. During this period, nodes added after a scale-up will not be deleted.
priority	Integer	Weight of a node pool. A node pool with a higher weight has a higher priority during scaling.

**Table 4-597** NodeManagement

Parameter	Type	Description
serverGroupReference	String	Cloud server group ID. If this field is specified, all nodes in the node pool will be created in this group. The group ID can be specified only when you create the node pool and cannot be modified. When you specify a cloud server group, the number of nodes in the node pool cannot exceed the group quota.

**Table 4-598** SecurityID

Parameter	Type	Description
id	String	Security group ID

**Table 4-599** ExtensionScaleGroup

Parameter	Type	Description
metadata	<a href="#">ExtensionScaleGroupMetadata</a> object	Basic information about an extended scaling group

Parameter	Type	Description
spec	<a href="#">ExtensionScaleGroupSpec</a> object	Configurations of an extended scaling group, which carry different configurations from those of the default scaling group

**Table 4-600** ExtensionScaleGroupMetadata

Parameter	Type	Description
uid	String	UUID of an extended scaling group, which is automatically generated by the system.
name	String	Name of an extended scaling group. The value cannot be <b>default</b> and can contain a maximum of 55 characters. Only digits, lowercase letters, and hyphens (-) are allowed.

**Table 4-601** ExtensionScaleGroupSpec

Parameter	Type	Description
flavor	String	Node flavor
az	String	AZ of a node. If this parameter is not specified or left blank, the default scaling group configurations take effect.
capacityReservationSpecification	<a href="#">CapacityReservationSpecification</a> object	Capacity reservation configurations of an extended scaling group
autoscaling	<a href="#">ScaleGroupAutoscaling</a> object	Auto scaling configurations of an extended scaling group

**Table 4-602** CapacityReservationSpecification

Parameter	Type	Description
id	String	Private pool ID. The parameter value can be ignored when <b>preference</b> is set to <b>none</b> .
preference	String	Capacity of a private storage pool. If the value is <b>none</b> , the capacity reservation is not specified. If the value is <b>targeted</b> , the capacity reservation is specified. In this case, the ID cannot be left blank.



**Table 4-603** ScaleGroupAutoscaling

Parameter	Type	Description
enable	Boolean	Whether to enable auto scaling for a scaling group. This function is disabled by default.
extensionPriority	Integer	Priority of a scaling group. If no value is specified for this parameter, it defaults to <b>0</b> . A higher value indicates a greater priority.
minNodeCount	Integer	Minimum number of nodes in a scaling group during auto scaling. The value must be greater than <b>0</b> .
maxNodeCount	Integer	Maximum number of nodes that can be retained in a scaling group during auto scaling. The value must be greater than or equal to that of <b>minNodeCount</b> , and can neither be greater than the maximum number of nodes allowed by the cluster nor the maximum number of nodes in the node pool.

**Table 4-604** NodePoolStatus

Parameter	Type	Description
currentNode	Integer	Total number of nodes in the current node pool (excluding the nodes that are being deleted)
creatingNode	Integer	Number of nodes in the creation process in the node pool
deletingNode	Integer	Number of nodes being deleted in the current node pool.

Parameter	Type	Description
phase	String	<p>Node pool status.</p> <ul style="list-style-type: none"> <li>• Null: available (the number of current nodes in the node pool has reached the expected value, and no node scaling is being performed.)</li> <li>• Synchronizing: scaling in progress (the number of current nodes in the node pool does not reach the expected value and no node scaling is being performed.)</li> <li>• Synchronized: pending scaling (the number of current nodes in the node pool does not reach the expected value, or node scaling is being performed.)</li> <li>• SoldOut: The node pool cannot be scaled out. (This field is used in multiple scenarios, for example, resources in the node pool have been sold out or the resource quota is insufficient.)</li> </ul> <p><b>NOTE</b> This node pool status has been discarded and is reserved only for compatibility. You are not advised to use it. The replacements are as follows:</p> <ul style="list-style-type: none"> <li>• Node pool scaling status: You can obtain accurate status of the current node pool based on parameters such as <b>currentNode</b>, <b>creatingNode</b>, or <b>deletingNode</b>.</li> <li>• Node pool scale-out: You can use <b>conditions</b> to obtain the detailed status of a node pool. <b>Scalable</b> can replace <b>SoldOut</b>.</li> <li>• <b>Deleting</b>: The object is being deleted.</li> <li>• <b>Error</b>: An error occurs.</li> </ul>
jobId	String	ID of a job executed on the node pool. This field is available only when the node pool is being deleted.
conditions	Array of <b>NodePoolCondition</b> objects	Node pool status details. For details, see the definition of <b>Condition</b> .
scaleGroupStatuses	Array of <b>ScaleGroupStatus</b> objects	Detailed status of a scaling group. For details, see the definition of <b>ScaleGroupStatus</b> .

**Table 4-605** ScaleGroupStatus

Parameter	Type	Description
name	String	Scaling group name
uid	String	UUID of a scaling group
creationTimes tamp	String	The time when a scaling group was created
updateTimest amp	String	The time when a scaling group was updated
phase	String	<p>Status of a scaling group.</p> <ul style="list-style-type: none"> <li>• <b>Null:</b> The scaling group is available. (The number of nodes in the current scaling group has reached the expected value, and no node scaling is being performed.)</li> <li>• <b>Synchronizing:</b> The scaling group is being scaled. (The number of nodes in the current scaling group has not reached the expected value, and no node scaling is being performed.)</li> <li>• <b>Synchronized:</b> The scaling is pending. (The number of nodes in the current scaling group has not reached the expected value, or node scaling is being performed.)</li> <li>• <b>SoldOut:</b> The scaling group cannot be scaled out. (This field is used in multiple scenarios, for example, resources in the node pool have been sold out or the resource quota is insufficient.)</li> </ul> <p><b>NOTE</b> This scaling group status has been deprecated. It is reserved only for compatibility. Do not use it anymore. Use the following instead:</p> <ul style="list-style-type: none"> <li>• Scaling status of a scaling group: You can obtain the status of the current scaling group using parameters such as <b>desiredNodeCount</b>, <b>existingNodeCount</b>, or <b>upcomingNodeCount</b>.</li> <li>• Scaling group scale-out: You can use <b>conditions</b> to obtain the detailed status of a scaling group. <b>Scalable</b> can replace <b>SoldOut</b>.</li> <li>• <b>Deleting:</b> The scaling group being deleted.</li> <li>• <b>Error:</b> An error occurred in the scaling group.</li> </ul>
desiredNodeC ount	Integer	Expected number of nodes in a scaling group
unpaidScaleN odeCount	Integer	Number of unpaid nodes in an order

Parameter	Type	Description
existingNodeCount	<a href="#">existingNodeCount</a> object	Statistics on existing nodes in a scaling group
upcomingNodeCount	<a href="#">upcomingNodeCount</a> object	Statistics on the nodes to be created in a scaling group
scaleDownDisabledNodeCount	Integer	Number of nodes that cannot be deleted in a scaling group
conditions	Array of <a href="#">NodePoolCondition</a> objects	Scaling group status details. For details, see the definition of <b>Condition</b> .

**Table 4-606** existingNodeCount

Parameter	Type	Description
postPaid	Integer	Number of pay-per-use nodes
total	Integer	Total number of pay-per-use nodes

**Table 4-607** upcomingNodeCount

Parameter	Type	Description
postPaid	Integer	Number of pay-per-use nodes
total	Integer	Total number of pay-per-use nodes

**Table 4-608** NodePoolCondition

Parameter	Type	Description
type	String	Condition type. The options are as follows: <ul style="list-style-type: none"> <li>● <b>Scalable</b>: whether a node pool can be scaled. If the status is <b>False</b>, node pool scaling will not be triggered again.</li> <li>● <b>QuotaInsufficient</b>: Quotas on which node pool scaling depends are insufficient, affecting the node pool scaling status.</li> <li>● <b>ResourceInsufficient</b>: Resources on which node pool scaling depends are insufficient, affecting the node pool scaling status.</li> <li>● <b>UnexpectedError</b>: The node pool fails to be scaled out due to unexpected reasons, affecting the node pool scaling status.</li> <li>● <b>Error</b>: A node pool error occurs. A common trigger is deletion failure.</li> </ul>
status	String	Current status of <b>Condition</b> . The options are as follows: <ul style="list-style-type: none"> <li>● "True"</li> <li>● "False"</li> </ul>
lastProbeTime	String	Time when the status was last checked
lastTransitTime	String	Time when the status was last changed
reason	String	Reason why the status was last changed
message	String	Detailed condition description

## Example Requests

None

## Example Responses

**Status code: 200**

Information about all node pools in the cluster is successfully obtained.

```
{
  "kind" : "List",
  "apiVersion" : "v3",
  "items" : [ {
    "kind" : "NodePool",
    "apiVersion" : "v3",
    "metadata" : {
      "name" : "az1.dc1#s1.large#EulerOS 2.2",
      "uid" : "az1.dc1#s1.large#EulerOS 2.2"
    }
  },
}
```

```

"spec" : {
  "nodeTemplate" : {
    "flavor" : "s1.large",
    "az" : "az1.dc1",
    "os" : "EulerOS 2.2",
    "login" : {
      "sshKey" : "KeyPair-001"
    },
    "rootVolume" : { },
    "publicIP" : {
      "eip" : {
        "bandwidth" : { }
      }
    },
    "billingMode" : 0
  },
  "autoscaling" : {
    "enable" : true,
    "maxNodeCount" : 50
  }
},
"status" : {
  "currentNode" : 1
}
} ]
}

```

## Status Codes

Status Code	Description
200	Information about all node pools in the cluster is successfully obtained.

## Error Codes

See [Error Codes](#).

### 4.4.4 Updating a Specified Node Pool

#### Function

This API is used to update information about a specified node pool. This API can be called only when the cluster is in the Available, Scale Out, or Scale In state.

#### NOTE

- The URL for cluster management is in the following format: `https://Endpoint/uri`. `uri` indicates the resource path, that is, the API access path.
- Currently, only the node pool name can be updated. `initialNodeCount`, `k8sTags`, `taints`, `login` and `userTags` under `spec` are related to the scaling configuration of the node pool. If no value is set for this update, the initial value is used by default.

#### URI

PUT `/api/v3/projects/{project_id}/clusters/{cluster_id}/nodepools/{nodepool_id}`

**Table 4-609** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
nodepool_id	Yes	String	Node pool ID.

## Request Parameters

**Table 4-610** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

**Table 4-611** Request body parameters

Parameter	Mandatory	Type	Description
metadata	Yes	<a href="#">NodePoolMetadataUpdate</a> object	Metadata information of the node pool.
spec	Yes	<a href="#">NodePoolSpecUpdate</a> object	Node pool specifications. <b>NOTE</b> <ul style="list-style-type: none"> <li>This field is optional for updating a node pool.</li> </ul>

**Table 4-612** NodePoolMetadataUpdate

Parameter	Mandatory	Type	Description
name	Yes	String	Name of the node pool. <b>NOTE</b> Naming rules: <ul style="list-style-type: none"> <li>• Enter 1 to 50 characters, starting with a lowercase letter and not ending with a hyphen (-). Only lowercase letters, digits, and hyphens (-) are allowed.</li> <li>• You cannot create node pools named DefaultPool.</li> </ul>

**Table 4-613** NodePoolSpecUpdate

Parameter	Mandatory	Type	Description
nodeTemplate	Yes	<a href="#">NodeSpecUpdate</a> object	Detailed parameters of the node pool template.
initialNodeCount	Yes	Integer	Initial number of nodes for the node pool. When queried, the value is the number of target nodes in the node pool. The default value is <b>0</b> .
autoscaling	Yes	<a href="#">NodePoolNodeAutoscaling</a> object	Auto scaling parameters. Only pay-per-use node pools support auto scaling.
taintPolicyOnExistingNodes	No	String	Whether to synchronize Kubernetes taints. <ul style="list-style-type: none"> <li>• refresh: Kubernetes taint modifications will be automatically synchronized to existing nodes.</li> <li>• ignore: Kubernetes taint modifications will not be automatically synchronized to existing nodes.</li> </ul>



Parameter	Mandatory	Type	Description
labelPolicyOnExistingNodes	No	String	Whether to synchronize Kubernetes labels. <ul style="list-style-type: none"> <li>refresh: Kubernetes label modifications will be automatically synchronized to existing nodes.</li> <li>ignore: Kubernetes label modifications will not be automatically synchronized to existing nodes.</li> </ul>
userTagsPolicyOnExistingNodes	No	String	Whether to synchronize node pool tags. <ul style="list-style-type: none"> <li>refresh: Node pool tag modifications will be automatically synchronized to existing nodes.</li> <li>ignore: Node pool tag modifications will not be automatically synchronized to existing nodes.</li> </ul>
extensionScaleGroups	No	Array of <a href="#">ExtensionScaleGroup</a> objects	Configuration list of the extended scaling groups in a node pool. For details, see the definition of <a href="#">ExtensionScaleGroup</a> .

**Table 4-614** NodeSpecUpdate

Parameter	Mandatory	Type	Description
flavor	No	String	Node specifications. For details about the node specifications supported by CCE, see <a href="#">Node Flavor Description</a> .

Parameter	Mandatory	Type	Description
os	No	String	<p>The OS type of the node. For details about the supported OSs, see <a href="#">Node OS</a>.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• If the selected OS is not supported by the current cluster version, it will be automatically replaced with a supported OS of the same series.</li> <li>• If alpha.cce/NodeImageID in extendParam is specified during node creation, you do not need to configure this parameter.</li> <li>• If this parameter is not specified, CCE will automatically select a supported OS version based on the cluster version.</li> <li>• This parameter is mandatory when creating a node pool.</li> <li>• This parameter is mandatory if the disk initialization parameter storage is used and the name field of virtualSpaces in StorageGroups is set to share during node creation.</li> </ul>
login	No	<a href="#">Login</a> object	<p>Node login mode. Either the key pair or password must be used for login.</p> <p><b>NOTE</b></p> <p>This parameter cannot be modified when a node pool is being updated.</p>
rootVolumeUpdate	No	<a href="#">Volume</a> object	Disk information of a node.

Parameter	Mandatory	Type	Description
dataVolumes Update	No	Array of <b>Volume</b> objects	<p>Data disk parameters of a node. For DeC nodes, the parameter details are the same as those in rootVolume.</p> <p>A VM can have up to 16 disks attached, while a BMS can only have up to 10 disks attached. The maximum number of disks that can be attached to either a VM or BMS is limited by their respective maximum disk limits. (You can add more data disks to each of your nodes through the CCE console or APIs.)</p> <p>The data disk used by the container runtime and kubelet cannot be detached, or the node will become unavailable.</p>
storage	No	<b>Storage</b> object	<p>Disk initialization management parameter.</p> <p>This parameter is complex to configure. For details, see <a href="#">Attaching Disks to a Node</a>.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• If a node has both local and EVS disks attached, make sure to not keep the default value for this parameter, or it may result in unexpected disk partitions.</li> <li>• If you want to change the value range of a data disk to 20 to 32768, do not use the default value.</li> <li>• If you want to use the shared disk space (with the runtime and Kubernetes partitions cancelled), do not retain the default value of this parameter. For details about the shared disk space, see <a href="#">Data Disk Space Allocation</a>.</li> </ul>

Parameter	Mandatory	Type	Description
runtime	No	<b>Runtime</b> object	<p>Container runtime</p> <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is docker.</li> <li>• Clusters of v1.25 or later: The default value varies with the OSs.</li> <li>• For nodes running EulerOS 2.5 or EulerOS 2.8, the default container runtime is docker. For nodes running other OSs, the default container runtime is containerd.</li> </ul>

Parameter	Mandatory	Type	Description
taints	Yes	Array of <b>Taint</b> objects	<p>You can add taints to created nodes to configure anti-affinity. A maximum of 20 taints can be added. It is left blank by default. Each taint contains the following parameters:</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain can be prefixed to a key.</li> <li>• <b>Value:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed.</li> <li>• <b>Effect:</b> Available options are <b>NoSchedule</b>, <b>PreferNoSchedule</b>, and <b>NoExecute</b>.</li> </ul> <p>Example:</p> <pre>"taints": [{   "key": "status",   "value": "unavailable",   "effect": "NoSchedule" }, {   "key": "looks",   "value": "bad",   "effect": "NoSchedule" }]</pre> <p><b>NOTE</b> If this parameter is not specified or left empty, the custom taints of the node pool will be deleted.</p>

Parameter	Mandatory	Type	Description
k8sTags	Yes	Map<String,String>	<p>Defined in key-value pairs. A maximum of 20 key-value pairs are allowed. It is left blank by default.</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain containing a maximum of 253 characters (such as <b>example.com/my-key</b>) can be prefixed to a key.</li> <li>• <b>Value:</b> The value can be left blank or contain 1 to 63 characters that start and end with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed.</li> </ul> <p>Example:</p> <pre>"k8sTags": {   "key": "value" }</pre> <p><b>NOTE</b> If this parameter is not specified or left empty, the custom Kubernetes labels of the node pool will be deleted.</p>
ecsGroupId	No	String	<p>ECS group ID. If this parameter is specified, nodes will be created in the specific ECS group.</p> <p><b>NOTE</b> This configuration does not take effect when a node pool is created. To ensure that all nodes in a node pool are in the same ECS group, configure nodeManagement in the node pool.</p>

Parameter	Mandatory	Type	Description
userTags	Yes	Array of <b>UserTag</b> objects	<p>Cloud server tag (resource tag). The key of a tag must be unique. The maximum number of custom tags supported by CCE varies depending on regions and cannot exceed 8. It is left blank by default.</p> <p><b>NOTE</b> If this parameter is not specified or left empty, the custom cloud server tags of the node pool will be deleted.</p>
nodeNameTemplate	No	Object	<p>Defined in key-value pairs. Default value: null.</p> <ul style="list-style-type: none"> <li>• nodeNamePrefix: A prefix must start with a lowercase letter and can only contain lowercase letters, digits, periods (.), and hyphens (-). If a period (.) is used, it must be surrounded by lowercase letters or digits.</li> <li>• nodeNameSuffix: A suffix must end with a lowercase letter or digit and can only contain lowercase letters, digits, periods (.), and hyphens (-). If a period (.) is used, it must be surrounded by lowercase letters or digits.</li> </ul> <p>Example:  <pre>"k8sTags": {   "nodeNamePrefix": "test1",   "nodeNameSuffix": "test2" }</pre> </p> <p><b>NOTE</b> This field is optional for updating a node pool.</p>

Parameter	Mandatory	Type	Description
initializedConditions	No	Array of strings	<p>Custom initialization flag, which is left blank by default. Before CCE nodes are initialized, <b>node.cloudprovider.kubernetes.io/uninitialized</b> is added to the nodes to prevent pods from being scheduled to these nodes. When creating a node, you can configure the <b>initializedConditions</b> parameter to control the taint removal time. By default, the timeout period is not configured. Example:</p> <ol style="list-style-type: none"> <li>1. Create a node and add parameter "initializedConditions": ["CCEInitial", "CustomedInitial"].</li> <li>2. After custom initialization is complete, CCE will call a Kubernetes API (for example, <b>PATCH /v1/nodes/{node_ip}/status</b>) to update the node <b>conditions</b> by adding two labels of types <b>CCEInitial</b> and <b>CustomedInitial</b> and setting the status to <b>True</b>, as shown in the following:</li> </ol> <pre>status:   conditions:   - type: CCEInitial     status: 'True'   - type: CustomedInitial     status: 'True'</pre> <ol style="list-style-type: none"> <li>1. CCE polls <b>status.Conditions</b> of nodes to check whether there are conditions of types <b>CCEInitial</b> and <b>CustomedInitial</b>. If such conditions exist and the status is <b>True</b>, the node initialization is complete and the initialization taint is removed.</li> <li>2. <b>initializedConditions</b> allows you to configure a</li> </ol>



Parameter	Mandatory	Type	Description
			<p>timeout period during node creation. For example, "initializedConditions": ["CCEInitial:15m", "CustomedInitial:15m"] indicates that the timeout period is 15 minutes. After the timeout period is reached, the initialization conditions will be automatically ignored and the initialization taint will be removed when CCE polls the node.</p> <ul style="list-style-type: none"> <li>• The timeout interval in <b>initializedConditions</b> ranges from 1 to 99 seconds.</li> <li>• Use only letters and digits. Max. characters: 20.</li> <li>• Max. flags: 2.</li> <li>• The unit of the timeout period is minute (m).</li> </ul>
serverEnterpriseProjectID	No	String	<p>Enterprise project ID of a server. CCE does not support features provided by EPS. This field is used only for synchronizing enterprise project IDs of servers. To create a node or node pool, you can specify an existing enterprise project for it. If this field is left blank, it inherits the enterprise project attributes of the cluster that runs the node or node pool. When updating a node pool, any configuration modifications you made will apply only to new nodes. You need to modify the enterprise projects of existing nodes on the EPS console. If no value is specified during the update, this field will not be updated. If this parameter is left blank, the cluster enterprise project is returned.</p>

Parameter	Mandatory	Type	Description
nodeNicSpec Update	No	<a href="#">nodeNicSpec Update</a> object	Update the NIC of a node.
extendParam	No	<a href="#">NodePoolUpdateExtendParam</a> object	Extended parameters supported when a node pool is updated.

**Table 4-615** Login

Parameter	Mandatory	Type	Description
sshKey	No	String	Name of the key pair used for login.
userPassword	No	<a href="#">UserPassword</a> object	Password used for node login.

**Table 4-616** UserPassword

Parameter	Mandatory	Type	Description
username	No	String	Login account. The default value is <b>root</b> .

Parameter	Mandatory	Type	Description
password	Yes	String	<p>Login password. If a username and password are used to create a node, this field is shielded in the response body. A password must meet the following complexity requirements:</p> <ul style="list-style-type: none"> <li>• Contains 8 to 26 characters.</li> <li>• Contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters !@\$%^_+=+[{ }];,./?</li> <li>• Cannot contain the username or the username spelled backwards. The <b>password</b> field must be salted during node creation. For details, see <a href="#">Adding a Salt in the password Field When Creating a Node</a>.</li> </ul>

**Table 4-617** Volume

Parameter	Mandatory	Type	Description
size	Yes	Integer	<p>Disk size, in GiB.</p> <ul style="list-style-type: none"> <li>• System disk: 40 to 1024</li> <li>• Value range for data disks: 100 to 32768</li> </ul>
volumetype	Yes	String	<p>Disk type. For details about possible values, see the description of the <b>root_volume</b> parameter in the API used to create an ECS.</p> <ul style="list-style-type: none"> <li>• <b>SAS</b>: high I/O SAS disk</li> <li>• <b>SSD</b>: ultra-high I/O SSD disk</li> <li>• <b>SATA</b>: common I/O SATA disk SATA disks have been removed from EVS. You can find them attached only on existing nodes.</li> </ul>

Parameter	Mandatory	Type	Description
extendParam	No	Map<String, Object>	Extended disk parameters, defined in <b>extendparam</b> in the API used to create an ECS.
cluster_id	No	String	ID of the storage pool used by the ECS system disk. This field is used only for DeC clusters, which functions as <b>dssPoolID</b> , that is, the ID of the DSS storage pool.
cluster_type	No	String	Storage class of the cloud server system disk. The value is always <b>dss</b> . This field is used only for DeC clusters.
hw:passthrough	No	Boolean	<ul style="list-style-type: none"> <li>Pay attention to this field if your ECS is SDI-compliant. If the value of this field is <b>true</b>, the created disk is of SCSI type.</li> <li>If the node pool type is <b>ElasticBMS</b>, this field must be set to <b>true</b>.</li> <li>If a node specification involves both local and EVS disks, configure the disk initialization parameters. For details, see <a href="#">Attaching Disks to a Node</a>.</li> </ul>
metadata	No	<b>VolumeMetadata</b> object	EVS disk encryption information. This field is mandatory only when you need to encrypt the system disk or data disks of the node to be created.

**Table 4-618** VolumeMetadata

Parameter	Mandatory	Type	Description
<code>__system__encrypted</code>	No	String	Whether the EVS disk is encrypted. The value <b>0</b> indicates that the EVS disk is not encrypted, and the value <b>1</b> indicates that the EVS disk is encrypted.  If this parameter is not specified, EVS disks will not be encrypted by default.
<code>__system__cmkid</code>	No	String	CMK ID, which indicates encryption in <b>metadata</b> . This field is used with <b>__system__encrypted</b> .

**Table 4-619** Storage

Parameter	Mandatory	Type	Description
<code>storageSelectors</code>	Yes	Array of <b>StorageSelectors</b> objects	Disk selection. Matched disks are managed according to <b>matchLabels</b> and <b>storageType</b> .
<code>storageGroups</code>	Yes	Array of <b>StorageGroups</b> objects	A storage group consists of multiple storage devices. It is used to divide storage space.

**Table 4-620** StorageSelectors

Parameter	Mandatory	Type	Description
<code>name</code>	Yes	String	Selector name, used as the index of <b>selectorNames</b> in <b>storageGroup</b> . Therefore, the name of each selector must be unique.
<code>storageType</code>	Yes	String	Storage type. Only <b>evs</b> (EVS disks) and <b>local</b> (local disks) are supported. The local storage does not support disk selection. All local disks will form a VG, so only one <code>storageSelector</code> of the local type is allowed.

Parameter	Mandatory	Type	Description
matchLabels	No	<a href="#">matchLabels</a> object	Matching field of an EVS disk. The value can be <b>size</b> , <b>volumeType</b> , <b>metadataEncrypted</b> , <b>metadataCmkid</b> , or <b>count</b> in <b>DataVolume</b> .

**Table 4-621** matchLabels

Parameter	Mandatory	Type	Description
size	No	String	Matched disk size. If this parameter is left unspecified, the disk size is not limited. Example: 100
volumeType	No	String	EVS disk type. <b>SSD</b> , <b>GPSSD</b> , <b>SAS</b> , <b>ESSD</b> , and <b>SATA</b> are supported. If this parameter is left blank, there is no restriction on the disk type.
metadataEncrypted	No	String	Disk encryption identifier. If this parameter is set to <b>0</b> , it specifies that the disk is not encrypted. If this parameter is set to <b>1</b> , it specifies that the disk is encrypted. If this parameter is left blank, there is no restriction on the disk encryption identifier.
metadataCmkid	No	String	CMK ID of an encrypted disk. The value is a string of 36 bytes. If this parameter is left blank, there is no restriction on the disk key ID.
count	No	String	Number of disks to be selected. If this parameter is left blank, all disks of this type are selected.

**Table 4-622** StorageGroups

Parameter	Mandatory	Type	Description
name	Yes	String	Name of a virtual storage group, which must be unique. <b>NOTE</b> <ul style="list-style-type: none"> <li>If <b>cceManaged</b> is set to <b>true</b>, the name must be set to <b>vgpaas</b>.</li> <li>If the data disk is used as a temporary storage volume, the name must be <b>vg-everest-localvolume-ephemeral</b>.</li> <li>If the data disk is used as a persistent storage volume, the name must be <b>vg-everest-localvolume-persistent</b>.</li> </ul>
cceManaged	No	Boolean	Storage space for Kubernetes and runtime components. Only one group can be set to <b>true</b> . If this parameter is left blank, the default value <b>false</b> is used.
selectorNames	Yes	Array of strings	Corresponds to <b>name</b> in <b>storageSelectors</b> . One group can select multiple selectors, but one selector can be selected by only one group.
virtualSpaces	Yes	Array of <b>VirtualSpace</b> objects	Detailed management of space configuration in a group.

**Table 4-623** VirtualSpace

Parameter	Mandatory	Type	Description
name	Yes	String	<p>Name of a virtualSpace. Options: <b>share</b>, <b>kubernetes</b>, <b>runtime</b>, and <b>user</b></p> <ul style="list-style-type: none"> <li>• <b>share</b>: shared space configuration (with the runtime and Kubernetes partitions cancelled). <b>lvmConfig</b> needs to be configured.</li> <li>• <b>Kubernetes</b>: Kubernetes space configuration. <b>lvmConfig</b> needs to be configured.</li> <li>• <b>runtime</b>: runtime space configuration. <b>runtimeConfig</b> needs to be configured.</li> <li>• <b>user</b>: user space configuration. <b>lvmConfig</b> needs to be configured.</li> </ul>
size	Yes	String	<p>Size of a virtualSpace. The value must be an integer in percentage. Example: 90%.</p> <p><b>NOTE</b> The value ranges from 10 to 90. The sum of percentages of all virtualSpaces in a group cannot exceed 100%.</p>
lvmConfig	No	<b>LVMConfig</b> object	<p>LVM configuration managements, applicable to the <i>share</i>*, <b>kubernetes</b>, and <b>user</b> spaces. Note that one virtual space supports only one config.</p>
runtimeConfig	No	<b>RuntimeConfig</b> object	<p>runtime configurations, applicable to the <b>runtime</b> space. Note that one virtual space supports only one config.</p>



**Table 4-624** LVMConfig

Parameter	Mandatory	Type	Description
lvType	Yes	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.
path	No	String	Disk mount path. This parameter takes effect only in user configuration. The value is an absolute path. Digits, letters, periods (.), hyphens (-), and underscores (_) are allowed.  <b>NOTE</b> Do not leave the path empty or set it to a key OS path such as the root directory. Key OS paths that cannot be specified: /, /home, /home/, /bin, /bin/, /lib, /lib/, /root, /root/, /boot, /boot/, /dev, /dev/, /etc, /etc/, /lost+found, /lost+found/, /mnt, /mnt/, /proc, /proc/, /sbin, /sbin/, /srv, /srv/, /tmp, /tmp/, /var, /var/, /media, /media/, /opt, /opt/, /selinux, /selinux/, /sys, /sys/, /usr, /usr/, /opt/cloud/, /mnt/paas/, /home/paas/, /var/paas/, /var/lib/, /var/script/

**Table 4-625** RuntimeConfig

Parameter	Mandatory	Type	Description
lvType	Yes	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.

**Table 4-626** Runtime

Parameter	Mandatory	Type	Description
name	No	String	Container runtime. Default value: <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>

**Table 4-627** Taint

Parameter	Mandatory	Type	Description
key	Yes	String	Key.
value	No	String	Value.
effect	Yes	String	Effect.

**Table 4-628** UserTag

Parameter	Mandatory	Type	Description
key	No	String	Key of the cloud server label. The value cannot start with <b>CCE-</b> or <b>__type_baremetal</b> .
value	No	String	Value of the cloud server label.

**Table 4-629** nodeNicSpecUpdate

Parameter	Mandatory	Type	Description
primaryNic	No	<b>primaryNic</b> object	Description of the primary NIC

**Table 4-630** primaryNic

Parameter	Mandatory	Type	Description
subnetId	No	String	Network ID of the subnet to which a NIC belongs. If <b>subnetId</b> is not specified when a primary NIC is creating, the cluster subnet will be used. If <b>subnetList</b> is also configured for a node pool, the <b>subnetList</b> field is used for adding subnets to the node pool. When creating an extension NIC, you must specify <b>subnetId</b> .
subnetList	No	Array of strings	Network ID list of the subnet where the NIC resides. A maximum of 20 subnets can be configured for a node pool.

**Table 4-631** NodePoolUpdateExtendParam

Parameter	Mandatory	Type	Description
agency_name	No	String	Agency name. An agency is created by a tenant administrator on Identity and Access Management (IAM) to provide temporary credentials for CCE nodes to access cloud servers. This parameter is returned only when it is transferred during node creation.
alpha.cce/ preInstall	No	String	Pre-installation script. The input value must be encoded using Base64. The method is as follows: echo -n <b>"*Content to be encoded*"   base64</b>
alpha.cce/ postInstall	No	String	Post-installation script. The input value must be encoded using Base64. The method is as follows: echo -n <b>"*Content to be encoded*"   base64</b>

**Table 4-632** NodePoolNodeAutoscaling

Parameter	Mandatory	Type	Description
enable	No	Boolean	Whether to enable auto scaling.
minNodeCount	No	Integer	Minimum number of nodes allowed if auto scaling is enabled. The value cannot be greater than the maximum number of nodes allowed by the cluster specifications.
maxNodeCount	No	Integer	Maximum number of nodes allowed if auto scaling is enabled. This value must be greater than or equal to the value of <b>minNodeCount</b> and cannot exceed the maximum number of nodes in the cluster specifications.
scaleDownCooldownTime	No	Integer	Interval between two scaling operations, in minutes. During this period, nodes added after a scale-up will not be deleted.
priority	No	Integer	Weight of a node pool. A node pool with a higher weight has a higher priority during scaling.

**Table 4-633** ExtensionScaleGroup

Parameter	Mandatory	Type	Description
metadata	No	<a href="#">ExtensionScaleGroupMetadata</a> object	Basic information about an extended scaling group
spec	No	<a href="#">ExtensionScaleGroupSpec</a> object	Configurations of an extended scaling group, which carry different configurations from those of the default scaling group

**Table 4-634** ExtensionScaleGroupMetadata

Parameter	Mandatory	Type	Description
uid	No	String	UUID of an extended scaling group, which is automatically generated by the system.
name	No	String	Name of an extended scaling group. The value cannot be <b>default</b> and can contain a maximum of 55 characters. Only digits, lowercase letters, and hyphens (-) are allowed.

**Table 4-635** ExtensionScaleGroupSpec

Parameter	Mandatory	Type	Description
flavor	No	String	Node flavor
az	No	String	AZ of a node. If this parameter is not specified or left blank, the default scaling group configurations take effect.
capacityReservationSpecification	No	<a href="#">CapacityReservationSpecification</a> object	Capacity reservation configurations of an extended scaling group
autoscaling	No	<a href="#">ScaleGroupAutoscaling</a> object	Auto scaling configurations of an extended scaling group

**Table 4-636** CapacityReservationSpecification

Parameter	Mandatory	Type	Description
id	No	String	Private pool ID. The parameter value can be ignored when <b>preference</b> is set to <b>none</b> .
preference	No	String	Capacity of a private storage pool. If the value is <b>none</b> , the capacity reservation is not specified. If the value is <b>targeted</b> , the capacity reservation is specified. In this case, the ID cannot be left blank.

**Table 4-637** ScaleGroupAutoscaling

Parameter	Mandatory	Type	Description
enable	No	Boolean	Whether to enable auto scaling for a scaling group. This function is disabled by default.
extensionPriority	No	Integer	Priority of a scaling group. If no value is specified for this parameter, it defaults to <b>0</b> . A higher value indicates a greater priority.
minNodeCount	No	Integer	Minimum number of nodes in a scaling group during auto scaling. The value must be greater than <b>0</b> .
maxNodeCount	No	Integer	Maximum number of nodes that can be retained in a scaling group during auto scaling. The value must be greater than or equal to that of <b>minNodeCount</b> , and can neither be greater than the maximum number of nodes allowed by the cluster nor the maximum number of nodes in the node pool.

## Response Parameters

Status code: 200

**Table 4-638** Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at <b>NodePool</b> .
apiVersion	String	API version. The value is fixed at <b>v3</b> .
metadata	<b>NodePoolMetadata</b> object	Metadata of the node pool.
spec	<b>NodePoolSpec</b> object	Node pool specifications. <b>NOTE</b> <ul style="list-style-type: none"> <li>This field is optional for updating a node pool.</li> </ul>

Parameter	Type	Description
status	<a href="#">UpdateNodePoolStatus</a> object	Node pool status.

**Table 4-639** NodePoolMetadata

Parameter	Type	Description
name	String	Node pool name. <b>NOTE</b> Naming rules: Enter 1 to 50 characters, starting with a lowercase letter and not ending with a hyphen (-). Only lowercase letters, digits, and hyphens (-) are allowed. <ul style="list-style-type: none"> <li>You cannot create node pools named DefaultPool.</li> </ul>
uid	String	UID of the node pool. The value is automatically generated after the object is updated. A user-defined value will not take effect.
annotations	Map<String,String>	Annotations of a node pool in key-value pairs. This parameter is used only for query and cannot be input through a request. Entered data of this parameter is invalid.
updateTimestamp	String	Update time.
creationTimestamp	String	Creation time.

**Table 4-640** NodePoolSpec

Parameter	Type	Description
type	String	Node pool type. If this parameter is left blank, the value <b>vm</b> is used by default. <ul style="list-style-type: none"> <li><b>vm</b>: ECS</li> <li><b>ElasticBMS</b>: C6 general computing-plus BMS. An example flavor is <b>c6.22xlarge.2.physical</b>.</li> <li><b>bms</b>: BMS</li> </ul>
nodeTemplate	<a href="#">NodeSpec</a> object	Detailed parameters of the node pool template.

Parameter	Type	Description
initialNodeCount	Integer	Initial number of nodes for the node pool. When queried, the value is the number of target nodes in the node pool.
autoscaling	<a href="#">NodePoolNodeAutoscaling</a> object	Auto scaling parameters
nodeManagement	<a href="#">NodeManagement</a> object	Node management configuration
podSecurityGroups	Array of <a href="#">SecurityID</a> objects	Security group configuration. This parameter is available only for CCE Turbo clusters.
extensionScaleGroups	Array of <a href="#">ExtensionScaleGroup</a> objects	Configuration list of the extended scaling groups in a node pool. For details, see the definition of <a href="#">ExtensionScaleGroup</a> .
customSecurityGroups	Array of strings	Custom security group settings for a node pool. New nodes scaled out in a node pool can be bound to a specified security group. <ul style="list-style-type: none"> <li>Specifying no security group ID will add the new nodes to the default security group of the worker nodes.</li> <li>Specifying a valid security group ID will put new nodes in that security group.</li> <li>When specifying a security group, do not modify the rules of the port on which CCE running depends.</li> </ul>

**Table 4-641** NodeSpec

Parameter	Type	Description
flavor	String	Node specifications. For details about the node specifications supported by CCE, see <a href="#">Node Flavor Description</a> .
az	String	AZ where the node to be created is located. You need to specify the AZ name. A node created using APIs does not support random AZs. <b>NOTE</b> <ul style="list-style-type: none"> <li>When you are creating a node pool and configuring a scaling group, this parameter cannot be set to <b>random</b>.</li> </ul>



Parameter	Type	Description
os	String	<p>Node OS. For details about the supported OSs, see <a href="#">Node OS</a>.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>The system automatically selects the supported OS based on the cluster version. If the current cluster version does not support the OS, an error will be reported.</li> <li>If <code>alpha.cce/NodeImageID</code> in <code>extendParam</code> is specified during node creation, you do not need to configure this parameter.</li> <li>This parameter is mandatory when creating a node pool.</li> <li>If you are using shared disk space when creating a node, this parameter is mandatory. Specifically, the disk initialization configuration management parameter <code>storage</code> is used and the <code>name</code> field of <code>virtualSpaces</code> in <code>StorageGroups</code> is set to <code>share</code>.</li> </ul>
login	<a href="#">Login</a> object	<p>Node login mode. Either the key pair or password must be used for login.</p> <p><b>NOTE</b> This parameter cannot be modified when a node pool is being updated.</p>
rootVolume	<a href="#">Volume</a> object	Information about disks on the node
dataVolumes	Array of <a href="#">Volume</a> objects	<p>Data disk parameters of a node. A VM can have up to 16 disks attached, while a BMS can only have up to 10 disks attached. The maximum number of disks that can be attached to either a VM or BMS is limited by their respective maximum disk limits. (You can add more data disks to each of your nodes through the CCE console and APIs.) The data disk used by the container runtime and kubelet cannot be detached, or the node will become unavailable. For DeC nodes, the parameter details are the same as those for <b>rootVolume</b>.</p>

Parameter	Type	Description
storage	<a href="#">Storage</a> object	<p>Disk initialization management parameter. This parameter is complex to configure. For details, see <a href="#">Attaching Disks to a Node</a>.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If a node has both local and EVS disks attached, make sure to not keep the default value for this parameter, or it may result in unexpected disk partitions.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If you want to change the value range of a data disk to 20 to 32768, do not use the default value.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If you want to use the shared disk space (with the runtime and Kubernetes partitions cancelled), do not retain the default value of this parameter. For details about the shared disk space, see <a href="#">Data Disk Space Allocation</a>.</li> </ul>
publicIP	<a href="#">NodePublicIP</a> object	<p>EIP of a node.</p> <p><b>NOTE</b></p> <p>This parameter is not supported when you add a node to a node pool.</p>
nodeNicSpec	<a href="#">NodeNicSpec</a> object	NIC of the node
count	Integer	Number of nodes to be created in a batch. The value must be a positive integer greater than or equal to 1 and less than or equal to the defined limit. This parameter can be left blank when it is used for a node pool.
billingMode	Integer	<p>Node billing mode.</p> <ul style="list-style-type: none"> <li><b>0</b>: pay-per-use</li> </ul>

Parameter	Type	Description
taints	Array of <b>Taint</b> objects	<p>You can add taints to created nodes to configure anti-affinity. A maximum of 20 taints can be added. Each taint contains the following parameters:</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain can be prefixed to a key.</li> <li>• <b>Value:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed.</li> <li>• <b>Effect:</b> Available options are <b>NoSchedule</b>, <b>PreferNoSchedule</b>, and <b>NoExecute</b>. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</li> </ul> <p>Example:</p> <pre data-bbox="815 1108 1428 1339">"taints": [{   "key": "status",   "value": "unavailable",   "effect": "NoSchedule" }, {   "key": "looks",   "value": "bad",   "effect": "NoSchedule" }]</pre>

Parameter	Type	Description
k8sTags	Map<String,String>	<p>Defined in key-value pairs. A maximum of 20 key-value pairs are allowed.</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain containing a maximum of 253 characters (such as <b>example.com/my-key</b>) can be prefixed to a key.</li> <li>• <b>Value:</b> The value can be left blank or contain 1 to 63 characters that start and end with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</li> </ul> <p>Example:</p> <pre>"k8sTags": {   "key": "value" }</pre>
ecsGroupId	String	<p>ECS group ID. If this parameter is configured, nodes will be created in the specified ECS group.</p> <p><b>NOTE</b> This configuration does not take effect when a node pool is created. To ensure that all nodes in a node pool are in the same ECS group, configure <b>nodeManagement</b> in the node pool.</p>
dedicatedHostId	String	<p>ID of the DeH to which the node is scheduled.</p> <p><b>NOTE</b> This parameter is not supported when you add a node during node pool creation.</p>

Parameter	Type	Description
userTags	Array of <a href="#">UserTag</a> objects	<p>Cloud server tag (resource tag). The key of a tag must be unique. The maximum number of custom tags supported by CCE varies depending on regions and cannot exceed 8. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</p> <p><b>NOTE</b> A tag key can have a maximum of 36 characters. It can only contain letters, digits, Unicode characters, and special characters (-_).</p>
runtime	<a href="#">Runtime</a> object	<p>Container runtime. Default value:</p> <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>

Parameter	Type	Description
initializedConditions	Array of strings	<p>Custom initialization flag, which is left blank by default.</p> <p>Before CCE nodes are initialized, <b>node.cloudprovider.kubernetes.io/uninitialized</b> is added to the nodes to prevent pods from being scheduled to these nodes. When creating a node, you can configure the <b>initializedConditions</b> parameter to control the taint removal time. By default, the timeout period is not configured.</p> <p>Example:</p> <ol style="list-style-type: none"> <li>1. Create a node and add parameter "initializedConditions": ["CCEInitial", "CustomedInitial"].</li> <li>2. After custom initialization is complete, CCE will call a Kubernetes API (for example, <b>PATCH /v1/nodes/{node_ip}/status</b>) to update the node <b>conditions</b> by adding two labels of types <b>CCEInitial</b> and <b>CustomedInitial</b> and setting the status to <b>True</b>, as shown in the following:</li> </ol> <pre>status:   conditions:   - type: CCEInitial     status: 'True'   - type: CustomedInitial     status: 'True'</pre> <ol style="list-style-type: none"> <li>1. CCE polls <b>status.Conditions</b> of nodes to check whether there are conditions of types <b>CCEInitial</b> and <b>CustomedInitial</b>. If such conditions exist and the status is <b>True</b>, the node initialization is complete and the initialization taint is removed.</li> <li>2. <b>initializedConditions</b> allows you to configure a timeout period during node creation. For example, "initializedConditions": ["CCEInitial:15m", "CustomedInitial:15m"] indicates that the timeout period is 15 minutes. After the timeout period is reached, the initialization conditions will be automatically ignored and the initialization taint will be removed when CCE polls the node. <ul style="list-style-type: none"> <li>• The timeout interval in <b>initializedConditions</b> ranges from 1 to 99 seconds.</li> <li>• Use only letters and digits. Max. characters: 20.</li> <li>• Max. flags: 2.</li> </ul> </li> </ol>

Parameter	Type	Description
		<ul style="list-style-type: none"> <li>The unit of the timeout period is minute (m).</li> </ul>
extendParam	<b>NodeExtendParam</b> object	Extended parameters for creating a node.
hostnameConfig	<b>HostnameConfig</b> object	Kubernetes node name configuration parameter, which is supported by clusters of v1.23.6-r0 to v1.25 or clusters of v1.25.2-r0 or later versions.
serverEnterpriseProjectID	String	Enterprise project ID of a server. CCE does not support features provided by EPS. This field is used only for synchronizing enterprise project IDs of servers. To create a node or node pool, you can specify an existing enterprise project for it. If this field is left blank, it inherits the enterprise project attributes of the cluster that runs the node or node pool. When updating a node pool, any configuration modifications you made will apply only to new nodes. You need to modify the enterprise projects of existing nodes on the EPS console. If no value is specified during the update, this field will not be updated. If this parameter is left blank, the cluster enterprise project is returned.

**Table 4-642** Login

Parameter	Type	Description
sshKey	String	Name of the key pair used for login.
userPassword	<b>UserPassword</b> object	Password used for node login.

**Table 4-643** UserPassword

Parameter	Type	Description
username	String	Login account. The default value is <b>root</b> .

Parameter	Type	Description
password	String	<p>Login password. If a username and password are used to create a node, this field is shielded in the response body. A password must meet the following complexity requirements:</p> <ul style="list-style-type: none"> <li>• Contains 8 to 26 characters.</li> <li>• Contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters !@\$%^-_=+[{ }];,./?</li> <li>• Cannot contain the username or the username spelled backwards. The <b>password</b> field must be salted during node creation. For details, see <a href="#">Adding a Salt in the password Field When Creating a Node</a>.</li> </ul>

Table 4-644 Volume

Parameter	Type	Description
size	Integer	<p>Disk size, in GiB.</p> <ul style="list-style-type: none"> <li>• System disk: 40 to 1024</li> <li>• Value range for data disks: 100 to 32768</li> </ul>
volumetype	String	<p>Disk type. For details about possible values, see the description of the <b>root_volume</b> parameter in the API used to create an ECS.</p> <ul style="list-style-type: none"> <li>• <b>SAS</b>: high I/O SAS disk</li> <li>• <b>SSD</b>: ultra-high I/O SSD disk</li> <li>• <b>SATA</b>: common I/O SATA disk SATA disks have been removed from EVS. You can find them attached only on existing nodes.</li> </ul>
extendParam	Map<String,Object>	<p>Extended disk parameters, defined in <b>extendparam</b> in the API used to create an ECS.</p>
cluster_id	String	<p>ID of the storage pool used by the ECS system disk. This field is used only for DeC clusters, which functions as <b>dssPoolID</b>, that is, the ID of the DSS storage pool.</p>
cluster_type	String	<p>Storage class of the cloud server system disk. The value is always <b>dss</b>. This field is used only for DeC clusters.</p>



Parameter	Type	Description
hw:passthrough	Boolean	<ul style="list-style-type: none"> <li>Pay attention to this field if your ECS is SDI-compliant. If the value of this field is <b>true</b>, the created disk is of SCSI type.</li> <li>If the node pool type is <b>ElasticBMS</b>, this field must be set to <b>true</b>.</li> <li>If a node specification involves both local and EVS disks, configure the disk initialization parameters. For details, see <a href="#">Attaching Disks to a Node</a>.</li> </ul>
metadata	<b>VolumeMetadata</b> object	EVS disk encryption information. This field is mandatory only when you need to encrypt the system disk or data disks of the node to be created.

**Table 4-645** VolumeMetadata

Parameter	Type	Description
__system__encrypted	String	Whether the EVS disk is encrypted. The value <b>0</b> indicates that the EVS disk is not encrypted, and the value <b>1</b> indicates that the EVS disk is encrypted. If this parameter is not specified, EVS disks will not be encrypted by default.
__system__cmkid	String	CMK ID, which indicates encryption in <b>metadata</b> . This field is used with <b>__system__encrypted</b> .

**Table 4-646** Storage

Parameter	Type	Description
storageSelectors	Array of <b>StorageSelectors</b> objects	Disk selection. Matched disks are managed according to <b>matchLabels</b> and <b>storageType</b> .
storageGroups	Array of <b>StorageGroups</b> objects	A storage group consists of multiple storage devices. It is used to divide storage space.

**Table 4-647** StorageSelectors

Parameter	Type	Description
name	String	Selector name, used as the index of <b>selectorNames</b> in <b>storageGroup</b> . Therefore, the name of each selector must be unique.
storageType	String	Storage type. Only <b>evs</b> (EVS disks) and <b>local</b> (local disks) are supported. The local storage does not support disk selection. All local disks will form a VG, so only one storageSelector of the local type is allowed.
matchLabels	<b>matchLabels</b> object	Matching field of an EVS disk. The value can be <b>size</b> , <b>volumeType</b> , <b>metadataEncrypted</b> , <b>metadataCmkid</b> , or <b>count</b> in <b>DataVolume</b> .

**Table 4-648** matchLabels

Parameter	Type	Description
size	String	Matched disk size. If this parameter is left unspecified, the disk size is not limited. Example: 100
volumeType	String	EVS disk type. <b>SSD</b> , <b>GPSSD</b> , <b>SAS</b> , <b>ESSD</b> , and <b>SATA</b> are supported. If this parameter is left blank, there is no restriction on the disk type.
metadataEncrypted	String	Disk encryption identifier. If this parameter is set to <b>0</b> , it specifies that the disk is not encrypted. If this parameter is set to <b>1</b> , it specifies that the disk is encrypted. If this parameter is left blank, there is no restriction on the disk encryption identifier.
metadataCmkid	String	CMK ID of an encrypted disk. The value is a string of 36 bytes. If this parameter is left blank, there is no restriction on the disk key ID.
count	String	Number of disks to be selected. If this parameter is left blank, all disks of this type are selected.

**Table 4-649** StorageGroups

Parameter	Type	Description
name	String	Name of a virtual storage group, which must be unique. <b>NOTE</b> <ul style="list-style-type: none"> <li>If <b>cceManaged</b> is set to <b>true</b>, the name must be set to <b>vgpaas</b>.</li> <li>If the data disk is used as a temporary storage volume, the name must be <b>vg-everest-localvolume-ephemeral</b>.</li> <li>If the data disk is used as a persistent storage volume, the name must be <b>vg-everest-localvolume-persistent</b>.</li> </ul>
cceManaged	Boolean	Storage space for Kubernetes and runtime components. Only one group can be set to <b>true</b> . If this parameter is left blank, the default value <b>false</b> is used.
selectorNames	Array of strings	Corresponds to <b>name</b> in <b>storageSelectors</b> . One group can select multiple selectors, but one selector can be selected by only one group.
virtualSpaces	Array of <a href="#">VirtualSpace</a> objects	Detailed management of space configuration in a group.

**Table 4-650** VirtualSpace

Parameter	Type	Description
name	String	Name of a virtualSpace. Options: <b>share</b> , <b>kubernetes</b> , <b>runtime</b> , and <b>user</b> <ul style="list-style-type: none"> <li><b>share</b>: shared space configuration (with the runtime and Kubernetes partitions cancelled). <b>lvmConfig</b> needs to be configured.</li> <li><b>Kubernetes</b>: Kubernetes space configuration. <b>lvmConfig</b> needs to be configured.</li> <li><b>runtime</b>: runtime space configuration. <b>runtimeConfig</b> needs to be configured.</li> <li><b>user</b>: user space configuration. <b>lvmConfig</b> needs to be configured.</li> </ul>

Parameter	Type	Description
size	String	Size of a virtualSpace. The value must be an integer in percentage. Example: 90%. <b>NOTE</b> The value ranges from 10 to 90. The sum of percentages of all virtualSpaces in a group cannot exceed 100%.
lvmConfig	<b>LVMConfig</b> object	LVM configuration managements, applicable to the <i>share*</i> , <b>kubernetes</b> , and <b>user</b> spaces. Note that one virtual space supports only one config.
runtimeConfig	<b>RuntimeConfig</b> object	runtime configurations, applicable to the <b>runtime</b> space. Note that one virtual space supports only one config.

**Table 4-651** LVMConfig

Parameter	Type	Description
lvType	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.
path	String	Disk mount path. This parameter takes effect only in user configuration. The value is an absolute path. Digits, letters, periods (.), hyphens (-), and underscores (_) are allowed. <b>NOTE</b> Do not leave the path empty or set it to a key OS path such as the root directory. Key OS paths that cannot be specified: /, /home, /home/, /bin, /bin/, /lib, /lib/, /root, /root/, /boot, /boot/, /dev, /dev/, /etc, /etc/, /lost+found, /lost+found/, /mnt, /mnt/, /proc, /proc/, /sbin, /sbin/, /srv, /srv/, /tmp, /tmp/, /var, /var/, /media, /media/, /opt, /opt/, /selinux, /selinux/, /sys, /sys/, /usr, /usr/, /opt/cloud/, /mnt/paas/, /home/paas/, /var/paas/, /var/lib/, /var/script/

**Table 4-652** RuntimeConfig

Parameter	Type	Description
lvType	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.

**Table 4-653** NodePublicIP

Parameter	Type	Description
ids	Array of strings	IDs of existing EIPs. The quantity cannot be greater than the number of nodes to be created. <b>NOTE</b> If <b>ids</b> has been set, you do not need to set <b>count</b> and <b>eip</b> .
count	Integer	Number of EIPs to be dynamically created. <b>NOTE</b> <b>count</b> and <b>eip</b> must be set at the same time.
eip	<b>NodeEIPSpec</b> object	EIP configuration. <b>NOTE</b> This parameter is not supported when creating a node pool.

**Table 4-654** NodeEIPSpec

Parameter	Type	Description
iptype	String	EIP type, specified in <b>publicip.type</b> in the API for assigning an EIP.
bandwidth	<b>NodeBandwidth</b> object	Bandwidth parameters of the EIP

**Table 4-655** NodeBandwidth

Parameter	Type	Description
chargemode	String	<p>Bandwidth billing mode.</p> <ul style="list-style-type: none"> <li>• If this field is not specified, the billing is based on bandwidth.</li> <li>• If the field is null, the billing is based on bandwidth.</li> <li>• If the field value is <b>traffic</b>, the billing is based on traffic.</li> <li>• If the value is out of the preceding options, the cloud server will fail to be created.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• Billed by bandwidth: The billing will be based on the data transmission rate (in Mbps) of public networks. This billing mode is recommended if your bandwidth usage is higher than 10%.</li> <li>• Billed by traffic: The billing is based on the total amount of data (in GB) transmitted over the public network. This mode is available only when you are creating a pay-per-use node. This billing mode is recommended if your bandwidth usage is lower than 10%.</li> </ul>
size	Integer	Bandwidth size, specified in <b>bandwidth.size</b> in the API for assigning an EIP.
sharetype	String	Bandwidth sharing type. Value options: <b>PER</b> (exclusive bandwidth)

**Table 4-656** NodeNicSpec

Parameter	Type	Description
primaryNic	<b>NicSpec</b> object	Description of the primary NIC.
extNics	Array of <b>NicSpec</b> objects	<p>Extension NIC</p> <p><b>NOTE</b> This parameter is not supported when you add a node to a node pool.</p>

**Table 4-657 NicSpec**

Parameter	Type	Description
subnetId	String	Network ID of the subnet to which a NIC belongs. If <b>subnetId</b> is not specified when a primary NIC is creating, the cluster subnet will be used. If <b>subnetList</b> is also configured for a node pool, the <b>subnetList</b> field is used for adding subnets to the node pool. When creating an extension NIC, you must specify <b>subnetId</b> .
fixedIps	Array of strings	The IP address of the primary ENI is specified using <b>fixedIps</b> . The number of IP addresses cannot be greater than the number of created nodes. Either <b>fixedIps</b> or <b>ipBlock</b> can be specified. <b>fixedIps</b> cannot be specified for extension ENIs.
ipBlock	String	CIDR format of the primary NIC IP range. The IP address of the created node falls in this range. <b>fixedIps</b> and <b>ipBlock</b> cannot be specified at the same time.
subnetList	Array of strings	Network ID list of the subnet where the NIC resides. A maximum of 20 subnets can be configured for a node pool.

**Table 4-658 Taint**

Parameter	Type	Description
key	String	Key.
value	String	Value.
effect	String	Effect.

**Table 4-659 UserTag**

Parameter	Type	Description
key	String	Key of the cloud server label. The value cannot start with <b>CCE-</b> or <b>__type_baremetal</b> .
value	String	Value of the cloud server label.

**Table 4-660** Runtime

Parameter	Type	Description
name	String	<p>Container runtime. Default value:</p> <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>

**Table 4-661** NodeExtendParam

Parameter	Type	Description
ecs:performancetype	String	ECS flavor types. This field is returned in the response.
maxPods	Integer	<p>Maximum number of pods that can be created on a node, including the default system pods. Value range: 16 to 256.</p> <p>This limit prevents the node from being overloaded of pods.</p> <p>The number of pods that can be created on a node is determined by multiple parameters. For details, see <a href="#">Maximum Number of Pods That Can Be Created on a Node</a>.</p>



Parameter	Type	Description
DockerLVMConfigOverride	String	<p>Docker data disk configuration item. (This parameter has been discarded. Use the <b>storage</b> field instead.) Example default configuration: "DockerLVMConfigOverride":"dockerThinpool=vgpaas/90%VG;kubernetesLV=vgpaas/10%VG;diskType=evs;lvType=linear"</p> <p>By default, if no VD disk is available, an error occurs because the data disk fails to be found. Set <b>diskType</b> based on the actual drive letter type. The following fields are included:</p> <ul style="list-style-type: none"> <li>• <b>userLV</b> (optional): size of the user space, for example, <b>vgpaas/20%VG</b>.</li> <li>• <b>userPath</b> (optional): mount path of the user space, for example, <b>/home/wqt-test</b>.</li> <li>• <b>diskType</b>: disk type. Currently, only <b>evs</b>, <b>hdd</b>, and <b>ssd</b> are supported.</li> <li>• <b>lvType</b>: type of a logic volume. The value can be <b>linear</b> or <b>striped</b>.</li> <li>• <b>dockerThinpool</b>: Docker space size, for example, <b>vgpaas/60%VG</b>.</li> <li>• <b>kubernetesLV</b>: kubelet space size, for example, <b>vgpaas/20%VG</b>.</li> </ul>
dockerBaseSize	Integer	<p>Available disk space of a single container on a node, in GB. This parameter has been discarded. Use <b>containerBaseSize</b>.</p> <p>If this parameter is left blank or is set to <b>0</b>, the default value is used. In Device Mapper mode, the default value is <b>10</b>. In OverlayFS mode, the available space of a single container is not limited by default, and the <b>dockerBaseSize</b> setting takes effect only on nodes running EulerOS in clusters of the new versions.</p> <p>For details about how to allocate the space for the container runtime, see <a href="#">Data Disk Space Allocation</a>.</p> <p>When Device Mapper is used, it is recommended that <b>dockerBaseSize</b> be set to a value less than or equal to 80 GB. If the value is too large, the container runtime may fail to be started due to long initialization. If there are special requirements for the container disk space, you can mount an external or local storage device.</p>

Parameter	Type	Description
containerBaseSize	Integer	<p>Available disk space of a single container on a node, in GiB.</p> <p>If this parameter is left blank or is set to <b>0</b>, the default value is used. In OverlayFS mode, the available space of a single container is not limited by default. In Device Mapper mode, the default value is <b>10</b>, and the <b>containerBaseSize</b> setting takes effect only on nodes running EulerOS in clusters v1.23.14-r0, v1.25.9-r0, v1.27.6-r0, v1.28.4-r0 or later.</p> <p>For details about how to allocate the space for the container runtime, see <a href="#">Data Disk Space Allocation</a>.</p> <p>When Device Mapper is used, it is recommended that <b>containerBaseSize</b> be set to a value less than or equal to 80 GB. If the value is too large, the container runtime may fail to be started due to long initialization. If there are special requirements for the container disk space, you can mount an external or local storage device. In new versions, the Device Mapper mode is used only by BMSs in the same resource pool and is being discarded.</p> <p><b>NOTE</b> This parameter cannot be updated when a node pool is being updated.</p>
publicKey	String	Public key of a node.
alpha.cce/preInstall	String	<p>Pre-installation scripts. The total number of characters after transcoding cannot exceed 10240.</p> <p><b>NOTE</b> The input value must be encoded using Base64. The method is as follows:</p> <pre>echo -n "Content to be encoded"   base64</pre>
alpha.cce/postInstall	String	<p>Post-installation scripts. The total number of characters after transcoding cannot exceed 10240.</p> <p><b>NOTE</b> The input value must be encoded using Base64. The method is as follows:</p> <pre>echo -n "Content to be encoded"   base64</pre>
alpha.cce/NodeImageID	String	This parameter is required when a custom image is used to create a BMS node.

Parameter	Type	Description
chargingMode	Integer	Billing mode of a node. This parameter has been deprecated. Use the <b>billingMode</b> parameter in <b>NodeSpec</b> .
agency_name	String	Name of an agency An agency is created by a tenant administrator on Identity and Access Management (IAM) to provide temporary credentials for CCE nodes to access cloud servers. This parameter is returned only when it is transferred during node creation.
kubeReservedMem	Integer	Reserved node memory, which is reserved for Kubernetes components.
systemReservedMem	Integer	Reserved node memory, which is reserved for system components.
init-node-password	String	Node password. If this parameter is used as a response parameter, asterisks (*) will be displayed.

**Table 4-662** HostnameConfig

Parameter	Type	Description
type	String	Configuration type of the Kubernetes node name. The default value is <b>privatelp</b> . <ul style="list-style-type: none"> <li>• <b>privatelp</b>: The Kubernetes node is named after its IP address.</li> <li>• <b>cceNodeName</b>: The Kubernetes node is named after the CCE node.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• For a node which is configured using <b>cceNodeName</b>, the name is the same as the Kubernetes node name and the ECS name. The node name cannot be changed. If the ECS name is changed on the ECS console, the node name will retain unchanged after ECS synchronization.</li> <li>• For a node which is configured using <b>cceNodeName</b>, to avoid a conflict between Kubernetes nodes, the system automatically adds a suffix to each node name. The suffix is in the format of <b>A hyphen (-) Five random characters</b>. The value of the random characters is a lowercase letter or a digit ranging from 0 to 9.</li> </ul>

**Table 4-663** NodePoolNodeAutoscaling

Parameter	Type	Description
enable	Boolean	Whether to enable auto scaling.
minNodeCount	Integer	Minimum number of nodes allowed if auto scaling is enabled. The value cannot be greater than the maximum number of nodes allowed by the cluster specifications.
maxNodeCount	Integer	Maximum number of nodes allowed if auto scaling is enabled. This value must be greater than or equal to the value of <b>minNodeCount</b> and cannot exceed the maximum number of nodes in the cluster specifications.
scaleDownCooldownTime	Integer	Interval between two scaling operations, in minutes. During this period, nodes added after a scale-up will not be deleted.
priority	Integer	Weight of a node pool. A node pool with a higher weight has a higher priority during scaling.

**Table 4-664** NodeManagement

Parameter	Type	Description
serverGroupReference	String	Cloud server group ID. If this field is specified, all nodes in the node pool will be created in this group. The group ID can be specified only when you create the node pool and cannot be modified. When you specify a cloud server group, the number of nodes in the node pool cannot exceed the group quota.

**Table 4-665** SecurityID

Parameter	Type	Description
id	String	Security group ID

**Table 4-666** ExtensionScaleGroup

Parameter	Type	Description
metadata	<a href="#">ExtensionScaleGroupMetadata</a> object	Basic information about an extended scaling group

Parameter	Type	Description
spec	<a href="#">ExtensionScaleGroupSpec</a> object	Configurations of an extended scaling group, which carry different configurations from those of the default scaling group

**Table 4-667** ExtensionScaleGroupMetadata

Parameter	Type	Description
uid	String	UUID of an extended scaling group, which is automatically generated by the system.
name	String	Name of an extended scaling group. The value cannot be <b>default</b> and can contain a maximum of 55 characters. Only digits, lowercase letters, and hyphens (-) are allowed.

**Table 4-668** ExtensionScaleGroupSpec

Parameter	Type	Description
flavor	String	Node flavor
az	String	AZ of a node. If this parameter is not specified or left blank, the default scaling group configurations take effect.
capacityReservationSpecification	<a href="#">CapacityReservationSpecification</a> object	Capacity reservation configurations of an extended scaling group
autoscaling	<a href="#">ScaleGroupAutoscaling</a> object	Auto scaling configurations of an extended scaling group

**Table 4-669** CapacityReservationSpecification

Parameter	Type	Description
id	String	Private pool ID. The parameter value can be ignored when <b>preference</b> is set to <b>none</b> .
preference	String	Capacity of a private storage pool. If the value is <b>none</b> , the capacity reservation is not specified. If the value is <b>targeted</b> , the capacity reservation is specified. In this case, the ID cannot be left blank.

**Table 4-670** ScaleGroupAutoscaling

Parameter	Type	Description
enable	Boolean	Whether to enable auto scaling for a scaling group. This function is disabled by default.
extensionPriority	Integer	Priority of a scaling group. If no value is specified for this parameter, it defaults to <b>0</b> . A higher value indicates a greater priority.
minNodeCount	Integer	Minimum number of nodes in a scaling group during auto scaling. The value must be greater than <b>0</b> .
maxNodeCount	Integer	Maximum number of nodes that can be retained in a scaling group during auto scaling. The value must be greater than or equal to that of <b>minNodeCount</b> , and can neither be greater than the maximum number of nodes allowed by the cluster nor the maximum number of nodes in the node pool.

**Table 4-671** UpdateNodePoolStatus

Parameter	Type	Description
currentNode	Integer	Total number of nodes in the current node pool (excluding the nodes that are being deleted).
creatingNode	Integer	Number of nodes that are being created in the node pool.
deletingNode	Integer	Number of nodes that are being deleted in the current node pool.

Parameter	Type	Description
phase	String	<p>Node pool status. Options:</p> <ul style="list-style-type: none"> <li>• <b>Null:</b> The node pool is available. (The number of nodes in the current node pool has reached the expected value, and no node scaling is being performed.)</li> <li>• <b>Synchronizing:</b> The node pool is being scaled. (The number of nodes in the current node pool has not reached the expected value, and no node scaling is being performed.)</li> <li>• <b>Synchronized:</b> The node pool scaling is pending. (The number of nodes in the current node pool has not reached the expected value, or node scaling is being performed.)</li> <li>• <b>SoldOut:</b> The node pool cannot be scaled out. (This field is used in multiple scenarios, for example, resources in the node pool have been sold out or the resource quota is insufficient.)</li> </ul> <p><b>NOTE</b> This node pool status has been deprecated. It is reserved only for compatibility. Do not use it anymore. Use the following instead:</p> <ul style="list-style-type: none"> <li>• Node pool scaling status: You can obtain the status of the current node pool using parameters such as <b>currentNode</b>, <b>creatingNode</b>, or <b>deletingNode</b>.</li> <li>• Node pool scale-out: You can use <b>conditions</b> to obtain the detailed status of a node pool. <b>Scalable</b> can replace <b>SoldOut</b>.</li> <li>• <b>Deleting:</b> The node pool is being deleted.</li> <li>• <b>Error:</b> An error occurred in the node pool.</li> </ul>
conditions	Array of <b>NodePoolCondition</b> objects	Node pool status details. For details, see the definition of <b>Condition</b> .
scaleGroupStatuses	Array of <b>ScaleGroupStatus</b> objects	Detailed status of a scaling group. For details, see the definition of <b>ScaleGroupStatus</b> .

**Table 4-672** ScaleGroupStatus

Parameter	Type	Description
name	String	Scaling group name

Parameter	Type	Description
uid	String	UUID of a scaling group
creationTimes tamp	String	The time when a scaling group was created
updateTimest amp	String	The time when a scaling group was updated
phase	String	<p>Status of a scaling group.</p> <ul style="list-style-type: none"> <li>• <b>Null:</b> The scaling group is available. (The number of nodes in the current scaling group has reached the expected value, and no node scaling is being performed.)</li> <li>• <b>Synchronizing:</b> The scaling group is being scaled. (The number of nodes in the current scaling group has not reached the expected value, and no node scaling is being performed.)</li> <li>• <b>Synchronized:</b> The scaling is pending. (The number of nodes in the current scaling group has not reached the expected value, or node scaling is being performed.)</li> <li>• <b>SoldOut:</b> The scaling group cannot be scaled out. (This field is used in multiple scenarios, for example, resources in the node pool have been sold out or the resource quota is insufficient.)</li> </ul> <p><b>NOTE</b> This scaling group status has been deprecated. It is reserved only for compatibility. Do not use it anymore. Use the following instead:</p> <ul style="list-style-type: none"> <li>• Scaling status of a scaling group: You can obtain the status of the current scaling group using parameters such as <b>desiredNodeCount</b>, <b>existingNodeCount</b>, or <b>upcomingNodeCount</b>.</li> <li>• Scaling group scale-out: You can use <b>conditions</b> to obtain the detailed status of a scaling group. <b>Scalable</b> can replace <b>SoldOut</b>.</li> <li>• <b>Deleting:</b> The scaling group being deleted.</li> <li>• <b>Error:</b> An error occurred in the scaling group.</li> </ul>
desiredNodeC ount	Integer	Expected number of nodes in a scaling group
unpaidScaleN odeCount	Integer	Number of unpaid nodes in an order
existingNodeC ount	<b>existingNode Count</b> object	Statistics on existing nodes in a scaling group



Parameter	Type	Description
upcomingNodeCount	<a href="#">upcomingNodeCount</a> object	Statistics on the nodes to be created in a scaling group
scaleDownDisabledNodeCount	Integer	Number of nodes that cannot be deleted in a scaling group
conditions	Array of <a href="#">NodePoolCondition</a> objects	Scaling group status details. For details, see the definition of <b>Condition</b> .

**Table 4-673** existingNodeCount

Parameter	Type	Description
postPaid	Integer	Number of pay-per-use nodes
total	Integer	Total number of pay-per-use nodes

**Table 4-674** upcomingNodeCount

Parameter	Type	Description
postPaid	Integer	Number of pay-per-use nodes
total	Integer	Total number of pay-per-use nodes

**Table 4-675** NodePoolCondition

Parameter	Type	Description
type	String	Condition type. The options are as follows: <ul style="list-style-type: none"> <li>● <b>Scalable</b>: whether a node pool can be scaled. If the status is <b>False</b>, node pool scaling will not be triggered again.</li> <li>● <b>QuotaInsufficient</b>: Quotas on which node pool scaling depends are insufficient, affecting the node pool scaling status.</li> <li>● <b>ResourceInsufficient</b>: Resources on which node pool scaling depends are insufficient, affecting the node pool scaling status.</li> <li>● <b>UnexpectedError</b>: The node pool fails to be scaled out due to unexpected reasons, affecting the node pool scaling status.</li> <li>● <b>Error</b>: A node pool error occurs. A common trigger is deletion failure.</li> </ul>
status	String	Current status of <b>Condition</b> . The options are as follows: <ul style="list-style-type: none"> <li>● "True"</li> <li>● "False"</li> </ul>
lastProbeTime	String	Time when the status was last checked
lastTransitTime	String	Time when the status was last changed
reason	String	Reason why the status was last changed
message	String	Detailed condition description

### Example Requests

- Change the number of nodes in the node pool to 1.  
PUT /api/v3/projects/{project\_id}/clusters/{cluster\_id}/nodepools/{nodepool\_id}

```
{
  "metadata": {
    "name": "lc-it-nodepool-3"
  },
  "spec": {
    "nodeTemplate": {
      "k8sTags": { },
      "taints": [ {
        "key": "status",
        "value": "unavailable",
        "effect": "NoSchedule"
      } ],
      "userTags": [ ]
    },
    "autoscaling": {
      "enable": false,
      "minNodeCount": 0,

```

```

    "maxNodeCount" : 0,
    "scaleDownCooldownTime" : 0,
    "priority" : 0
  },
  "initialNodeCount" : 1
}
}

```

- **Modify the key pair.**

PUT /api/v3/projects/{project\_id}/clusters/{cluster\_id}/nodepools/{nodepool\_id}

```

{
  "metadata" : {
    "name" : "lc-it-nodepool-2"
  },
  "spec" : {
    "nodeTemplate" : {
      "k8sTags" : { },
      "taints" : [ {
        "key" : "status",
        "value" : "unavailable",
        "effect" : "NoSchedule"
      } ],
      "userTags" : [ ],
      "login" : {
        "sshKey" : "KeyPair-IES"
      }
    },
    "autoscaling" : {
      "enable" : false,
      "minNodeCount" : 0,
      "maxNodeCount" : 0,
      "scaleDownCooldownTime" : 0,
      "priority" : 0
    },
    "initialNodeCount" : 1
  }
}

```

## Example Responses

### Status code: 200

Information about the specified node pool is successfully updated.

```

{
  "kind" : "NodePool",
  "apiVersion" : "v3",
  "metadata" : {
    "name" : "lc-it-nodepool-3",
    "uid" : "1deef848-690d-11ea-a11b-0255ac1001b7"
  },
  "spec" : {
    "initialNodeCount" : 1,
    "type" : "vm",
    "nodeTemplate" : {
      "flavor" : "Sit3.xlarge.2",
      "az" : "*****",
      "os" : "EulerOS 2.5",
      "login" : {
        "sshKey" : "KeyPair-001"
      }
    },
    "rootVolume" : {
      "volumetype" : "SAS",
      "size" : 40
    },
    "dataVolumes" : [ {
      "volumetype" : "SAS",
      "size" : 100,

```

```

"extendParam" : {
  "useType" : "docker"
},
},
"publicIP" : {
  "eip" : {
    "bandwidth" : { }
  }
},
"nodeNicSpec" : {
  "primaryNic" : {
    "subnetId" : "7e767d10-7548-4df5-ad72-aeac1d08bd8a"
  }
},
"billingMode" : 0,
"extendParam" : {
  "maxPods" : 110
},
"k8sTags" : {
  "cce.cloud.com/cce-nodepool" : "lc-it-nodepool-3"
}
},
"autoscaling" : { },
"nodeManagement" : { }
},
"status" : {
  "phase" : ""
}
}

```

## Status Codes

Status Code	Description
200	Information about the specified node pool is successfully updated.

## Error Codes

See [Error Codes](#).

## 4.4.5 Deleting a Node Pool

### Function

This API is used to delete a specified node pool.

#### NOTE

The URL for cluster management is in the format of `https://Endpoint/uri`. In the URL, **uri** indicates the resource path, that is, the path for API access.

### URI

DELETE `/api/v3/projects/{project_id}/clusters/{cluster_id}/nodepools/{nodepool_id}`

**Table 4-676** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
nodepool_id	Yes	String	Node pool ID.

## Request Parameters

**Table 4-677** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

## Response Parameters

**Status code: 200**

**Table 4-678** Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at <b>NodePool</b> .
apiVersion	String	API version. The value is fixed at <b>v3</b> .
metadata	<a href="#">NodePoolMetadata</a> object	Metadata of the node pool.
spec	<a href="#">NodePoolSpec</a> object	Node pool specifications. <b>NOTE</b> <ul style="list-style-type: none"> <li>This field is optional for updating a node pool.</li> </ul>

Parameter	Type	Description
status	<a href="#">DeleteNodePoolStatus</a> object	Node pool status.

**Table 4-679** NodePoolMetadata

Parameter	Type	Description
name	String	Node pool name. <b>NOTE</b> Naming rules: Enter 1 to 50 characters, starting with a lowercase letter and not ending with a hyphen (-). Only lowercase letters, digits, and hyphens (-) are allowed. <ul style="list-style-type: none"> <li>You cannot create node pools named DefaultPool.</li> </ul>
uid	String	UID of the node pool. The value is automatically generated after the object is updated. A user-defined value will not take effect.
annotations	Map<String,String>	Annotations of a node pool in key-value pairs. This parameter is used only for query and cannot be input through a request. Entered data of this parameter is invalid.
updateTimestamp	String	Update time.
creationTimestamp	String	Creation time.

**Table 4-680** NodePoolSpec

Parameter	Type	Description
type	String	Node pool type. If this parameter is left blank, the value <b>vm</b> is used by default. <ul style="list-style-type: none"> <li><b>vm</b>: ECS</li> <li><b>ElasticBMS</b>: C6 general computing-plus BMS. An example flavor is <b>c6.22xlarge.2.physical</b>.</li> <li><b>bms</b>: BMS</li> </ul>
nodeTemplate	<a href="#">NodeSpec</a> object	Detailed parameters of the node pool template.

Parameter	Type	Description
initialNodeCount	Integer	Initial number of nodes for the node pool. When queried, the value is the number of target nodes in the node pool.
autoscaling	<a href="#">NodePoolNodeAutoscaling</a> object	Auto scaling parameters
nodeManagement	<a href="#">NodeManagement</a> object	Node management configuration
podSecurityGroups	Array of <a href="#">SecurityID</a> objects	Security group configuration. This parameter is available only for CCE Turbo clusters.
extensionScaleGroups	Array of <a href="#">ExtensionScaleGroup</a> objects	Configuration list of the extended scaling groups in a node pool. For details, see the definition of <a href="#">ExtensionScaleGroup</a> .
customSecurityGroups	Array of strings	Custom security group settings for a node pool. New nodes scaled out in a node pool can be bound to a specified security group. <ul style="list-style-type: none"> <li>• Specifying no security group ID will add the new nodes to the default security group of the worker nodes.</li> <li>• Specifying a valid security group ID will put new nodes in that security group.</li> <li>• When specifying a security group, do not modify the rules of the port on which CCE running depends.</li> </ul>

**Table 4-681** NodeSpec

Parameter	Type	Description
flavor	String	Node specifications. For details about the node specifications supported by CCE, see <a href="#">Node Flavor Description</a> .
az	String	AZ where the node to be created is located. You need to specify the AZ name. A node created using APIs does not support random AZs. <b>NOTE</b> <ul style="list-style-type: none"> <li>• When you are creating a node pool and configuring a scaling group, this parameter cannot be set to <b>random</b>.</li> </ul>

Parameter	Type	Description
os	String	<p>Node OS. For details about the supported OSs, see <a href="#">Node OS</a>.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>The system automatically selects the supported OS based on the cluster version. If the current cluster version does not support the OS, an error will be reported.</li> <li>If <code>alpha.cce/NodeImageID</code> in <code>extendParam</code> is specified during node creation, you do not need to configure this parameter.</li> <li>This parameter is mandatory when creating a node pool.</li> <li>If you are using shared disk space when creating a node, this parameter is mandatory. Specifically, the disk initialization configuration management parameter <code>storage</code> is used and the <code>name</code> field of <code>virtualSpaces</code> in <code>StorageGroups</code> is set to <code>share</code>.</li> </ul>
login	<a href="#">Login</a> object	<p>Node login mode. Either the key pair or password must be used for login.</p> <p><b>NOTE</b> This parameter cannot be modified when a node pool is being updated.</p>
rootVolume	<a href="#">Volume</a> object	Information about disks on the node
dataVolumes	Array of <a href="#">Volume</a> objects	<p>Data disk parameters of a node. A VM can have up to 16 disks attached, while a BMS can only have up to 10 disks attached. The maximum number of disks that can be attached to either a VM or BMS is limited by their respective maximum disk limits. (You can add more data disks to each of your nodes through the CCE console and APIs.) The data disk used by the container runtime and kubelet cannot be detached, or the node will become unavailable. For DeC nodes, the parameter details are the same as those for <b>rootVolume</b>.</p>



Parameter	Type	Description
storage	<a href="#">Storage</a> object	<p>Disk initialization management parameter. This parameter is complex to configure. For details, see <a href="#">Attaching Disks to a Node</a>.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If a node has both local and EVS disks attached, make sure to not keep the default value for this parameter, or it may result in unexpected disk partitions.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If you want to change the value range of a data disk to 20 to 32768, do not use the default value.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>If you want to use the shared disk space (with the runtime and Kubernetes partitions cancelled), do not retain the default value of this parameter. For details about the shared disk space, see <a href="#">Data Disk Space Allocation</a>.</li> </ul>
publicIP	<a href="#">NodePublicIP</a> object	<p>EIP of a node.</p> <p><b>NOTE</b></p> <p>This parameter is not supported when you add a node to a node pool.</p>
nodeNicSpec	<a href="#">NodeNicSpec</a> object	NIC of the node
count	Integer	Number of nodes to be created in a batch. The value must be a positive integer greater than or equal to 1 and less than or equal to the defined limit. This parameter can be left blank when it is used for a node pool.
billingMode	Integer	<p>Node billing mode.</p> <ul style="list-style-type: none"> <li><b>0</b>: pay-per-use</li> </ul>

Parameter	Type	Description
taints	Array of <b>Taint</b> objects	<p>You can add taints to created nodes to configure anti-affinity. A maximum of 20 taints can be added. Each taint contains the following parameters:</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain can be prefixed to a key.</li> <li>• <b>Value:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed.</li> <li>• <b>Effect:</b> Available options are <b>NoSchedule</b>, <b>PreferNoSchedule</b>, and <b>NoExecute</b>. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</li> </ul> <p>Example:</p> <pre data-bbox="815 1111 1430 1339">"taints": [{   "key": "status",   "value": "unavailable",   "effect": "NoSchedule" }, {   "key": "looks",   "value": "bad",   "effect": "NoSchedule" }]</pre>

Parameter	Type	Description
k8sTags	Map<String,String>	<p>Defined in key-value pairs. A maximum of 20 key-value pairs are allowed.</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting and ending with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain containing a maximum of 253 characters (such as <b>example.com/my-key</b>) can be prefixed to a key.</li> <li>• <b>Value:</b> The value can be left blank or contain 1 to 63 characters that start and end with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</li> </ul> <p>Example:</p> <pre>"k8sTags": {   "key": "value" }</pre>
ecsGroupId	String	<p>ECS group ID. If this parameter is configured, nodes will be created in the specified ECS group.</p> <p><b>NOTE</b> This configuration does not take effect when a node pool is created. To ensure that all nodes in a node pool are in the same ECS group, configure <b>nodeManagement</b> in the node pool.</p>
dedicatedHostId	String	<p>ID of the DeH to which the node is scheduled.</p> <p><b>NOTE</b> This parameter is not supported when you add a node during node pool creation.</p>

Parameter	Type	Description
userTags	Array of <a href="#">UserTag</a> objects	<p>Cloud server tag (resource tag). The key of a tag must be unique. The maximum number of custom tags supported by CCE varies depending on regions and cannot exceed 8. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In node pools, the initial value can be specified in the node template and this field is returned during query. In other scenarios, this field is not returned during query.</p> <p><b>NOTE</b> A tag key can have a maximum of 36 characters. It can only contain letters, digits, Unicode characters, and special characters (-_).</p>
runtime	<a href="#">Runtime</a> object	<p>Container runtime. Default value:</p> <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>

Parameter	Type	Description
initializedConditions	Array of strings	<p>Custom initialization flag, which is left blank by default.</p> <p>Before CCE nodes are initialized, <b>node.cloudprovider.kubernetes.io/uninitialized</b> is added to the nodes to prevent pods from being scheduled to these nodes. When creating a node, you can configure the <b>initializedConditions</b> parameter to control the taint removal time. By default, the timeout period is not configured.</p> <p>Example:</p> <ol style="list-style-type: none"> <li>1. Create a node and add parameter "initializedConditions": ["CCEInitial", "CustomedInitial"].</li> <li>2. After custom initialization is complete, CCE will call a Kubernetes API (for example, <b>PATCH /v1/nodes/{node_ip}/status</b>) to update the node <b>conditions</b> by adding two labels of types <b>CCEInitial</b> and <b>CustomedInitial</b> and setting the status to <b>True</b>, as shown in the following:</li> </ol> <pre data-bbox="817 1055 1430 1205"> status:   conditions:   - type: CCEInitial     status: 'True'   - type: CustomedInitial     status: 'True'                     </pre> <ol style="list-style-type: none"> <li>1. CCE polls <b>status.Conditions</b> of nodes to check whether there are conditions of types <b>CCEInitial</b> and <b>CustomedInitial</b>. If such conditions exist and the status is <b>True</b>, the node initialization is complete and the initialization taint is removed.</li> <li>2. <b>initializedConditions</b> allows you to configure a timeout period during node creation. For example, "initializedConditions": ["CCEInitial:15m", "CustomedInitial:15m"] indicates that the timeout period is 15 minutes. After the timeout period is reached, the initialization conditions will be automatically ignored and the initialization taint will be removed when CCE polls the node. <ul style="list-style-type: none"> <li>● The timeout interval in <b>initializedConditions</b> ranges from 1 to 99 seconds.</li> <li>● Use only letters and digits. Max. characters: 20.</li> <li>● Max. flags: 2.</li> </ul> </li> </ol>

Parameter	Type	Description
		<ul style="list-style-type: none"> <li>The unit of the timeout period is minute (m).</li> </ul>
extendParam	<a href="#">NodeExtendParam</a> object	Extended parameters for creating a node.
hostnameConfig	<a href="#">HostnameConfig</a> object	Kubernetes node name configuration parameter, which is supported by clusters of v1.23.6-r0 to v1.25 or clusters of v1.25.2-r0 or later versions.
serverEnterpriseProjectID	String	Enterprise project ID of a server. CCE does not support features provided by EPS. This field is used only for synchronizing enterprise project IDs of servers. To create a node or node pool, you can specify an existing enterprise project for it. If this field is left blank, it inherits the enterprise project attributes of the cluster that runs the node or node pool. When updating a node pool, any configuration modifications you made will apply only to new nodes. You need to modify the enterprise projects of existing nodes on the EPS console. If no value is specified during the update, this field will not be updated. If this parameter is left blank, the cluster enterprise project is returned.

**Table 4-682** Login

Parameter	Type	Description
sshKey	String	Name of the key pair used for login.
userPassword	<a href="#">UserPassword</a> object	Password used for node login.

**Table 4-683** UserPassword

Parameter	Type	Description
username	String	Login account. The default value is <b>root</b> .

Parameter	Type	Description
password	String	<p>Login password. If a username and password are used to create a node, this field is shielded in the response body. A password must meet the following complexity requirements:</p> <ul style="list-style-type: none"> <li>• Contains 8 to 26 characters.</li> <li>• Contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters !@\$%^-_=+[{ }];,./?</li> <li>• Cannot contain the username or the username spelled backwards. The <b>password</b> field must be salted during node creation. For details, see <a href="#">Adding a Salt in the password Field When Creating a Node</a>.</li> </ul>

**Table 4-684** Volume

Parameter	Type	Description
size	Integer	<p>Disk size, in GiB.</p> <ul style="list-style-type: none"> <li>• System disk: 40 to 1024</li> <li>• Value range for data disks: 100 to 32768</li> </ul>
volumetype	String	<p>Disk type. For details about possible values, see the description of the <b>root_volume</b> parameter in the API used to create an ECS.</p> <ul style="list-style-type: none"> <li>• <b>SAS</b>: high I/O SAS disk</li> <li>• <b>SSD</b>: ultra-high I/O SSD disk</li> <li>• <b>SATA</b>: common I/O SATA disk. SATA disks have been removed from EVS. You can find them attached only on existing nodes.</li> </ul>
extendParam	Map<String, Object>	<p>Extended disk parameters, defined in <b>extendparam</b> in the API used to create an ECS.</p>
cluster_id	String	<p>ID of the storage pool used by the ECS system disk. This field is used only for DeC clusters, which functions as <b>dssPoolID</b>, that is, the ID of the DSS storage pool.</p>
cluster_type	String	<p>Storage class of the cloud server system disk. The value is always <b>dss</b>. This field is used only for DeC clusters.</p>

Parameter	Type	Description
hw:passthrough	Boolean	<ul style="list-style-type: none"> <li>Pay attention to this field if your ECS is SDI-compliant. If the value of this field is <b>true</b>, the created disk is of SCSI type.</li> <li>If the node pool type is <b>ElasticBMS</b>, this field must be set to <b>true</b>.</li> <li>If a node specification involves both local and EVS disks, configure the disk initialization parameters. For details, see <a href="#">Attaching Disks to a Node</a>.</li> </ul>
metadata	<b>VolumeMetadata</b> object	EVS disk encryption information. This field is mandatory only when you need to encrypt the system disk or data disks of the node to be created.

**Table 4-685** VolumeMetadata

Parameter	Type	Description
__system__encrypted	String	Whether the EVS disk is encrypted. The value <b>0</b> indicates that the EVS disk is not encrypted, and the value <b>1</b> indicates that the EVS disk is encrypted.  If this parameter is not specified, EVS disks will not be encrypted by default.
__system__cmkid	String	CMK ID, which indicates encryption in <b>metadata</b> . This field is used with <b>__system__encrypted</b> .

**Table 4-686** Storage

Parameter	Type	Description
storageSelectors	Array of <b>StorageSelectors</b> objects	Disk selection. Matched disks are managed according to <b>matchLabels</b> and <b>storageType</b> .
storageGroups	Array of <b>StorageGroups</b> objects	A storage group consists of multiple storage devices. It is used to divide storage space.



**Table 4-687** StorageSelectors

Parameter	Type	Description
name	String	Selector name, used as the index of <b>selectorNames</b> in <b>storageGroup</b> . Therefore, the name of each selector must be unique.
storageType	String	Storage type. Only <b>evs</b> (EVS disks) and <b>local</b> (local disks) are supported. The local storage does not support disk selection. All local disks will form a VG, so only one storageSelector of the local type is allowed.
matchLabels	<b>matchLabels</b> object	Matching field of an EVS disk. The value can be <b>size</b> , <b>volumeType</b> , <b>metadataEncrypted</b> , <b>metadataCmkid</b> , or <b>count</b> in <b>DataVolume</b> .

**Table 4-688** matchLabels

Parameter	Type	Description
size	String	Matched disk size. If this parameter is left unspecified, the disk size is not limited. Example: 100
volumeType	String	EVS disk type. <b>SSD</b> , <b>GPSSD</b> , <b>SAS</b> , <b>ESSD</b> , and <b>SATA</b> are supported. If this parameter is left blank, there is no restriction on the disk type.
metadataEncrypted	String	Disk encryption identifier. If this parameter is set to <b>0</b> , it specifies that the disk is not encrypted. If this parameter is set to <b>1</b> , it specifies that the disk is encrypted. If this parameter is left blank, there is no restriction on the disk encryption identifier.
metadataCmkid	String	CMK ID of an encrypted disk. The value is a string of 36 bytes. If this parameter is left blank, there is no restriction on the disk key ID.
count	String	Number of disks to be selected. If this parameter is left blank, all disks of this type are selected.

**Table 4-689** StorageGroups

Parameter	Type	Description
name	String	Name of a virtual storage group, which must be unique. <b>NOTE</b> <ul style="list-style-type: none"> <li>If <b>cceManaged</b> is set to <b>true</b>, the name must be set to <b>vgpaas</b>.</li> <li>If the data disk is used as a temporary storage volume, the name must be <b>vg-everest-localvolume-ephemeral</b>.</li> <li>If the data disk is used as a persistent storage volume, the name must be <b>vg-everest-localvolume-persistent</b>.</li> </ul>
cceManaged	Boolean	Storage space for Kubernetes and runtime components. Only one group can be set to <b>true</b> . If this parameter is left blank, the default value <b>false</b> is used.
selectorNames	Array of strings	Corresponds to <b>name</b> in <b>storageSelectors</b> . One group can select multiple selectors, but one selector can be selected by only one group.
virtualSpaces	Array of <a href="#">VirtualSpace</a> objects	Detailed management of space configuration in a group.

**Table 4-690** VirtualSpace

Parameter	Type	Description
name	String	Name of a virtualSpace. Options: <b>share</b> , <b>kubernetes</b> , <b>runtime</b> , and <b>user</b> <ul style="list-style-type: none"> <li><b>share</b>: shared space configuration (with the runtime and Kubernetes partitions cancelled). <b>lvmConfig</b> needs to be configured.</li> <li><b>Kubernetes</b>: Kubernetes space configuration. <b>lvmConfig</b> needs to be configured.</li> <li><b>runtime</b>: runtime space configuration. <b>runtimeConfig</b> needs to be configured.</li> <li><b>user</b>: user space configuration. <b>lvmConfig</b> needs to be configured.</li> </ul>

Parameter	Type	Description
size	String	Size of a virtualSpace. The value must be an integer in percentage. Example: 90%. <b>NOTE</b> The value ranges from 10 to 90. The sum of percentages of all virtualSpaces in a group cannot exceed 100%.
lvmConfig	<b>LVMConfig</b> object	LVM configuration managements, applicable to the <i>share*</i> , <b>kubernetes</b> , and <b>user</b> spaces. Note that one virtual space supports only one config.
runtimeConfig	<b>RuntimeConfig</b> object	runtime configurations, applicable to the <b>runtime</b> space. Note that one virtual space supports only one config.

**Table 4-691** LVMConfig

Parameter	Type	Description
lvType	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.
path	String	Disk mount path. This parameter takes effect only in user configuration. The value is an absolute path. Digits, letters, periods (.), hyphens (-), and underscores (_) are allowed. <b>NOTE</b> Do not leave the path empty or set it to a key OS path such as the root directory. Key OS paths that cannot be specified: /, /home, /home/, /bin, /bin/, /lib, /lib/, /root, /root/, /boot, /boot/, /dev, /dev/, /etc, /etc/, /lost+found, /lost+found/, /mnt, /mnt/, /proc, /proc/, /sbin, /sbin/, /srv, /srv/, /tmp, /tmp/, /var, /var/, /media, /media/, /opt, /opt/, /selinux, /selinux/, /sys, /sys/, /usr, /usr/, /opt/cloud/, /mnt/paas/, /home/paas/, /var/paas/, /var/lib/, /var/script/

**Table 4-692** RuntimeConfig

Parameter	Type	Description
lvType	String	LVM write mode. <b>linear</b> indicates the linear mode. <b>striped</b> indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.

**Table 4-693** NodePublicIP

Parameter	Type	Description
ids	Array of strings	IDs of existing EIPs. The quantity cannot be greater than the number of nodes to be created. <b>NOTE</b> If <b>ids</b> has been set, you do not need to set <b>count</b> and <b>eip</b> .
count	Integer	Number of EIPs to be dynamically created. <b>NOTE</b> <b>count</b> and <b>eip</b> must be set at the same time.
eip	<b>NodeEIPSpec</b> object	EIP configuration. <b>NOTE</b> This parameter is not supported when creating a node pool.

**Table 4-694** NodeEIPSpec

Parameter	Type	Description
iptype	String	EIP type, specified in <b>publicip.type</b> in the API for assigning an EIP.
bandwidth	<b>NodeBandwidth</b> object	Bandwidth parameters of the EIP

**Table 4-695** NodeBandwidth

Parameter	Type	Description
chargemode	String	<p>Bandwidth billing mode.</p> <ul style="list-style-type: none"> <li>• If this field is not specified, the billing is based on bandwidth.</li> <li>• If the field is null, the billing is based on bandwidth.</li> <li>• If the field value is <b>traffic</b>, the billing is based on traffic.</li> <li>• If the value is out of the preceding options, the cloud server will fail to be created.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• Billed by bandwidth: The billing will be based on the data transmission rate (in Mbps) of public networks. This billing mode is recommended if your bandwidth usage is higher than 10%.</li> <li>• Billed by traffic: The billing is based on the total amount of data (in GB) transmitted over the public network. This mode is available only when you are creating a pay-per-use node. This billing mode is recommended if your bandwidth usage is lower than 10%.</li> </ul>
size	Integer	Bandwidth size, specified in <b>bandwidth.size</b> in the API for assigning an EIP.
sharetype	String	Bandwidth sharing type. Value options: <b>PER</b> (exclusive bandwidth)

**Table 4-696** NodeNicSpec

Parameter	Type	Description
primaryNic	<b>NicSpec</b> object	Description of the primary NIC.
extNics	Array of <b>NicSpec</b> objects	<p>Extension NIC</p> <p><b>NOTE</b> This parameter is not supported when you add a node to a node pool.</p>

**Table 4-697 NicSpec**

Parameter	Type	Description
subnetId	String	Network ID of the subnet to which a NIC belongs. If <b>subnetId</b> is not specified when a primary NIC is creating, the cluster subnet will be used. If <b>subnetList</b> is also configured for a node pool, the <b>subnetList</b> field is used for adding subnets to the node pool. When creating an extension NIC, you must specify <b>subnetId</b> .
fixedIps	Array of strings	The IP address of the primary ENI is specified using <b>fixedIps</b> . The number of IP addresses cannot be greater than the number of created nodes. Either <b>fixedIps</b> or <b>ipBlock</b> can be specified. <b>fixedIps</b> cannot be specified for extension ENIs.
ipBlock	String	CIDR format of the primary NIC IP range. The IP address of the created node falls in this range. <b>fixedIps</b> and <b>ipBlock</b> cannot be specified at the same time.
subnetList	Array of strings	Network ID list of the subnet where the NIC resides. A maximum of 20 subnets can be configured for a node pool.

**Table 4-698 Taint**

Parameter	Type	Description
key	String	Key.
value	String	Value.
effect	String	Effect.

**Table 4-699 UserTag**

Parameter	Type	Description
key	String	Key of the cloud server label. The value cannot start with <b>CCE-</b> or <b>__type_baremetal</b> .
value	String	Value of the cloud server label.

**Table 4-700** Runtime

Parameter	Type	Description
name	String	<p>Container runtime. Default value:</p> <ul style="list-style-type: none"> <li>• Clusters earlier than v1.25: The default value is <b>docker</b>.</li> <li>• Clusters of v1.25 or later: The default container runtime varies depending on the OS. For nodes running EulerOS 2.5, the default container runtime is <b>docker</b>; for nodes running other OSs, the default container runtime is <b>containerd</b>.</li> </ul>

**Table 4-701** NodeExtendParam

Parameter	Type	Description
ecs:performancetype	String	ECS flavor types. This field is returned in the response.
maxPods	Integer	<p>Maximum number of pods that can be created on a node, including the default system pods. Value range: 16 to 256.</p> <p>This limit prevents the node from being overloaded of pods.</p> <p>The number of pods that can be created on a node is determined by multiple parameters. For details, see <a href="#">Maximum Number of Pods That Can Be Created on a Node</a>.</p>

Parameter	Type	Description
DockerLVMConfigOverride	String	<p>Docker data disk configuration item. (This parameter has been discarded. Use the <b>storage</b> field instead.) Example default configuration: "DockerLVMConfigOverride":"dockerThinpool=vgpaas/90%VG;kubernetesLV=vgpaas/10%VG;diskType=evs;lvType=linear"</p> <p>By default, if no VD disk is available, an error occurs because the data disk fails to be found. Set <b>diskType</b> based on the actual drive letter type. The following fields are included:</p> <ul style="list-style-type: none"> <li>• <b>userLV</b> (optional): size of the user space, for example, <b>vgpaas/20%VG</b>.</li> <li>• <b>userPath</b> (optional): mount path of the user space, for example, <b>/home/wqt-test</b>.</li> <li>• <b>diskType</b>: disk type. Currently, only <b>evs</b>, <b>hdd</b>, and <b>ssd</b> are supported.</li> <li>• <b>lvType</b>: type of a logic volume. The value can be <b>linear</b> or <b>striped</b>.</li> <li>• <b>dockerThinpool</b>: Docker space size, for example, <b>vgpaas/60%VG</b>.</li> <li>• <b>kubernetesLV</b>: kubelet space size, for example, <b>vgpaas/20%VG</b>.</li> </ul>
dockerBaseSize	Integer	<p>Available disk space of a single container on a node, in GB. This parameter has been discarded. Use <b>containerBaseSize</b>.</p> <p>If this parameter is left blank or is set to <b>0</b>, the default value is used. In Device Mapper mode, the default value is <b>10</b>. In OverlayFS mode, the available space of a single container is not limited by default, and the <b>dockerBaseSize</b> setting takes effect only on nodes running EulerOS in clusters of the new versions.</p> <p>For details about how to allocate the space for the container runtime, see <a href="#">Data Disk Space Allocation</a>.</p> <p>When Device Mapper is used, it is recommended that <b>dockerBaseSize</b> be set to a value less than or equal to 80 GB. If the value is too large, the container runtime may fail to be started due to long initialization. If there are special requirements for the container disk space, you can mount an external or local storage device.</p>



Parameter	Type	Description
containerBaseSize	Integer	<p>Available disk space of a single container on a node, in GiB.</p> <p>If this parameter is left blank or is set to <b>0</b>, the default value is used. In OverlayFS mode, the available space of a single container is not limited by default. In Device Mapper mode, the default value is <b>10</b>, and the <b>containerBaseSize</b> setting takes effect only on nodes running EulerOS in clusters v1.23.14-r0, v1.25.9-r0, v1.27.6-r0, v1.28.4-r0 or later.</p> <p>For details about how to allocate the space for the container runtime, see <a href="#">Data Disk Space Allocation</a>.</p> <p>When Device Mapper is used, it is recommended that <b>containerBaseSize</b> be set to a value less than or equal to 80 GB. If the value is too large, the container runtime may fail to be started due to long initialization. If there are special requirements for the container disk space, you can mount an external or local storage device. In new versions, the Device Mapper mode is used only by BMSs in the same resource pool and is being discarded.</p> <p><b>NOTE</b> This parameter cannot be updated when a node pool is being updated.</p>
publicKey	String	Public key of a node.
alpha.cce/preInstall	String	<p>Pre-installation scripts. The total number of characters after transcoding cannot exceed 10240.</p> <p><b>NOTE</b> The input value must be encoded using Base64. The method is as follows:</p> <pre>echo -n "Content to be encoded"   base64</pre>
alpha.cce/postInstall	String	<p>Post-installation scripts. The total number of characters after transcoding cannot exceed 10240.</p> <p><b>NOTE</b> The input value must be encoded using Base64. The method is as follows:</p> <pre>echo -n "Content to be encoded"   base64</pre>
alpha.cce/NodeImageID	String	This parameter is required when a custom image is used to create a BMS node.

Parameter	Type	Description
chargingMode	Integer	Billing mode of a node. This parameter has been deprecated. Use the <b>billingMode</b> parameter in <b>NodeSpec</b> .
agency_name	String	Name of an agency An agency is created by a tenant administrator on Identity and Access Management (IAM) to provide temporary credentials for CCE nodes to access cloud servers. This parameter is returned only when it is transferred during node creation.
kubeReservedMem	Integer	Reserved node memory, which is reserved for Kubernetes components.
systemReservedMem	Integer	Reserved node memory, which is reserved for system components.
init-node-password	String	Node password. If this parameter is used as a response parameter, asterisks (*) will be displayed.

**Table 4-702** HostnameConfig

Parameter	Type	Description
type	String	Configuration type of the Kubernetes node name. The default value is <b>privatelp</b> . <ul style="list-style-type: none"> <li>• <b>privatelp</b>: The Kubernetes node is named after its IP address.</li> <li>• <b>cceNodeName</b>: The Kubernetes node is named after the CCE node.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• For a node which is configured using <b>cceNodeName</b>, the name is the same as the Kubernetes node name and the ECS name. The node name cannot be changed. If the ECS name is changed on the ECS console, the node name will retain unchanged after ECS synchronization.</li> <li>• For a node which is configured using <b>cceNodeName</b>, to avoid a conflict between Kubernetes nodes, the system automatically adds a suffix to each node name. The suffix is in the format of <b>A hyphen (-) Five random characters</b>. The value of the random characters is a lowercase letter or a digit ranging from 0 to 9.</li> </ul>

**Table 4-703** NodePoolNodeAutoscaling

Parameter	Type	Description
enable	Boolean	Whether to enable auto scaling.
minNodeCount	Integer	Minimum number of nodes allowed if auto scaling is enabled. The value cannot be greater than the maximum number of nodes allowed by the cluster specifications.
maxNodeCount	Integer	Maximum number of nodes allowed if auto scaling is enabled. This value must be greater than or equal to the value of <b>minNodeCount</b> and cannot exceed the maximum number of nodes in the cluster specifications.
scaleDownCooldownTime	Integer	Interval between two scaling operations, in minutes. During this period, nodes added after a scale-up will not be deleted.
priority	Integer	Weight of a node pool. A node pool with a higher weight has a higher priority during scaling.

**Table 4-704** NodeManagement

Parameter	Type	Description
serverGroupReference	String	Cloud server group ID. If this field is specified, all nodes in the node pool will be created in this group. The group ID can be specified only when you create the node pool and cannot be modified. When you specify a cloud server group, the number of nodes in the node pool cannot exceed the group quota.

**Table 4-705** SecurityID

Parameter	Type	Description
id	String	Security group ID

**Table 4-706** ExtensionScaleGroup

Parameter	Type	Description
metadata	<a href="#">ExtensionScaleGroupMetadata</a> object	Basic information about an extended scaling group

Parameter	Type	Description
spec	<a href="#">ExtensionScaleGroupSpec</a> object	Configurations of an extended scaling group, which carry different configurations from those of the default scaling group

**Table 4-707** ExtensionScaleGroupMetadata

Parameter	Type	Description
uid	String	UUID of an extended scaling group, which is automatically generated by the system.
name	String	Name of an extended scaling group. The value cannot be <b>default</b> and can contain a maximum of 55 characters. Only digits, lowercase letters, and hyphens (-) are allowed.

**Table 4-708** ExtensionScaleGroupSpec

Parameter	Type	Description
flavor	String	Node flavor
az	String	AZ of a node. If this parameter is not specified or left blank, the default scaling group configurations take effect.
capacityReservationSpecification	<a href="#">CapacityReservationSpecification</a> object	Capacity reservation configurations of an extended scaling group
autoscaling	<a href="#">ScaleGroupAutoscaling</a> object	Auto scaling configurations of an extended scaling group

**Table 4-709** CapacityReservationSpecification

Parameter	Type	Description
id	String	Private pool ID. The parameter value can be ignored when <b>preference</b> is set to <b>none</b> .
preference	String	Capacity of a private storage pool. If the value is <b>none</b> , the capacity reservation is not specified. If the value is <b>targeted</b> , the capacity reservation is specified. In this case, the ID cannot be left blank.

**Table 4-710** ScaleGroupAutoscaling

Parameter	Type	Description
enable	Boolean	Whether to enable auto scaling for a scaling group. This function is disabled by default.
extensionPriority	Integer	Priority of a scaling group. If no value is specified for this parameter, it defaults to <b>0</b> . A higher value indicates a greater priority.
minNodeCount	Integer	Minimum number of nodes in a scaling group during auto scaling. The value must be greater than <b>0</b> .
maxNodeCount	Integer	Maximum number of nodes that can be retained in a scaling group during auto scaling. The value must be greater than or equal to that of <b>minNodeCount</b> , and can neither be greater than the maximum number of nodes allowed by the cluster nor the maximum number of nodes in the node pool.

**Table 4-711** DeleteNodePoolStatus

Parameter	Type	Description
currentNode	Integer	Total number of nodes in the current node pool (excluding the nodes that are being deleted).
creatingNode	Integer	Number of nodes that are being created in the node pool.
deletingNode	Integer	Number of nodes that are being deleted in the current node pool.

Parameter	Type	Description
phase	String	<p>Node pool status. Options:</p> <ul style="list-style-type: none"> <li>• <b>Null:</b> The node pool is available. (The number of nodes in the current node pool has reached the expected value, and no node scaling is being performed.)</li> <li>• <b>Synchronizing:</b> The node pool is being scaled. (The number of nodes in the current node pool has not reached the expected value, and no node scaling is being performed.)</li> <li>• <b>Synchronized:</b> The node pool scaling is pending. (The number of nodes in the current node pool has not reached the expected value, or node scaling is being performed.)</li> <li>• <b>SoldOut:</b> The node pool cannot be scaled out. (This field is used in multiple scenarios, for example, resources in the node pool have been sold out or the resource quota is insufficient.)</li> </ul> <p><b>NOTE</b> This node pool status has been deprecated. It is reserved only for compatibility. Do not use it anymore. Use the following instead:</p> <ul style="list-style-type: none"> <li>• Node pool scaling status: You can obtain the status of the current node pool using parameters such as <b>currentNode</b>, <b>creatingNode</b>, or <b>deletingNode</b>.</li> <li>• Node pool scale-out: You can use <b>conditions</b> to obtain the detailed status of a node pool. <b>Scalable</b> can replace <b>SoldOut</b>.</li> <li>• <b>Deleting:</b> The node pool is being deleted.</li> <li>• <b>Error:</b> An error occurred in the node pool.</li> </ul>
jobId	String	ID of a job executed on the node pool.
conditions	Array of <b>NodePoolCondition</b> objects	Node pool status details. For details, see the definition of <b>Condition</b> .
scaleGroupStatuses	Array of <b>ScaleGroupStatus</b> objects	Detailed status of a scaling group. For details, see the definition of <b>ScaleGroupStatus</b> .

**Table 4-712** ScaleGroupStatus

Parameter	Type	Description
name	String	Scaling group name
uid	String	UUID of a scaling group
creationTimes tamp	String	The time when a scaling group was created
updateTimest amp	String	The time when a scaling group was updated
phase	String	<p>Status of a scaling group.</p> <ul style="list-style-type: none"> <li>• <b>Null:</b> The scaling group is available. (The number of nodes in the current scaling group has reached the expected value, and no node scaling is being performed.)</li> <li>• <b>Synchronizing:</b> The scaling group is being scaled. (The number of nodes in the current scaling group has not reached the expected value, and no node scaling is being performed.)</li> <li>• <b>Synchronized:</b> The scaling is pending. (The number of nodes in the current scaling group has not reached the expected value, or node scaling is being performed.)</li> <li>• <b>SoldOut:</b> The scaling group cannot be scaled out. (This field is used in multiple scenarios, for example, resources in the node pool have been sold out or the resource quota is insufficient.)</li> </ul> <p><b>NOTE</b> This scaling group status has been deprecated. It is reserved only for compatibility. Do not use it anymore. Use the following instead:</p> <ul style="list-style-type: none"> <li>• Scaling status of a scaling group: You can obtain the status of the current scaling group using parameters such as <b>desiredNodeCount</b>, <b>existingNodeCount</b>, or <b>upcomingNodeCount</b>.</li> <li>• Scaling group scale-out: You can use <b>conditions</b> to obtain the detailed status of a scaling group. <b>Scalable</b> can replace <b>SoldOut</b>.</li> <li>• <b>Deleting:</b> The scaling group being deleted.</li> <li>• <b>Error:</b> An error occurred in the scaling group.</li> </ul>
desiredNodeC ount	Integer	Expected number of nodes in a scaling group
unpaidScaleN odeCount	Integer	Number of unpaid nodes in an order

Parameter	Type	Description
existingNodeCount	<a href="#">existingNodeCount</a> object	Statistics on existing nodes in a scaling group
upcomingNodeCount	<a href="#">upcomingNodeCount</a> object	Statistics on the nodes to be created in a scaling group
scaleDownDisabledNodeCount	Integer	Number of nodes that cannot be deleted in a scaling group
conditions	Array of <a href="#">NodePoolCondition</a> objects	Scaling group status details. For details, see the definition of <b>Condition</b> .

**Table 4-713** existingNodeCount

Parameter	Type	Description
postPaid	Integer	Number of pay-per-use nodes
total	Integer	Total number of pay-per-use nodes

**Table 4-714** upcomingNodeCount

Parameter	Type	Description
postPaid	Integer	Number of pay-per-use nodes
total	Integer	Total number of pay-per-use nodes



**Table 4-715** NodePoolCondition

Parameter	Type	Description
type	String	Condition type. The options are as follows: <ul style="list-style-type: none"> <li>• <b>Scalable</b>: whether a node pool can be scaled. If the status is <b>False</b>, node pool scaling will not be triggered again.</li> <li>• <b>QuotaInsufficient</b>: Quotas on which node pool scaling depends are insufficient, affecting the node pool scaling status.</li> <li>• <b>ResourceInsufficient</b>: Resources on which node pool scaling depends are insufficient, affecting the node pool scaling status.</li> <li>• <b>UnexpectedError</b>: The node pool fails to be scaled out due to unexpected reasons, affecting the node pool scaling status.</li> <li>• <b>Error</b>: A node pool error occurs. A common trigger is deletion failure.</li> </ul>
status	String	Current status of <b>Condition</b> . The options are as follows: <ul style="list-style-type: none"> <li>• "True"</li> <li>• "False"</li> </ul>
lastProbeTime	String	Time when the status was last checked
lastTransitTime	String	Time when the status was last changed
reason	String	Reason why the status was last changed
message	String	Detailed condition description

## Example Requests

None

## Example Responses

**Status code: 200**

The job for deleting the node pool is successfully delivered.

```
{
  "kind": "NodePool",
  "apiVersion": "v3",
  "metadata": {
    "name": "lc-it-nodepool-79796",
    "uid": "99addaa2-69eb-11ea-a592-0255ac1001bb"
  },
  "spec": {
    "type": "vm",
    "nodeTemplate": {
```

```

"flavor": "s6.large.2",
"az": "*****",
"os": "EulerOS 2.5",
"login": {
  "sshKey": "KeyPair-001"
},
"rootVolume": {
  "volumetype": "SAS",
  "size": 40
},
"dataVolumes": [ {
  "volumetype": "SAS",
  "size": 100,
  "extendParam": {
    "useType": "docker"
  }
} ],
"publicIP": {
  "eip": {
    "bandwidth": { }
  }
},
"nodeNicSpec": {
  "primaryNic": {
    "subnetId": "7e767d10-7548-4df5-ad72-aeac1d08bd8a"
  }
},
"billingMode": 0,
"extendParam": {
  "maxPods": 110
},
"k8sTags": {
  "cce.cloud.com/cce-nodepool": "lc-it-nodepool-79796"
},
"autoscaling": { },
"nodeManagement": { }
},
"status": {
  "phase": "Deleting",
  "jobId": "3281fa02-69ee-11ea-a592-0255ac1001bb"
}
}

```

## Status Codes

Status Code	Description
200	The job for deleting the node pool is successfully delivered.

## Error Codes

See [Error Codes](#).

## 4.4.6 Scaling a Node Pool

### Function

This API is used to scale in or out a node pool.

 NOTE

The URL for cluster management is in the format of **https://Endpoint/uri**, where **uri** specifies the resource path for API access.

## URI

POST /api/v3/projects/{project\_id}/clusters/{cluster\_id}/nodepools/{nodepool\_id}/operation/scale

**Table 4-716** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
nodepool_id	Yes	String	Node pool ID.

## Request Parameters

**Table 4-717** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

**Table 4-718** Request body parameters

Parameter	Mandatory	Type	Description
kind	Yes	String	API type. The value is fixed at <b>NodePool</b> .

Parameter	Mandatory	Type	Description
apiVersion	Yes	String	API version. The value is fixed at <b>v3</b> .
spec	Yes	<a href="#">ScaleNodePoolSpec</a> object	Parameters in the request for scaling a node pool

**Table 4-719** ScaleNodePoolSpec

Parameter	Mandatory	Type	Description
desiredNodeCount	Yes	Integer	Desired number of nodes in a node pool
scaleGroups	Yes	Array of strings	Node pool to be scaled in or out. Only one scaling group can be specified. To scale the default scaling group, set the value to <b>default</b> .
options	No	<a href="#">ScaleNodePoolOptions</a> object	Configurations of node pool scaling options

**Table 4-720** ScaleNodePoolOptions

Parameter	Mandatory	Type	Description
scalableChecking	No	String	Scale-out status check policy. The value can be <b>instant</b> (synchronous checks) or <b>async</b> (asynchronous checks). It defaults to <b>instant</b> .
billingConfigOverride	No	<a href="#">ScaleUpBillingConfigOverride</a> object	Overwrites the default billing mode configuration of a node during node pool scaling.

**Table 4-721** ScaleUpBillingConfigOverride

Parameter	Mandatory	Type	Description
billingMode	Yes	Integer	Billing mode of a node. The value can be <b>0</b> (pay-per-use).

## Response Parameters

None

## Example Requests

- Scale out nodes in the default scaling group of a node pool (pay-per-use).

```
{
  "kind": "NodePool",
  "apiVersion": "v3",
  "spec": {
    "desiredNodeCount": 1,
    "scaleGroups": [ "default" ]
  }
}
```

## Example Responses

None

## Status Codes

Status Code	Description
202	The node pool scaling is accepted. Nodes in the node pool will be added or deleted based on the expected number of nodes in the node pool after scaling.

## Error Codes

See [Error Codes](#).

## 4.4.7 Synchronizing Node Pools

### Function

This API is used to synchronize the configurations of existing nodes in a node pool.

#### NOTE

The URL for cluster management is in the format of **https://Endpoint/uri**, where **uri** specifies the resource path for API access.

### URI

POST /api/v3/projects/:projectid/clusters/{cluster\_id}/nodepools/{nodepool\_id}/operation/upgrade

**Table 4-722** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
nodepool_id	Yes	String	Node pool ID.

## Request Parameters

**Table 4-723** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

**Table 4-724** Request body parameters

Parameter	Mandatory	Type	Description
spec	Yes	<a href="#">UpgradeNodePoolSpec</a> object	Parameters in the request for synchronizing a node pool

**Table 4-725** UpgradeNodePoolSpec

Parameter	Mandatory	Type	Description
maxUnavailable	Yes	Integer	Maximum number of nodes to be synchronized in each batch. The value ranges from 1 to 20. It specifies the maximum number of nodes that can be unavailable during node upgrade. Nodes will be unavailable during synchronization by resetting the nodes. Properly configure this parameter to prevent pod scheduling failures caused by too many unavailable nodes in the cluster.
nodeIDs	No	Array of strings	List of node IDs selected in the node pool to be synchronized. For details about how to obtain the IDs, see <a href="#">How to Obtain Parameters in the API URI</a> .
nodePoolID	Yes	String	Node pool ID. For details about how to obtain the ID, see <a href="#">How to Obtain Parameters in the API URI</a> .
nodeTemplate	No	<a href="#">UpgradeNodePoolSpecNodeTemplate</a> object	Parameters for synchronizing a node pool template

**Table 4-726** UpgradeNodePoolSpecNodeTemplate

Parameter	Mandatory	Type	Description
lifecycle	Yes	<a href="#">NodeLifecycleConfig</a> object	Customized node lifecycle configurations
login	Yes	<a href="#">Login</a> object	Node login mode. You can use both key pair and password login simultaneously.  <b>NOTE</b> This parameter cannot be modified when a node pool is being updated.

**Table 4-727** NodeLifecycleConfig

Parameter	Mandatory	Type	Description
preInstall	No	String	Pre-installation scripts. The total number of characters after transcoding cannot exceed 10240. <b>NOTE</b> The input value must be encoded using Base64. The method is as follows: <code>echo -n "Content to be encoded"   base64</code>
postInstall	No	String	Post-installation scripts. The total number of characters after transcoding cannot exceed 10240. <b>NOTE</b> The input value must be encoded using Base64. The method is as follows: <code>echo -n "Content to be encoded"   base64</code>

**Table 4-728** Login

Parameter	Mandatory	Type	Description
sshKey	No	String	Name of the key pair used for login.
userPassword	No	<b>UserPassword</b> object	Password used for node login.

**Table 4-729** UserPassword

Parameter	Mandatory	Type	Description
username	No	String	Login account. The default value is <b>root</b> .



Parameter	Mandatory	Type	Description
password	Yes	String	<p>Login password. If a username and password are used to create a node, this field is shielded in the response body. A password must meet the following complexity requirements:</p> <ul style="list-style-type: none"> <li>• Contains 8 to 26 characters.</li> <li>• Contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters !@\$%^_-=+[{ }];,./?</li> <li>• Cannot contain the username or the username spelled backwards. The <b>password</b> field must be salted during node creation. For details, see <a href="#">Adding a Salt in the password Field When Creating a Node</a>.</li> </ul>

## Response Parameters

None

## Example Requests

- Synchronize an on-premises node pool.

```
{
  "spec": {
    "maxUnavailable": 1,
    "nodeIDs": [ "aaaaaaaa-aaaa-aaaa-aaaaaaaaaaaa", "bbbbbbb-bbbb-bbbb-bbbb-
    bbbbbbbbbb" ],
    "nodePoolID": "ccccccc-cccc-cccc-cccccccccccc"
  }
}
```

- Synchronize the default node pool.

```
{
  "spec": {
    "maxUnavailable": 1,
    "nodeIDs": [ "aaaaaaaa-aaaa-aaaa-aaaaaaaaaaaa", "bbbbbbb-bbbb-bbbb-bbbb-
    bbbbbbbbbb" ],
    "nodePoolID": "DefaultPool",
    "nodeTemplate": {
      "lifeCycle": {
        "preInstall": "ZWNobyAx",
        "postInstall": "ZWNobyAx"
      },
      "login": {
        "userPassword": {
```

```

    "password" : "xxxxx",
    "username" : "root"
  }
}
}
}
}

```

## Example Responses

None

## Status Codes

Status Code	Description
200	The node pool synchronization has been accepted.

## Error Codes

See [Error Codes](#).

# 4.5 Storage Management

## 4.5.1 Creating a PVC (to be discarded)

### Function

This API is used to create a PersistentVolumeClaim (PVC) in a specified namespace. A PVC is a request for PersistentVolume (PV). The PV carries details of real cloud storage such as EVS disks, SFS file systems, and OBS buckets. This API is to be deprecated. Use the corresponding Kubernetes PVC API instead.

#### NOTE

The URL for storage management is in the format of **https://{clusterid}.Endpoint/uri**. In the URL, *{clusterid}* indicates the cluster ID, and *uri* indicates the resource path, that is, the path for API access. If *https://Endpoint/uri* is used, the **X-Cluster-ID** parameter in the request header must be specified.

### URI

POST /api/v1/namespaces/{namespace}/cloudpersistentvolumeclaims

**Table 4-730** Path Parameters

Parameter	Mandatory	Type	Description
namespace	Yes	String	<p>Namespace where the PVC is located.</p> <p>Constraints:</p> <ul style="list-style-type: none"> <li>• Custom namespace. You must create a namespace in the cluster before using it.</li> <li>• The namespace provided by the system is namespace <b>default</b>.</li> <li>• Namespaces <b>kube-system</b> and <b>kube-public</b> cannot be used.</li> </ul>

## Request Parameters

**Table 4-731** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	<p>Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a>.</p>
X-Cluster-ID	No	String	<p>Cluster ID. This field is mandatory when the URL format <b>https://Endpoint/uri</b> is used. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a>.</p>

**Table 4-732** Request body parameters

Parameter	Mandatory	Type	Description
apiVersion	Yes	String	API version. The value is fixed at <b>v1</b> .

Parameter	Mandatory	Type	Description
kind	Yes	String	API type. The value is fixed at <b>PersistentVolumeClaim</b> .
metadata	Yes	<b>PersistentVolumeClaimMetadata</b> object	Cluster object metadata, which is a collection of attributes.
spec	Yes	<b>PersistentVolumeClaimSpec</b> object	Detailed description of the cluster object. CCE creates or updates objects by defining or updating <b>spec</b> .
status	No	<b>PersistentVolumeClaimStatus</b> object	PVC status. The request to create a PVC does not need to carry this parameter.

**Table 4-733** PersistentVolumeClaimMetadata

Parameter	Mandatory	Type	Description
name	Yes	String	PVC name. Enter 1 to 253 characters starting and ending with a letter or digit. Only lowercase letters, digits, hyphens (-), and periods (.) are allowed. PVC names must be unique in a namespace.
labels	No	String	PVC labels, in the format of key-value pairs. <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain containing a maximum of 253 characters (such as <b>example.com/my-key</b>) can be prefixed to a key.</li> <li>• <b>Value:</b> The value can be left blank or contain 1 to 63 characters that start and end with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed.</li> </ul>

**Table 4-734** PersistentVolumeClaimSpec

Parameter	Mandatory	Type	Description
volumeID	Yes	String	<p>ID of an existing storage volume.</p> <ul style="list-style-type: none"> <li>• If an SFS, EVS, or SFS Turbo volume is used, set this parameter to the ID of the volume.</li> <li>• If an OBS bucket is used, set this parameter to the OBS bucket name.</li> </ul>
storageType	Yes	String	<p>Cloud storage class. This parameter is used together with <b>volumeID</b>. That is, <b>volumeID</b> and <b>storageType</b> must be configured at the same time.</p> <ul style="list-style-type: none"> <li>• <b>bs</b>: EVS</li> <li>• <b>nfs</b>: SFS</li> <li>• <b>obs</b>: OBS</li> <li>• <b>efs</b>: SFS Turbo</li> </ul>
accessModes	Yes	Array of strings	<p>Access mode of the volume. Only the first value in all selected options is valid.</p> <ul style="list-style-type: none"> <li>• <b>ReadWriteOnce</b>: The volume can be mounted as read-write by a single node.</li> </ul> <p><b>NOTE</b> This function is supported only when the cluster version is v1.13.10 and the storage-driver version is 1.0.19.</p> <ul style="list-style-type: none"> <li>• <b>ReadOnlyMany</b> (default): The volume can be mounted as read-only by many nodes.</li> <li>• <b>ReadWriteMany</b>: The volume can be mounted as read-write by many nodes.</li> </ul>
storageClassName	No	String	Storage class name of the PVC.
volumeName	No	String	Name of the PV bound to the PVC.

Parameter	Mandatory	Type	Description
resources	No	<a href="#">ResourceRequirements</a> object	Resource requests and limits.
volumeMode	No	String	PV type specified by the PVC.

**Table 4-735** ResourceRequirements

Parameter	Mandatory	Type	Description
limits	No	Map<String,String>	Resource limits. This parameter is invalid during creation.
requests	No	Map<String,String>	Resource requests. This parameter is invalid during creation.

**Table 4-736** PersistentVolumeClaimStatus

Parameter	Mandatory	Type	Description
accessModes	No	Array of strings	Access mode of the PV.
capacity	No	String	Actual resources and capacity of the PV.
phase	No	String	Current status of the PVC.

## Response Parameters

Status code: 201

**Table 4-737** Response body parameters

Parameter	Type	Description
apiVersion	String	API version. The value is fixed at <b>v1</b> .
kind	String	API type. The value is fixed at <b>PersistentVolumeClaim</b> .
metadata	<a href="#">PersistentVolumeClaimMetadata</a> object	Cluster object metadata, which is a collection of attributes.

Parameter	Type	Description
spec	<a href="#">PersistentVolumeClaimSpec</a> object	Detailed description of the cluster object. CCE creates or updates objects by defining or updating <b>spec</b> .
status	<a href="#">PersistentVolumeClaimStatus</a> object	PVC status. The request to create a PVC does not need to carry this parameter.

**Table 4-738** PersistentVolumeClaimMetadata

Parameter	Type	Description
name	String	PVC name. Enter 1 to 253 characters starting and ending with a letter or digit. Only lowercase letters, digits, hyphens (-), and periods (.) are allowed. PVC names must be unique in a namespace.
labels	String	PVC labels, in the format of key-value pairs. <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain containing a maximum of 253 characters (such as <b>example.com/my-key</b>) can be prefixed to a key.</li> <li>• <b>Value:</b> The value can be left blank or contain 1 to 63 characters that start and end with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed.</li> </ul>

**Table 4-739** PersistentVolumeClaimSpec

Parameter	Type	Description
volumeID	String	ID of an existing storage volume. <ul style="list-style-type: none"> <li>• If an SFS, EVS, or SFS Turbo volume is used, set this parameter to the ID of the volume.</li> <li>• If an OBS bucket is used, set this parameter to the OBS bucket name.</li> </ul>

Parameter	Type	Description
storageType	String	Cloud storage class. This parameter is used together with <b>volumeID</b> . That is, <b>volumeID</b> and <b>storageType</b> must be configured at the same time. <ul style="list-style-type: none"> <li>• <b>bs</b>: EVS</li> <li>• <b>nfs</b>: SFS</li> <li>• <b>obs</b>: OBS</li> <li>• <b>efs</b>: SFS Turbo</li> </ul>
accessModes	Array of strings	Access mode of the volume. Only the first value in all selected options is valid. <ul style="list-style-type: none"> <li>• <b>ReadWriteOnce</b>: The volume can be mounted as read-write by a single node.</li> </ul> <p><b>NOTE</b> This function is supported only when the cluster version is v1.13.10 and the storage-driver version is 1.0.19.</p> <ul style="list-style-type: none"> <li>• <b>ReadOnlyMany</b> (default): The volume can be mounted as read-only by many nodes.</li> <li>• <b>ReadWriteMany</b>: The volume can be mounted as read-write by many nodes.</li> </ul>
storageClassName	String	Storage class name of the PVC.
volumeName	String	Name of the PV bound to the PVC.
resources	<b>ResourceRequirements</b> object	Resource requests and limits.
volumeMode	String	PV type specified by the PVC.

**Table 4-740** ResourceRequirements

Parameter	Type	Description
limits	Map<String,String>	Resource limits. This parameter is invalid during creation.
requests	Map<String,String>	Resource requests. This parameter is invalid during creation.



**Table 4-741** PersistentVolumeClaimStatus

Parameter	Type	Description
accessModes	Array of strings	Access mode of the PV.
capacity	String	Actual resources and capacity of the PV.
phase	String	Current status of the PVC.

## Example Requests

Specifying an EVS volume ID and creating a PVC

POST /api/v1/namespaces/default/cloudpersistentvolumeclaims

```
{
  "apiVersion": "v1",
  "kind": "PersistentVolumeClaim",
  "metadata": {
    "name": "csms-dev-create",
    "namespace": "default"
  },
  "spec": {
    "volumeID": "86b29e16-23db-11e7-9c83-fa163ec08232",
    "storageType": "bs",
    "accessModes": [ "ReadWriteMany" ]
  }
}
```

## Example Responses

**Status code: 201**

The job for creating a PVC is successfully delivered.

```
{
  "kind": "PersistentVolumeClaim",
  "apiVersion": "v1",
  "metadata": {
    "name": " csms-dev-create ",
    "namespace": "default",
    "selfLink": "/api/v1/namespaces/default/persistentvolumeclaims/db-mysql-0",
    "uid": "86b29e16-23db-11e7-9c83-fa163ec08232",
    "resourceVersion": "1793115",
    "creationTimestamp": "2017-04-18T02:05:42Z"
  },
  "spec": {
    "volumeName": "csms-dev-create ",
    "accessModes": [ "ReadWriteMany" ],
    "resources": {
      "requests": {
        "storage": "1Gi"
      }
    }
  },
  "status": {
    "phase": "Pending",
    "accessModes": [ "ReadWriteMany" ]
  }
}
```

## Status Codes

Status Code	Description
201	The job for creating a PVC is successfully delivered.

## Error Codes

See [Error Codes](#).

## 4.5.2 Deleting a PVC (to be discarded)

### Function

This API is used to delete a PVC from a specified namespace. Using this API, you can determine whether to retain the cloud storage volume associated with the PVC. This API has been deprecated. Use the corresponding Kubernetes PV API instead.

#### NOTE

The URL for storage management is in the format of **https://{clusterid}.Endpoint/uri**. In the URL, *{clusterid}* indicates the cluster ID, and *uri* indicates the resource path, that is, the path for API access. If *https://Endpoint/uri* is used, the **X-Cluster-ID** parameter in the request header must be specified.

### URI

DELETE /api/v1/namespaces/{namespace}/cloudpersistentvolumeclaims/{name}

**Table 4-742** Path Parameters

Parameter	Mandatory	Type	Description
name	Yes	String	Name of the PVC to be deleted.
namespace	Yes	String	Namespace where the PVC is located.

**Table 4-743** Query Parameters

Parameter	Mandatory	Type	Description
deleteVolume	No	String	Whether to delete associated backend cloud storage volume after the PVC is deleted. The value can be <b>true</b> (delete) or <b>false</b> (not delete). The value defaults to <b>false</b> .
storageType	No	String	Cloud storage type. This parameter is used together with deleteVolume. That is, <b>deleteVolume</b> and <b>storageType</b> must be configured at the same time. <ul style="list-style-type: none"> <li>• <b>bs</b>: EVS</li> <li>• <b>nfs</b>: SFS</li> <li>• <b>obs</b>: OBS</li> <li>• <b>efs</b>: SFS Turbo</li> </ul>

## Request Parameters

**Table 4-744** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .
X-Cluster-ID	No	String	Cluster ID. This field is mandatory when the URL format <b>https://Endpoint/uri</b> is used. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

## Response Parameters

Status code: 200

**Table 4-745** Response body parameters

Parameter	Type	Description
apiVersion	String	API version. The value is fixed at <b>v1</b> .
kind	String	API type. The value is fixed at <b>PersistentVolumeClaim</b> .
metadata	<b>PersistentVolumeClaimMetadata</b> object	Cluster object metadata, which is a collection of attributes.
spec	<b>PersistentVolumeClaimSpec</b> object	Detailed description of the cluster object. CCE creates or updates objects by defining or updating <b>spec</b> .
status	<b>PersistentVolumeClaimStatus</b> object	PVC status. The request to create a PVC does not need to carry this parameter.

**Table 4-746** PersistentVolumeClaimMetadata

Parameter	Type	Description
name	String	PVC name. Enter 1 to 253 characters starting and ending with a letter or digit. Only lowercase letters, digits, hyphens (-), and periods (.) are allowed. PVC names must be unique in a namespace.
labels	String	PVC labels, in the format of key-value pairs. <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain containing a maximum of 253 characters (such as <b>example.com/my-key</b>) can be prefixed to a key.</li> <li>• <b>Value:</b> The value can be left blank or contain 1 to 63 characters that start and end with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed.</li> </ul>

**Table 4-747** PersistentVolumeClaimSpec

Parameter	Type	Description
volumeID	String	ID of an existing storage volume. <ul style="list-style-type: none"> <li>If an SFS, EVS, or SFS Turbo volume is used, set this parameter to the ID of the volume.</li> <li>If an OBS bucket is used, set this parameter to the OBS bucket name.</li> </ul>
storageType	String	Cloud storage class. This parameter is used together with <b>volumeID</b> . That is, <b>volumeID</b> and <b>storageType</b> must be configured at the same time. <ul style="list-style-type: none"> <li><b>bs</b>: EVS</li> <li><b>nfs</b>: SFS</li> <li><b>obs</b>: OBS</li> <li><b>efs</b>: SFS Turbo</li> </ul>
accessModes	Array of strings	Access mode of the volume. Only the first value in all selected options is valid. <ul style="list-style-type: none"> <li><b>ReadWriteOnce</b>: The volume can be mounted as read-write by a single node.</li> </ul> <p><b>NOTE</b> This function is supported only when the cluster version is v1.13.10 and the storage-driver version is 1.0.19.</p> <ul style="list-style-type: none"> <li><b>ReadOnlyMany</b> (default): The volume can be mounted as read-only by many nodes.</li> <li><b>ReadWriteMany</b>: The volume can be mounted as read-write by many nodes.</li> </ul>
storageClassName	String	Storage class name of the PVC.
volumeName	String	Name of the PV bound to the PVC.
resources	<a href="#">ResourceRequirements</a> object	Resource requests and limits.
volumeMode	String	PV type specified by the PVC.

**Table 4-748** ResourceRequirements

Parameter	Type	Description
limits	Map<String,String>	Resource limits. This parameter is invalid during creation.

Parameter	Type	Description
requests	Map<String,String>	Resource requests. This parameter is invalid during creation.

**Table 4-749** PersistentVolumeClaimStatus

Parameter	Type	Description
accessModes	Array of strings	Access mode of the PV.
capacity	String	Actual resources and capacity of the PV.
phase	String	Current status of the PVC.

## Example Requests

None

## Example Responses

None

## Status Codes

Status Code	Description
200	The job for deleting a PVC is successfully delivered.

## Error Codes

See [Error Codes](#).

# 4.6 Add-on Management

## 4.6.1 Installing an Add-on Instance

### Function

This API is used to install an add-on instance by using an add-on template.

### URI

POST /api/v3/addons

## Request Parameters

**Table 4-750** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

**Table 4-751** Request body parameters

Parameter	Mandatory	Type	Description
kind	Yes	String	API type. The value is fixed at <b>Addon</b> and cannot be changed. Any user-defined value is invalid.
apiVersion	Yes	String	API version. The value is fixed at <b>v3</b> and cannot be changed. Any user-defined value is invalid.
metadata	Yes	<a href="#">AddonMetadata</a> object	Basic information about the object. Metadata is a collection of attributes.
spec	Yes	<a href="#">InstanceRequestSpec</a> object	Detailed description of add-on installation or upgrade.

**Table 4-752** AddonMetadata

Parameter	Mandatory	Type	Description
uid	No	String	Unique ID
name	No	String	Add-on name
alias	No	String	Add-on alias
labels	No	Map<String,String>	Add-on labels in key-value pairs. This is a reserved field and does not take effect.

Parameter	Mandatory	Type	Description
annotations	No	Map<String,String>	Add-on annotations in the format of key-value pairs. <ul style="list-style-type: none"> <li>For add-on installation, the value is fixed at <b>{"addon.install/type":"install"}</b>.</li> <li>For add-on upgrade, the value is fixed at <b>{"addon.upgrade/type":"upgrade"}</b>.</li> </ul>
updateTimestamp	No	String	Updated at
creationTimestamp	No	String	Created at

**Table 4-753** InstanceRequestSpec

Parameter	Mandatory	Type	Description
version	No	String	Version of the add-on to install or upgrade, for example, <b>1.0.0</b> . <ul style="list-style-type: none"> <li>Installation: This parameter is optional. If not specified, the latest version supported by the cluster is used.</li> <li>Upgrade: This parameter is mandatory. The version number must be specified.</li> </ul>
clusterID	Yes	String	Cluster ID.
values	Yes	Map<String,Object>	Add-on template installation parameters (varying depending on the add-on). During the add-on upgrade, you need to specify all the installation parameters. If the parameters are not specified, the default values in the add-on template are used. The current add-on installation parameters can be obtained through the API for querying add-on instances.



Parameter	Mandatory	Type	Description
addonTemplateName	Yes	String	Name of the add-on template to be installed, for example, <b>coredns</b> .

## Response Parameters

Status code: 201

**Table 4-754** Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at <b>Addon</b> and cannot be changed.
apiVersion	String	API version. The value is fixed at <b>v3</b> and cannot be changed.
metadata	<a href="#">AddonMetadata</a> object	Basic information about the object. Metadata is a collection of attributes.
spec	<a href="#">InstanceSpec</a> object	Detailed description of the add-on instance.
status	<a href="#">AddonInstanceStatus</a> object	Add-on instance status.

**Table 4-755** AddonMetadata

Parameter	Type	Description
uid	String	Unique ID
name	String	Add-on name
alias	String	Add-on alias
labels	Map<String,String>	Add-on labels in key-value pairs. This is a reserved field and does not take effect.
annotations	Map<String,String>	Add-on annotations in the format of key-value pairs. <ul style="list-style-type: none"> <li>For add-on installation, the value is fixed at <b>{"addon.install/type":"install"}</b>.</li> <li>For add-on upgrade, the value is fixed at <b>{"addon.upgrade/type":"upgrade"}</b>.</li> </ul>

Parameter	Type	Description
updateTimestamp	String	Updated at
creationTimestamp	String	Created at

**Table 4-756** InstanceSpec

Parameter	Type	Description
clusterID	String	Cluster ID.
version	String	Add-on template version, for example, <b>1.0.0</b> .
addonTemplateName	String	Add-on template name, for example, <b>coredns</b> .
addonTemplateType	String	Add-on template type.
addonTemplateLogo	String	URL for obtaining the add-on template logo.
addonTemplateLabels	Array of strings	Labels of the add-on template.
description	String	Add-on template description.
values	Map<String, Object>	Add-on template installation parameters (varying depending on the add-on). Set the parameters accordingly.

**Table 4-757** AddonInstanceStatus

Parameter	Type	Description
status	String	<p>Add-on instance status. Options:</p> <ul style="list-style-type: none"> <li>● <b>running</b>: All of the add-on instances are running. This specifies that the add-on runs properly.</li> <li>● <b>abnormal</b>: The add-on instances are abnormal and the add-on cannot be used. You can click the add-on name to view exceptions.</li> <li>● <b>installing</b>: The add-on is being installed.</li> <li>● <b>installFailed</b>: Installing the add-on failed. In this case, uninstall the add-on and then reinstall it.</li> <li>● <b>upgrading</b>: The add-on is being upgraded.</li> <li>● <b>upgradeFailed</b>: Upgrading the add-on failed. In this case, upgrade the add-on again or uninstall the add-on and reinstall it.</li> <li>● <b>deleting</b>: The add-on is being deleted.</li> <li>● <b>deleteFailed</b>: Deleting the add-on failed. In this case, uninstall the add-on again.</li> <li>● <b>deleteSuccess</b>: Deleting the add-on succeeded.</li> <li>● <b>available</b>: Only some instances of the add-on are running. This specifies that some functions of the add-on are available.</li> <li>● <b>rollbacking</b>: The add-on is being rolled back.</li> <li>● <b>rollbackFailed</b>: Rolling back the add-on failed. In this case, roll back the add-on again or uninstall the add-on and reinstall it.</li> <li>● <b>unknown</b>: The add-on chart instance does not exist.</li> </ul>
Reason	String	Cause of the add-on installation failure.
message	String	Installation error details.
targetVersions	Array of strings	Versions to which the current add-on version can be upgraded.
currentVersion	<b>Versions</b> object	Information about the current add-on version.
isRollbackable	Boolean	Whether the add-on version can be rolled back to the source version.

Parameter	Type	Description
previousVersion	String	The add-on version before upgrade or rollback

**Table 4-758** Versions

Parameter	Type	Description
version	String	Add-on version.
input	Object	Add-on installation parameters.
stable	Boolean	Whether the add-on version is a stable release.
translate	Object	Translation information used by the GUI.
supportVersions	Array of <a href="#">SupportVersions</a> objects	Cluster versions that support the add-on.
creationTimestamp	String	Creation time.
updateTimestamp	String	Update time.

**Table 4-759** SupportVersions

Parameter	Type	Description
clusterType	String	Cluster type that supports the add-on.
clusterVersion	Array of strings	Cluster versions that support the add-on. The value is a regular expression.

## Example Requests

Install the coredns add-on of version 1.17.15. The add-on specification is 2500 QPS, and the number of add-on instances is 2.

```
{
  "kind": "Addon",
  "apiVersion": "v3",
  "metadata": {
    "annotations": {
      "addon.install/type": "install"
    }
  },
  "spec": {
    "clusterID": "1b2ec02d-a3b2-11ec-b0d0-0255ac100099",
    "version": "1.17.15",
    "addonTemplateName": "coredns",
    "values": {
```

```

"basic" : {
  "cluster_ip" : "10.247.3.10",
  "image_version" : "1.17.15",
  "platform" : "linux-amd64",
  "swr_addr" : "<Replace_SWR_address>",
  "swr_user" : "hwofficial",
  "rbac_enabled" : true
},
"flavor" : {
  "name" : 2500,
  "replicas" : 2,
  "resources" : [ {
    "limitsCpu" : "500m",
    "limitsMem" : "512Mi",
    "name" : "coredns",
    "requestsCpu" : "500m",
    "requestsMem" : "512Mi"
  } ]
},
"custom" : {
  "stub_domains" : { },
  "upstream_nameservers" : [ ],
  "cluster_id" : "1b2ec02d-a3b2-11ec-b0d0-0255ac100099",
  "tenant_id" : "0504201b6c80256b2f08c0099f0c8fe4"
}
}
}
}

```

## Example Responses

**Status code: 201**

OK

```

{
  "kind" : "Addon",
  "apiVersion" : "v3",
  "metadata" : {
    "uid" : "b748aaea-a984-11ec-987b-0255ac1000bc",
    "name" : "coredns",
    "alias" : "coredns",
    "creationTimestamp" : "2022-03-22T02:06:41Z",
    "updateTimestamp" : "2022-03-22T02:06:41Z"
  },
  "spec" : {
    "clusterID" : "1b2ec02d-a3b2-11ec-b0d0-0255ac100099",
    "version" : "1.17.15",
    "addonTemplateName" : "coredns",
    "addonTemplateType" : "helm",
    "addonTemplateLogo" : "",
    "addonTemplateLabels" : [ "ServiceDiscovery" ],
    "description" : "CoreDNS is a DNS server that chains plugins and provides Kubernetes DNS Services",
    "values" : {
      "basic" : {
        "cluster_ip" : "10.247.3.10",
        "image_version" : "1.17.15",
        "platform" : "linux-amd64",
        "rbac_enabled" : true,
        "swr_addr" : "",
        "swr_user" : "hwofficial"
      },
      "custom" : {
        "cluster_id" : "1b2ec02d-a3b2-11ec-b0d0-0255ac100099",
        "stub_domains" : { },
        "tenant_id" : "0504201b6c80256b2f08c0099f0c8fe4",
        "upstream_nameservers" : [ ]
      },
      "flavor" : {

```

```

"name" : 2500,
"replicas" : 2,
"resources" : [ {
  "limitsCpu" : "500m",
  "limitsMem" : "512Mi",
  "name" : "coredns",
  "requestsCpu" : "500m",
  "requestsMem" : "512Mi"
} ]
}
},
"status" : {
  "status" : "installing",
  "Reason" : "",
  "message" : "",
  "targetVersions" : null,
  "currentVersion" : {
    "version" : "1.17.15",
    "input" : {
      "basic" : {
        "cluster_ip" : "10.247.3.10",
        "image_version" : "1.17.15",
        "platform" : "linux-amd64",
        "swr_addr" : "",
        "swr_user" : "hwofficial"
      },
      "parameters" : {
        "custom" : {
          "stub_domains" : "",
          "upstream_nameservers" : ""
        },
        "flavor1" : {
          "name" : 2500,
          "replicas" : 2,
          "resources" : [ {
            "limitsCpu" : "500m",
            "limitsMem" : "512Mi",
            "name" : "coredns",
            "requestsCpu" : "500m",
            "requestsMem" : "512Mi"
          } ]
        },
        "flavor2" : {
          "name" : 5000,
          "replicas" : 2,
          "resources" : [ {
            "limitsCpu" : "1000m",
            "limitsMem" : "1024Mi",
            "name" : "coredns",
            "requestsCpu" : "1000m",
            "requestsMem" : "1024Mi"
          } ]
        },
        "flavor3" : {
          "name" : 10000,
          "replicas" : 2,
          "resources" : [ {
            "limitsCpu" : "2000m",
            "limitsMem" : "2048Mi",
            "name" : "coredns",
            "requestsCpu" : "2000m",
            "requestsMem" : "2048Mi"
          } ]
        },
        "flavor4" : {
          "name" : 20000,
          "replicas" : 4,
          "resources" : [ {

```

```

        "limitsCpu" : "2000m",
        "limitsMem" : "2048Mi",
        "name" : "coredns",
        "requestsCpu" : "2000m",
        "requestsMem" : "2048Mi"
    } ]
}
},
"stable" : true,
"translate" : {
    "en_US" : {
        "addon" : {
            "changeLog" : "Supported CCE clusters of v1.21.",
            "description" : "CoreDNS is a DNS server that chains plugins and provides Kubernetes DNS Services"
        },
        "description" : {
            "Parameters.custom.stub_domains" : "The target nameserver may itself be a Kubernetes service. For instance, you can run your own copy of dnsmasq to export custom DNS names into the ClusterDNS namespace, a JSON map using a DNS suffix key (e.g. \"acme.local\") and a value consisting of a JSON array of DNS IPs.",
            "Parameters.custom.upstream_nameservers" : "If specified, then the values specified replace the nameservers taken by default from the node's /etc/resolv.conf. Limits:a maximum of three upstream nameservers can be specified, A JSON array of DNS IPs.",
            "Parameters.flavor1.description" : "Concurrent domain name resolution ability - External domain name: 2500 qps, Internal domain name: 10000 qps",
            "Parameters.flavor1.name" : 2500,
            "Parameters.flavor2.description" : "Concurrent domain name resolution ability - External domain name: 5000 qps, Internal domain name: 20000 qps",
            "Parameters.flavor2.name" : 5000,
            "Parameters.flavor3.description" : "Concurrent domain name resolution ability - External domain name: 10000 qps, Internal domain name: 40000 qps",
            "Parameters.flavor3.name" : 10000,
            "Parameters.flavor4.description" : "Concurrent domain name resolution ability - External domain name: 20000 qps, Internal domain name: 80000 qps",
            "Parameters.flavor4.name" : 20000
        },
        "key" : {
            "Parameters.custom.stub_domains" : "stub domain",
            "Parameters.custom.upstream_nameservers" : "upstream nameservers"
        }
    },
    "fr_FR" : {
        "addon" : {
            "changeLog" : "Prise en charge du cluster 1.21.",
            "description" : "Un serveur DNS qui enchaîne les plug-ins et fournit des services DNS Kubernetes."
        },
        "description" : {
            "Parameters.custom.stub_domains" : "Le serveur de noms cible peut lui-même être un service Kubernetes. Par exemple, vous pouvez exécuter votre propre copie de dnsmasq pour exporter des noms DNS personnalisés dans l'espace de noms ClusterDNS, une carte JSON à l'aide d'une clé de suffixe DNS (par exemple, «acme.local») et une valeur constituée d'un tableau JSON d'adresses IP DNS.",
            "Parameters.custom.upstream_nameservers" : "Si spécifié, les valeurs spécifiées remplacent les serveurs de noms pris par défaut dans le fichier /etc/resolv.conf du nœud. Limites: un maximum de trois serveurs de noms en amont peuvent être spécifiés, un tableau JSON d'adresses IP DNS.",
            "Parameters.flavor1.description" : "Capacité de résolution de nom de domaine simultanée - Nom de domaine externe: 2500 qps, Nom de domaine interne: 10000 qp",
            "Parameters.flavor1.name" : 2500,
            "Parameters.flavor2.description" : "Capacité de résolution de nom de domaine simultanée - Nom de domaine externe: 5000 qps, Nom de domaine interne: 20000 qp",
            "Parameters.flavor2.name" : 5000,
            "Parameters.flavor3.description" : "Capacité de résolution de nom de domaine simultanée - Nom de domaine externe: 10000 qps, Nom de domaine interne: 40000 qp",
            "Parameters.flavor3.name" : 10000,
            "Parameters.flavor4.description" : "Capacité de résolution de nom de domaine simultanée - Nom de domaine externe: 20000 qps, Nom de domaine interne: 80000 qp",
            "Parameters.flavor4.name" : 20000
        },
        "key" : {

```

```

    "Parameters.custom.stub_domains" : "domaine stub",
    "Parameters.custom.upstream_nameservers" : "serveurs de noms en amont"
  }
},
"zh_CN" : {
  "addon" : {
    "changeLog" : "",
    "description" : ""
  },
  "description" : {
    "Parameters.custom.stub_domains" : "",
    "Parameters.custom.upstream_nameservers" : "",
    "Parameters.flavor1.description" : "",
    "Parameters.flavor1.name" : 2500,
    "Parameters.flavor2.description" : "",
    "Parameters.flavor2.name" : 5000,
    "Parameters.flavor3.description" : "",
    "Parameters.flavor3.name" : 10000,
    "Parameters.flavor4.description" : "",
    "Parameters.flavor4.name" : 20000
  },
  "key" : {
    "Parameters.custom.stub_domains" : "",
    "Parameters.custom.upstream_nameservers" : ""
  }
}
},
"supportVersions" : null,
"creationTimestamp" : "2021-12-14T13:43:15Z",
"updateTimestamp" : "2022-01-11T14:32:10Z"
}
}
}

```

## Status Codes

Status Code	Description
201	OK

## Error Codes

See [Error Codes](#).

### 4.6.2 Listing Add-on Templates

#### Function

This API is used to query add-on templates.

#### URI

GET /api/v3/addontemplates



**Table 4-760** Query Parameters

Parameter	Mandatory	Type	Description
addon_template_name	No	String	Add-on name or alias for search. If this parameter is not specified, all the add-ons will be searched.

## Request Parameters

**Table 4-761** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

## Response Parameters

Status code: 200

**Table 4-762** Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at <b>Addon</b> and cannot be changed.
apiVersion	String	API version. The value is fixed at <b>v3</b> and cannot be changed.
items	Array of <a href="#">AddonTemplate</a> objects	List of add-on templates.

**Table 4-763** AddonTemplate

Parameter	Type	Description
kind	String	API type. The value is fixed at <b>Addon</b> and cannot be changed.
apiVersion	String	API version. The value is fixed at <b>v3</b> and cannot be changed.
metadata	<a href="#">AddonMetadata</a> object	Basic information about the object. Metadata is a collection of attributes.
spec	<a href="#">TemplateSpec</a> object	Detailed description of the add-on template.

**Table 4-764** AddonMetadata

Parameter	Type	Description
uid	String	Unique ID
name	String	Add-on name
alias	String	Add-on alias
labels	Map<String,String>	Add-on labels in key-value pairs. This is a reserved field and does not take effect.
annotations	Map<String,String>	Add-on annotations in the format of key-value pairs. <ul style="list-style-type: none"> <li>For add-on installation, the value is fixed at <b>{"addon.install/type":"install"}</b>.</li> <li>For add-on upgrade, the value is fixed at <b>{"addon.upgrade/type":"upgrade"}</b>.</li> </ul>
updateTimestamp	String	Updated at
creationTimestamp	String	Created at

**Table 4-765** TemplateSpec

Parameter	Type	Description
type	String	Template type. The value is <b>helm</b> or <b>static</b> .
require	Boolean	Whether the add-on is mandatory.
labels	Array of strings	Group to which the template belongs.

Parameter	Type	Description
logoURL	String	URL of the logo image.
readmeURL	String	Add-on details and usage description.
description	String	Description
versions	Array of <b>Versions</b> objects	Template version details.

**Table 4-766** Versions

Parameter	Type	Description
version	String	Add-on version.
input	Object	Add-on installation parameters.
stable	Boolean	Whether the add-on version is a stable release.
translate	Object	Translation information used by the GUI.
supportVersions	Array of <b>SupportVersions</b> objects	Cluster versions that support the add-on.
creationTimestamp	String	Creation time.
updateTimestamp	String	Update time.

**Table 4-767** SupportVersions

Parameter	Type	Description
clusterType	String	Cluster type that supports the add-on.
clusterVersion	Array of strings	Cluster versions that support the add-on. The value is a regular expression.

## Example Requests

None

## Example Responses

**Status code: 200**

OK

```
{
  "kind": "Addon",
  "apiVersion": "v3",
  "items": [ {
    "kind": "Addon",
    "apiVersion": "v3",
    "metadata": {
      "uid": "coredns",
      "name": "coredns",
      "alias": "coredns",
      "creationTimestamp": "2018-11-04T16:15:56Z",
      "updateTimestamp": "2022-01-11T14:32:10Z"
    },
    "spec": {
      "type": "helm",
      "require": true,
      "labels": [ "ServiceDiscovery" ],
      "logoURL": "",
      "description": "CoreDNS is a DNS server that chains plugins and provides Kubernetes DNS Services",
      "versions": [ {
        "version": "1.13.6",
        "input": {
          "basic": {
            "cluster_ip": "10.247.3.10",
            "ipv6": false,
            "platform": "linux-amd64",
            "swr_addr": "100.125.16.65:20202",
            "swr_user": "hwofficial"
          },
          "parameters": {
            "custom": {
              "stub_domains": "",
              "upstream_nameservers": ""
            },
            "flavor1": {
              "name": 2500,
              "replicas": 2,
              "resources": [ {
                "limitsCpu": "500m",
                "limitsMem": "512Mi",
                "name": "coredns",
                "requestsCpu": "500m",
                "requestsMem": "512Mi"
              } ]
            },
            "flavor2": {
              "name": 5000,
              "replicas": 2,
              "resources": [ {
                "limitsCpu": "1000m",
                "limitsMem": "1024Mi",
                "name": "coredns",
                "requestsCpu": "1000m",
                "requestsMem": "1024Mi"
              } ]
            },
            "flavor3": {
              "name": 10000,
              "replicas": 2,
              "resources": [ {
                "limitsCpu": "2000m",
                "limitsMem": "2048Mi",
                "name": "coredns",
                "requestsCpu": "2000m",
                "requestsMem": "2048Mi"
              } ]
            },
            "flavor4": {
              "name": 20000,

```

```

    "replicas" : 4,
    "resources" : [ {
      "limitsCpu" : "2000m",
      "limitsMem" : "2048Mi",
      "name" : "coredns",
      "requestsCpu" : "2000m",
      "requestsMem" : "2048Mi"
    } ]
  }
},
"stable" : true,
"translate" : {
  "en_US" : {
    "addon" : {
      "changeLog" : "Support for clusters with new version",
      "description" : "CoreDNS is a DNS server that chains plugins and provides Kubernetes DNS
Services"
    },
    "description" : {
      "Parameters.custom.stub_domains" : "The target nameserver may itself be a Kubernetes service.
For instance, you can run your own copy of dnsmasq to export custom DNS names into the ClusterDNS
namespace, a JSON map using a DNS suffix key (e.g. \"acme.local\") and a value consisting of a JSON array
of DNS IPs.",
      "Parameters.custom.upstream_nameservers" : "If specified, then the values specified replace the
nameservers taken by default from the node's /etc/resolv.conf. Limits:a maximum of three upstream
nameservers can be specified, A JSON array of DNS IPs.",
      "Parameters.flavor1.description" : "Concurrent domain name resolution ability - External domain
name: 2500 qps, Internal domain name: 10000 qps",
      "Parameters.flavor1.name" : 2500,
      "Parameters.flavor2.description" : "Concurrent domain name resolution ability - External domain
name: 5000 qps, Internal domain name: 20000 qps",
      "Parameters.flavor2.name" : 5000,
      "Parameters.flavor3.description" : "Concurrent domain name resolution ability - External domain
name: 10000 qps, Internal domain name: 40000 qps",
      "Parameters.flavor3.name" : 10000,
      "Parameters.flavor4.description" : "Concurrent domain name resolution ability - External domain
name: 20000 qps, Internal domain name: 80000 qps",
      "Parameters.flavor4.name" : 20000
    },
    "key" : {
      "Parameters.custom.stub_domains" : "stub domain",
      "Parameters.custom.upstream_nameservers" : "upstream nameservers"
    }
  },
  "fr_FR" : {
    "addon" : {
      "changeLog" : "Prise en charge des clusters avec une nouvelle version",
      "description" : "Un serveur DNS qui enchaîne les plug-ins et fournit des services DNS Kubernetes."
    },
    "description" : {
      "Parameters.custom.stub_domains" : "Le serveur de noms cible peut lui-même être un service
Kubernetes. Par exemple, vous pouvez exécuter votre propre copie de dnsmasq pour exporter des noms
DNS personnalisés dans l'espace de noms ClusterDNS, une carte JSON à l'aide d'une clé de suffixe DNS (par
exemple, «acme.local») et une valeur constituée d'un tableau JSON d'adresses IP DNS.",
      "Parameters.custom.upstream_nameservers" : "Si spécifié, les valeurs spécifiées remplacent les
serveurs de noms pris par défaut dans le fichier /etc/resolv.conf du nœud. Limites: un maximum de trois
serveurs de noms en amont peuvent être spécifiés, un tableau JSON d'adresses IP DNS.",
      "Parameters.flavor1.description" : "Capacité de résolution de nom de domaine simultanée - Nom
de domaine externe: 2500 qps, Nom de domaine interne: 10000 qp",
      "Parameters.flavor1.name" : 2500,
      "Parameters.flavor2.description" : "Capacité de résolution de nom de domaine simultanée - Nom
de domaine externe: 5000 qps, Nom de domaine interne: 20000 qp",
      "Parameters.flavor2.name" : 5000,
      "Parameters.flavor3.description" : "Capacité de résolution de nom de domaine simultanée - Nom
de domaine externe: 10000 qps, Nom de domaine interne: 40000 qp",
      "Parameters.flavor3.name" : 10000,
      "Parameters.flavor4.description" : "Capacité de résolution de nom de domaine simultanée - Nom
de domaine externe: 20000 qps, Nom de domaine interne: 80000 qp",

```

```

    "Parameters.flavor4.name" : 20000
  },
  "key" : {
    "Parameters.custom.stub_domains" : "domaine stub",
    "Parameters.custom.upstream_nameservers" : "serveurs de noms en amont"
  }
},
"zh_CN" : {
  "addon" : {
    "changeLog" : "",
    "description" : ""
  },
  "description" : {
    "Parameters.custom.stub_domains" : "",
    "Parameters.custom.upstream_nameservers" : "",
    "Parameters.flavor1.description" : "",
    "Parameters.flavor1.name" : 2500,
    "Parameters.flavor2.description" : "",
    "Parameters.flavor2.name" : 5000,
    "Parameters.flavor3.description" : "",
    "Parameters.flavor3.name" : 10000,
    "Parameters.flavor4.description" : "",
    "Parameters.flavor4.name" : 20000
  },
  "key" : {
    "Parameters.custom.stub_domains" : "",
    "Parameters.custom.upstream_nameservers" : ""
  }
}
},
"supportVersions" : [ {
  "clusterType" : "VirtualMachine",
  "clusterVersion" : [ "v1.13.*" ]
}, {
  "clusterType" : "BareMetal",
  "clusterVersion" : [ "v1.13.*" ]
}, {
  "clusterType" : "ARM64",
  "clusterVersion" : [ "v1.13.*" ]
} ],
"creationTimestamp" : "2021-03-18T12:51:05Z",
"updateTimestamp" : "2021-03-18T12:51:05Z"
} ]
} ]
}

```

## Status Codes

Status Code	Description
200	OK

## Error Codes

See [Error Codes](#).

## 4.6.3 Updating an Add-on Instance

### Function

This API is used to update an add-on instance.

### URI

PUT /api/v3/addons/{id}

**Table 4-768** Path Parameters

Parameter	Mandatory	Type	Description
id	Yes	String	Add-on instance ID.

### Request Parameters

**Table 4-769** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

**Table 4-770** Request body parameters

Parameter	Mandatory	Type	Description
kind	Yes	String	API type. The value is fixed at <b>Addon</b> and cannot be changed. Any user-defined value is invalid.
apiVersion	Yes	String	API version. The value is fixed at <b>v3</b> and cannot be changed. Any user-defined value is invalid.

Parameter	Mandatory	Type	Description
metadata	Yes	<a href="#">AddonMetadata</a> object	Basic information about the object. Metadata is a collection of attributes.
spec	Yes	<a href="#">InstanceRequestSpec</a> object	Detailed description of add-on installation or upgrade.

**Table 4-771** AddonMetadata

Parameter	Mandatory	Type	Description
uid	No	String	Unique ID
name	No	String	Add-on name
alias	No	String	Add-on alias
labels	No	Map<String,String>	Add-on labels in key-value pairs. This is a reserved field and does not take effect.
annotations	No	Map<String,String>	Add-on annotations in the format of key-value pairs. <ul style="list-style-type: none"> <li>For add-on installation, the value is fixed at <b>{"addon.install/type":"install"}</b>.</li> <li>For add-on upgrade, the value is fixed at <b>{"addon.upgrade/type":"upgrade"}</b>.</li> </ul>
updateTimestamp	No	String	Updated at
creationTimestamp	No	String	Created at



**Table 4-772** InstanceRequestSpec

Parameter	Mandatory	Type	Description
version	No	String	Version of the add-on to install or upgrade, for example, <b>1.0.0</b> . <ul style="list-style-type: none"> <li>Installation: This parameter is optional. If not specified, the latest version supported by the cluster is used.</li> <li>Upgrade: This parameter is mandatory. The version number must be specified.</li> </ul>
clusterID	Yes	String	Cluster ID.
values	Yes	Map<String,Object>	Add-on template installation parameters (varying depending on the add-on). During the add-on upgrade, you need to specify all the installation parameters. If the parameters are not specified, the default values in the add-on template are used. The current add-on installation parameters can be obtained through the API for querying add-on instances.
addonTemplateName	Yes	String	Name of the add-on template to be installed, for example, <b>coredns</b> .

## Response Parameters

Status code: 200

**Table 4-773** Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at <b>Addon</b> and cannot be changed.
apiVersion	String	API version. The value is fixed at <b>v3</b> and cannot be changed.
metadata	<a href="#">AddonMetadata</a> object	Basic information about the object. Metadata is a collection of attributes.

Parameter	Type	Description
spec	<a href="#">InstanceSpec</a> object	Detailed description of the add-on instance.
status	<a href="#">AddonInstanceStatus</a> object	Add-on instance status.

**Table 4-774** AddonMetadata

Parameter	Type	Description
uid	String	Unique ID
name	String	Add-on name
alias	String	Add-on alias
labels	Map<String,String>	Add-on labels in key-value pairs. This is a reserved field and does not take effect.
annotations	Map<String,String>	Add-on annotations in the format of key-value pairs. <ul style="list-style-type: none"> <li>For add-on installation, the value is fixed at <b>{"addon.install/type":"install"}</b>.</li> <li>For add-on upgrade, the value is fixed at <b>{"addon.upgrade/type":"upgrade"}</b>.</li> </ul>
updateTimestamp	String	Updated at
creationTimestamp	String	Created at

**Table 4-775** InstanceSpec

Parameter	Type	Description
clusterID	String	Cluster ID.
version	String	Add-on template version, for example, <b>1.0.0</b> .
addonTemplateName	String	Add-on template name, for example, <b>coredns</b> .
addonTemplateType	String	Add-on template type.
addonTemplateLogo	String	URL for obtaining the add-on template logo.

Parameter	Type	Description
addonTemplateLabels	Array of strings	Labels of the add-on template.
description	String	Add-on template description.
values	Map<String, Object>	Add-on template installation parameters (varying depending on the add-on). Set the parameters accordingly.

**Table 4-776** AddonInstanceStatus

Parameter	Type	Description
status	String	<p>Add-on instance status. Options:</p> <ul style="list-style-type: none"> <li>● <b>running</b>: All of the add-on instances are running. This specifies that the add-on runs properly.</li> <li>● <b>abnormal</b>: The add-on instances are abnormal and the add-on cannot be used. You can click the add-on name to view exceptions.</li> <li>● <b>installing</b>: The add-on is being installed.</li> <li>● <b>installFailed</b>: Installing the add-on failed. In this case, uninstall the add-on and then reinstall it.</li> <li>● <b>upgrading</b>: The add-on is being upgraded.</li> <li>● <b>upgradeFailed</b>: Upgrading the add-on failed. In this case, upgrade the add-on again or uninstall the add-on and reinstall it.</li> <li>● <b>deleting</b>: The add-on is being deleted.</li> <li>● <b>deleteFailed</b>: Deleting the add-on failed. In this case, uninstall the add-on again.</li> <li>● <b>deleteSuccess</b>: Deleting the add-on succeeded.</li> <li>● <b>available</b>: Only some instances of the add-on are running. This specifies that some functions of the add-on are available.</li> <li>● <b>rollbacking</b>: The add-on is being rolled back.</li> <li>● <b>rollbackFailed</b>: Rolling back the add-on failed. In this case, roll back the add-on again or uninstall the add-on and reinstall it.</li> <li>● <b>unknown</b>: The add-on chart instance does not exist.</li> </ul>
Reason	String	Cause of the add-on installation failure.
message	String	Installation error details.
targetVersions	Array of strings	Versions to which the current add-on version can be upgraded.
currentVersion	<b>Versions</b> object	Information about the current add-on version.
isRollbackable	Boolean	Whether the add-on version can be rolled back to the source version.

Parameter	Type	Description
previousVersion	String	The add-on version before upgrade or rollback

**Table 4-777** Versions

Parameter	Type	Description
version	String	Add-on version.
input	Object	Add-on installation parameters.
stable	Boolean	Whether the add-on version is a stable release.
translate	Object	Translation information used by the GUI.
supportVersions	Array of <a href="#">SupportVersions</a> objects	Cluster versions that support the add-on.
creationTimestamp	String	Creation time.
updateTimestamp	String	Update time.

**Table 4-778** SupportVersions

Parameter	Type	Description
clusterType	String	Cluster type that supports the add-on.
clusterVersion	Array of strings	Cluster versions that support the add-on. The value is a regular expression.

## Example Requests

Update the everest add-on to version 2.1.30.

```
{
  "kind": "Addon",
  "apiVersion": "v3",
  "metadata": {
    "annotations": {
      "addon.upgrade/type": "upgrade"
    }
  },
  "spec": {
    "clusterID": "*****",
    "version": "2.1.30",
    "addonTemplateName": "everest",
    "values": {
      "basic": {
```

```

    "bms_url" : "*****",
    "driver_init_image_version" : "2.1.30",
    "ecsEndpoint" : "*****",
    "everest_image_version" : "2.1.30",
    "evs_url" : "*****",
    "iam_url" : "*****",
    "ims_url" : "*****",
    "obs_url" : "*****",
    "platform" : "linux-amd64",
    "sfs30_url" : "*****",
    "sfs_turbo_url" : "*****",
    "sfs_url" : "*****",
    "supportHcs" : false,
    "swr_addr" : "*****",
    "swr_user" : "hwofficial",
    "rbac_enabled" : true,
    "cluster_version" : "v1.23"
  },
  "flavor" : {
    "description" : "High available",
    "name" : "HA",
    "replicas" : 2,
    "resources" : [ {
      "limitsCpu" : "250m",
      "limitsMem" : "2000Mi",
      "name" : "everest-csi-controller",
      "requestsCpu" : "250m",
      "requestsMem" : "1500Mi"
    }, {
      "limitsCpu" : "500m",
      "limitsMem" : "300Mi",
      "name" : "everest-csi-driver",
      "requestsCpu" : "100m",
      "requestsMem" : "300Mi"
    } ],
    "category" : [ "CCE", "Turbo" ]
  },
  "custom" : {
    "cluster_id" : "*****",
    "csi_attacher_detach_worker_threads" : "60",
    "csi_attacher_worker_threads" : "60",
    "default_vpc_id" : "*****",
    "disable_auto_mount_secret" : false,
    "enable_node_attacher" : true,
    "flow_control" : { },
    "multiAZEnabled" : false,
    "over_subscription" : "80",
    "project_id" : "*****",
    "volume_attaching_flow_ctrl" : "0"
  }
}
}
}
}

```

## Example Responses

**Status code: 200**

OK

```

{
  "kind" : "Addon",
  "apiVersion" : "v3",
  "metadata" : {
    "uid" : "*****",
    "name" : "everest",
    "alias" : "everest",
    "creationTimestamp" : "2023-07-03T10:57:43Z",
    "updateTimestamp" : "2023-07-03T11:24:05Z"
  }
}

```

```

},
"spec" : {
  "clusterID" : "*****",
  "version" : "2.1.30",
  "addonTemplateName" : "everest",
  "addonTemplateType" : "helm",
  "addonTemplateLogo" : "*****",
  "addonTemplateLabels" : [ "Storage" ],
  "description" : "",
  "values" : {
    "basic" : {
      "bms_url" : "*****",
      "cluster_version" : "v1.23",
      "driver_init_image_version" : "2.1.30",
      "ecsEndpoint" : "*****",
      "everest_image_version" : "2.1.30",
      "evs_url" : "*****",
      "iam_url" : "*****",
      "ims_url" : "*****",
      "obs_url" : "*****",
      "platform" : "linux-amd64",
      "rbac_enabled" : true,
      "sfs30_url" : "*****",
      "sfs_turbo_url" : "*****",
      "sfs_url" : "*****",
      "supportHcs" : false,
      "swr_addr" : "*****",
      "swr_user" : "hwofficial"
    },
    "custom" : {
      "cluster_id" : "*****",
      "cluster_version" : "v1.23.8-r0",
      "csi_attacher_detach_worker_threads" : "60",
      "csi_attacher_worker_threads" : "60",
      "default_vpc_id" : "*****",
      "disable_auto_mount_secret" : false,
      "enable_node_attacher" : false,
      "flow_control" : { },
      "multiAZEnabled" : false,
      "over_subscription" : "80",
      "project_id" : "*****",
      "volume_attaching_flow_ctrl" : "0"
    },
    "flavor" : {
      "category" : [ "CCE", "Turbo" ],
      "description" : "Has only one instance",
      "name" : "Single",
      "replicas" : 1,
      "resources" : [ {
        "limitsCpu" : "250m",
        "limitsMem" : "600Mi",
        "name" : "everest-csi-controller",
        "requestsCpu" : "250m",
        "requestsMem" : "600Mi"
      }, {
        "limitsCpu" : "100m",
        "limitsMem" : "300Mi",
        "name" : "everest-csi-driver",
        "requestsCpu" : "100m",
        "requestsMem" : "300Mi"
      } ]
    },
    "systemAutoInject" : {
      "cluster" : {
        "clusterID" : "*****",
        "clusterNetworkMode" : "vpc-router",
        "clusterVersion" : "v1.23.8-r0"
      },
      "user" : {

```

```

    "projectID" : "*****"
  }
}
},
"status" : {
  "status" : "upgrading",
  "Reason" : "addon upgrading",
  "message" : "",
  "targetVersions" : null,
  "isRollbackable" : false,
  "currentVersion" : {
    "version" : "2.1.30",
    "input" : {
      "basic" : {
        "bms_url" : "*****",
        "driver_init_image_version" : "2.1.30",
        "ecsEndpoint" : "*****",
        "everest_image_version" : "2.1.30",
        "evs_url" : "*****",
        "iam_url" : "*****",
        "ims_url" : "*****",
        "obs_url" : "*****",
        "platform" : "*****",
        "sfs30_url" : "*****",
        "sfs_turbo_url" : "*****",
        "sfs_url" : "*****",
        "supportHcs" : false,
        "swr_addr" : "*****",
        "swr_user" : "hwofficial"
      },
      "parameters" : {
        "common" : {
          "defaultVPCId" : 1234567
        },
        "custom" : {
          "cluster_id" : "",
          "csi_attacher_detach_worker_threads" : "60",
          "csi_attacher_worker_threads" : "60",
          "default_vpc_id" : "",
          "disable_auto_mount_secret" : false,
          "enable_node_attacher" : false,
          "flow_control" : { },
          "multiAZEnabled" : false,
          "over_subscription" : "80",
          "project_id" : "",
          "volume_attaching_flow_ctrl" : "0"
        },
        "flavor1" : {
          "description" : "High available",
          "name" : "HA",
          "replicas" : 2,
          "resources" : [ {
            "limitsCpu" : "250m",
            "limitsMem" : "1500Mi",
            "name" : "everest-csi-controller",
            "requestsCpu" : "250m",
            "requestsMem" : "600Mi"
          }, {
            "limitsCpu" : "500m",
            "limitsMem" : "300Mi",
            "name" : "everest-csi-driver",
            "requestsCpu" : "100m",
            "requestsMem" : "300Mi"
          } ]
        },
        "flavor2" : {
          "description" : "Has only one instance",
          "name" : "Single",

```



```

    "replicas" : 1,
    "resources" : [ {
      "limitsCpu" : "250m",
      "limitsMem" : "600Mi",
      "name" : "everest-csi-controller",
      "requestsCpu" : "250m",
      "requestsMem" : "600Mi"
    }, {
      "limitsCpu" : "100m",
      "limitsMem" : "300Mi",
      "name" : "everest-csi-driver",
      "requestsCpu" : "100m",
      "requestsMem" : "300Mi"
    } ]
  },
  "flavor3" : {
    "description" : "custom resources",
    "name" : "custom-resources",
    "replicas" : 2,
    "resources" : [ {
      "limitsCpu" : "250m",
      "limitsMem" : "2000Mi",
      "name" : "everest-csi-controller",
      "requestsCpu" : "250m",
      "requestsMem" : "1500Mi"
    }, {
      "limitsCpu" : "500m",
      "limitsMem" : "300Mi",
      "name" : "everest-csi-driver",
      "requestsCpu" : "100m",
      "requestsMem" : "300Mi"
    } ]
  }
},
"stable" : true,
"translate" : {
  "en_US" : {
    "addon" : {
      "changeLog" : "",
      "description" : ""
    },
    "description" : {
      "Parameters.flavor1.description" : "Deploy the add-on with two instances, delivering high
availability but requiring more compute resources.",
      "Parameters.flavor1.name" : "HA",
      "Parameters.flavor2.description" : "Deploy the add-on with one instance.",
      "Parameters.flavor2.name" : "Standalone",
      "Parameters.flavor3.name" : "Custom"
    }
  },
  "fr_FR" : {
    "addon" : {
      "changeLog" : "",
      "description" : ""
    },
    "description" : {
      "Parameters.flavor1.description" : "Déployez avec deux instances, haute disponibilité.",
      "Parameters.flavor1.name" : "HA",
      "Parameters.flavor2.description" : "Déployez avec une seule instance.",
      "Parameters.flavor2.name" : "Célibataire",
      "Parameters.flavor3.name" : "Douane"
    }
  },
  "zh_CN" : {
    "addon" : {
      "changeLog" : "",
      "description" : ""
    }
  },

```

```

    "description" : {
      "Parameters.flavor1.description" : "Dual-instance deployment features high availability but requires
more compute resources.",
      "Parameters.flavor1.name" : "High availability",
      "Parameters.flavor2.description" : "Single-instance deployment",
      "Parameters.flavor2.name" : "Single instance",
      "Parameters.flavor3.description" : "Custom resource specifications for deployment",
      "Parameters.flavor3.name" : "Custom"
    }
  }
},
"supportVersions" : null,
"creationTimestamp" : "2023-05-12T16:10:05Z",
"updateTimestamp" : "2023-05-12T16:10:05Z"
}
}
}

```

### Status Codes

Status Code	Description
200	OK

### Error Codes

See [Error Codes](#).

## 4.6.4 Rolling Back an Add-on Instance

### Function

Roll back the add-on instance version to the source version. Only when the **status.isRollbackable** parameter is set to **true**, and the add-on instance is in the **running**, **available**, **abnormal**, **upgradeFailed**, or **rollbackFailed** state, the version can be rolled back to the source version.

### URI

POST /api/v3/addons/{id}/operation/rollback

**Table 4-779** Path Parameters

Parameter	Mandatory	Type	Description
id	Yes	String	Add-on instance ID

## Request Parameters

**Table 4-780** Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

**Table 4-781** Request body parameters

Parameter	Mandatory	Type	Description
clusterID	Yes	String	Cluster ID

## Response Parameters

**Status code: 200**

**Table 4-782** Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at <b>Addon</b> and cannot be changed.
apiVersion	String	API version. The value is fixed at <b>v3</b> and cannot be changed.
metadata	<a href="#">AddonMetadata</a> object	Basic information about the object. Metadata is a collection of attributes.
spec	<a href="#">InstanceSpec</a> object	Detailed description of the add-on instance.
status	<a href="#">AddonInstanceStatus</a> object	Add-on instance status.

**Table 4-783** AddonMetadata

Parameter	Type	Description
uid	String	Unique ID
name	String	Add-on name
alias	String	Add-on alias
labels	Map<String,String>	Add-on labels in key-value pairs. This is a reserved field and does not take effect.
annotations	Map<String,String>	Add-on annotations in the format of key-value pairs. <ul style="list-style-type: none"> <li>For add-on installation, the value is fixed at <b>{"addon.install/type":"install"}</b>.</li> <li>For add-on upgrade, the value is fixed at <b>{"addon.upgrade/type":"upgrade"}</b>.</li> </ul>
updateTimestamp	String	Updated at
creationTimestamp	String	Created at

**Table 4-784** InstanceSpec

Parameter	Type	Description
clusterID	String	Cluster ID.
version	String	Add-on template version, for example, <b>1.0.0</b> .
addonTemplateName	String	Add-on template name, for example, <b>coredns</b> .
addonTemplateType	String	Add-on template type.
addonTemplateLogo	String	URL for obtaining the add-on template logo.
addonTemplateLabels	Array of strings	Labels of the add-on template.
description	String	Add-on template description.
values	Map<String,Object>	Add-on template installation parameters (varying depending on the add-on). Set the parameters accordingly.

**Table 4-785** AddonInstanceStatus

Parameter	Type	Description
status	String	<p>Add-on instance status. Options:</p> <ul style="list-style-type: none"> <li>● <b>running</b>: All of the add-on instances are running. This specifies that the add-on runs properly.</li> <li>● <b>abnormal</b>: The add-on instances are abnormal and the add-on cannot be used. You can click the add-on name to view exceptions.</li> <li>● <b>installing</b>: The add-on is being installed.</li> <li>● <b>installFailed</b>: Installing the add-on failed. In this case, uninstall the add-on and then reinstall it.</li> <li>● <b>upgrading</b>: The add-on is being upgraded.</li> <li>● <b>upgradeFailed</b>: Upgrading the add-on failed. In this case, upgrade the add-on again or uninstall the add-on and reinstall it.</li> <li>● <b>deleting</b>: The add-on is being deleted.</li> <li>● <b>deleteFailed</b>: Deleting the add-on failed. In this case, uninstall the add-on again.</li> <li>● <b>deleteSuccess</b>: Deleting the add-on succeeded.</li> <li>● <b>available</b>: Only some instances of the add-on are running. This specifies that some functions of the add-on are available.</li> <li>● <b>rollbacking</b>: The add-on is being rolled back.</li> <li>● <b>rollbackFailed</b>: Rolling back the add-on failed. In this case, roll back the add-on again or uninstall the add-on and reinstall it.</li> <li>● <b>unknown</b>: The add-on chart instance does not exist.</li> </ul>
Reason	String	Cause of the add-on installation failure.
message	String	Installation error details.
targetVersions	Array of strings	Versions to which the current add-on version can be upgraded.
currentVersion	<b>Versions</b> object	Information about the current add-on version.
isRollbackable	Boolean	Whether the add-on version can be rolled back to the source version.

Parameter	Type	Description
previousVersion	String	The add-on version before upgrade or rollback

**Table 4-786** Versions

Parameter	Type	Description
version	String	Add-on version.
input	Object	Add-on installation parameters.
stable	Boolean	Whether the add-on version is a stable release.
translate	Object	Translation information used by the GUI.
supportVersions	Array of <a href="#">SupportVersions</a> objects	Cluster versions that support the add-on.
creationTimestamp	String	Creation time.
updateTimestamp	String	Update time.

**Table 4-787** SupportVersions

Parameter	Type	Description
clusterType	String	Cluster type that supports the add-on.
clusterVersion	Array of strings	Cluster versions that support the add-on. The value is a regular expression.

## Example Requests

```
{
  "clusterID" : "*****"
}
```

## Example Responses

**Status code: 200**

Add-on instance version is rolled back.

```
{
  "kind" : "Addon",
  "apiVersion" : "v3",
  "metadata" : {
    "uid" : "*****",

```

```

"name" : "everest",
"alias" : "everest",
"creationTimestamp" : "2023-03-15T02:48:01Z",
"updateTimestamp" : "2023-03-15T04:18:45Z"
},
"spec" : {
"clusterID" : "*****",
"version" : "2.1.16",
"addonTemplateName" : "everest",
"addonTemplateType" : "helm",
"addonTemplateLogo" : "*****",
"addonTemplateLabels" : [ "Storage" ],
"description" : "Everest is a cloud native container storage system based on CSI, used\erto support cloud
storages services for Kubernetes",
"values" : {
"basic" : {
"base_image" : "euleros",
"bms_url" : "*****",
"cluster_version" : "v1.25",
"driver_init_image_version" : "2.1.16",
"ecsEndpoint" : "*****",
"euleros_version" : "2.2.5",
"everest_image_version" : "2.1.16",
"evs_url" : "*****",
"iam_url" : "*****",
"ims_url" : "*****",
"obs_url" : "*****",
"platform" : "linux-amd64",
"rbac_enabled" : true,
"sfs30_url" : "*****",
"sfs_turbo_url" : "*****",
"sfs_url" : "*****",
"supportHcs" : false,
"swr_addr" : "*****",
"swr_user" : "*****"
},
"custom" : {
"cluster_id" : "*****",
"cluster_version" : "v1.25.3-r0",
"csi_attacher_detach_worker_threads" : "60",
"csi_attacher_worker_threads" : "60",
"default_vpc_id" : "*****",
"disable_auto_mount_secret" : false,
"enable_node_attacher" : false,
"flow_control" : { },
"multiAZEnabled" : false,
"over_subscription" : "80",
"project_id" : "*****",
"volume_attaching_flow_ctrl" : "0"
},
"flavor" : {
"category" : [ "CCE", "Turbo" ],
"description" : "High available",
"name" : "HA",
"replicas" : 2,
"resources" : [ {
"limitsCpu" : "250m",
"limitsMem" : "1500Mi",
"name" : "everest-csi-controller",
"requestsCpu" : "250m",
"requestsMem" : "600Mi"
}, {
"limitsCpu" : "500m",
"limitsMem" : "300Mi",
"name" : "everest-csi-driver",
"requestsCpu" : "100m",
"requestsMem" : "300Mi"
}
]
},
},

```

```

"multiAZPreferred" : {
  "podAntiAffinity" : {
    "preferredDuringSchedulingIgnoredDuringExecution" : [ {
      "podAffinityTerm" : {
        "labelSelector" : {
          "matchExpressions" : [ {
            "key" : "app",
            "operator" : "In",
            "values" : [ "everest-csi-controller" ]
          } ]
        },
        "topologyKey" : "topology.kubernetes.io/zone"
      },
      "weight" : 100
    } ]
  }
},
"multiAZRequired" : {
  "podAntiAffinity" : {
    "requiredDuringSchedulingIgnoredDuringExecution" : [ {
      "labelSelector" : {
        "matchExpressions" : [ {
          "key" : "app",
          "operator" : "In",
          "values" : [ "everest-csi-controller" ]
        } ]
      },
      "topologyKey" : "topology.kubernetes.io/zone"
    } ]
  }
},
"systemAutoInject" : {
  "cluster" : {
    "category" : "CCE",
    "clusterID" : "*****",
    "clusterNetworkMode" : "vpc-router",
    "clusterVersion" : "v1.25.3-r0"
  },
  "user" : {
    "projectID" : "*****"
  }
},
"tolerations" : [ {
  "effect" : "NoExecute",
  "key" : "node.kubernetes.io/not-ready",
  "operator" : "Exists",
  "tolerationSeconds" : 60
}, {
  "effect" : "NoExecute",
  "key" : "node.kubernetes.io/unreachable",
  "operator" : "Exists",
  "tolerationSeconds" : 60
} ]
},
"status" : {
  "status" : "rollbacking",
  "Reason" : "Rollback to 4",
  "message" : "",
  "targetVersions" : [ "2.1.18", "2.1.19" ],
  "isRollbackable" : false,
  "previousVersion" : "2.1.19",
  "currentVersion" : {
    "version" : "2.1.16",
    "input" : {
      "basic" : {
        "bms_url" : "*****",
        "driver_init_image_version" : "2.1.16",
        "ecsEndpoint" : "*****",

```



```

"everest_image_version" : "2.1.16",
"evs_url" : "*****",
"iam_url" : "*****",
"ims_url" : "*****",
"obs_url" : "*****",
"platform" : "linux-amd64",
"sfs30_url" : "*****",
"sfs_turbo_url" : "*****",
"sfs_url" : "*****",
"supportHcs" : false,
"swr_addr" : "*****",
"swr_user" : "*****"
},
"parameters" : {
  "common" : {
    "defaultVPCId" : 0
  },
  "custom" : {
    "cluster_id" : "",
    "csi_attacher_detach_worker_threads" : "60",
    "csi_attacher_worker_threads" : "60",
    "default_vpc_id" : "",
    "disable_auto_mount_secret" : false,
    "enable_node_attacher" : false,
    "flow_control" : { },
    "multiAZEnabled" : false,
    "over_subscription" : "80",
    "project_id" : "",
    "volume_attaching_flow_ctrl" : "0"
  },
  "flavor1" : {
    "description" : "High available",
    "name" : "HA",
    "replicas" : 2,
    "resources" : [ {
      "limitsCpu" : "250m",
      "limitsMem" : "1500Mi",
      "name" : "everest-csi-controller",
      "requestsCpu" : "250m",
      "requestsMem" : "600Mi"
    }, {
      "limitsCpu" : "500m",
      "limitsMem" : "300Mi",
      "name" : "everest-csi-driver",
      "requestsCpu" : "100m",
      "requestsMem" : "300Mi"
    } ]
  },
  "flavor2" : {
    "description" : "Has only one instance",
    "name" : "Single",
    "replicas" : 1,
    "resources" : [ {
      "limitsCpu" : "250m",
      "limitsMem" : "600Mi",
      "name" : "everest-csi-controller",
      "requestsCpu" : "250m",
      "requestsMem" : "600Mi"
    }, {
      "limitsCpu" : "100m",
      "limitsMem" : "300Mi",
      "name" : "everest-csi-driver",
      "requestsCpu" : "100m",
      "requestsMem" : "300Mi"
    } ]
  },
  "flavor3" : {
    "description" : "custom resources",
    "name" : "custom-resources",

```

```

"replicas" : 2,
"resources" : [ {
  "limitsCpu" : "250m",
  "limitsMem" : "2000Mi",
  "name" : "everest-csi-controller",
  "requestsCpu" : "250m",
  "requestsMem" : "1500Mi"
}, {
  "limitsCpu" : "500m",
  "limitsMem" : "300Mi",
  "name" : "everest-csi-driver",
  "requestsCpu" : "100m",
  "requestsMem" : "300Mi"
} ]
}
},
"stable" : true,
"translate" : {
  "en_US" : {
    "addon" : {
      "changeLog" : "*****",
      "description" : "*****"
    },
    "description" : {
      "Parameters.flavor1.description" : "*****",
      "Parameters.flavor1.name" : "*****",
      "Parameters.flavor2.description" : "*****",
      "Parameters.flavor2.name" : "*****",
      "Parameters.flavor3.name" : "*****"
    }
  },
  "zh_CN" : {
    "addon" : {
      "changeLog" : "*****",
      "description" : "*****"
    },
    "description" : {
      "Parameters.flavor1.description" : "*****",
      "Parameters.flavor1.name" : "*****",
      "Parameters.flavor2.description" : "*****",
      "Parameters.flavor2.name" : "*****",
      "Parameters.flavor3.description" : "*****",
      "Parameters.flavor3.name" : "*****"
    }
  }
},
"supportVersions" : null,
"creationTimestamp" : "2023-02-21T16:29:02Z",
"updateTimestamp" : "2023-02-22T06:49:50Z"
}
}
}

```

## Status Codes

Status Code	Description
200	Add-on instance version is rolled back.

## Error Codes

See [Error Codes](#).

## 4.6.5 Deleting an Add-on Instance

### Function

This API is used to delete an add-on instance.

### URI

DELETE /api/v3/addons/{id}

**Table 4-788** Path Parameters

Parameter	Mandatory	Type	Description
id	Yes	String	Add-on instance ID.

**Table 4-789** Query Parameters

Parameter	Mandatory	Type	Description
cluster_id	No	String	Cluster ID (deprecated). For details about how to obtain the cluster ID, see <a href="#">How Do I Obtain Parameters in the API URI</a> .

### Request Parameters

**Table 4-790** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

### Response Parameters

None

## Example Requests

None

## Example Responses

Status code: 200

OK

success

## Status Codes

Status Code	Description
200	OK

## Error Codes

See [Error Codes](#).

## 4.6.6 Querying an Add-on Instance

### Function

This API is used to obtain details about an add-on instance.

### URI

GET /api/v3/addons/{id}

**Table 4-791** Path Parameters

Parameter	Mandatory	Type	Description
id	Yes	String	Add-on instance ID.

**Table 4-792** Query Parameters

Parameter	Mandatory	Type	Description
cluster_id	No	String	Cluster ID (deprecated). For details about how to obtain the cluster ID, see <a href="#">How Do I Obtain Parameters in the API URI</a> .

## Request Parameters

**Table 4-793** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

## Response Parameters

Status code: 200

**Table 4-794** Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at <b>Addon</b> and cannot be changed.
apiVersion	String	API version. The value is fixed at <b>v3</b> and cannot be changed.
metadata	<a href="#">AddonMetadata</a> object	Basic information about the object. Metadata is a collection of attributes.
spec	<a href="#">InstanceSpec</a> object	Detailed description of the add-on instance.
status	<a href="#">AddonInstanceStatus</a> object	Add-on instance status.

**Table 4-795** AddonMetadata

Parameter	Type	Description
uid	String	Unique ID
name	String	Add-on name
alias	String	Add-on alias

Parameter	Type	Description
labels	Map<String,String>	Add-on labels in key-value pairs. This is a reserved field and does not take effect.
annotations	Map<String,String>	Add-on annotations in the format of key-value pairs. <ul style="list-style-type: none"> <li>For add-on installation, the value is fixed at <b>{"addon.install/type":"install"}</b>.</li> <li>For add-on upgrade, the value is fixed at <b>{"addon.upgrade/type":"upgrade"}</b>.</li> </ul>
updateTimestamp	String	Updated at
creationTimestamp	String	Created at

**Table 4-796** InstanceSpec

Parameter	Type	Description
clusterID	String	Cluster ID.
version	String	Add-on template version, for example, <b>1.0.0</b> .
addonTemplateName	String	Add-on template name, for example, <b>coredns</b> .
addonTemplateType	String	Add-on template type.
addonTemplateLogo	String	URL for obtaining the add-on template logo.
addonTemplateLabels	Array of strings	Labels of the add-on template.
description	String	Add-on template description.
values	Map<String,Object>	Add-on template installation parameters (varying depending on the add-on). Set the parameters accordingly.

**Table 4-797** AddonInstanceStatus

Parameter	Type	Description
status	String	<p>Add-on instance status. Options:</p> <ul style="list-style-type: none"> <li>● <b>running</b>: All of the add-on instances are running. This specifies that the add-on runs properly.</li> <li>● <b>abnormal</b>: The add-on instances are abnormal and the add-on cannot be used. You can click the add-on name to view exceptions.</li> <li>● <b>installing</b>: The add-on is being installed.</li> <li>● <b>installFailed</b>: Installing the add-on failed. In this case, uninstall the add-on and then reinstall it.</li> <li>● <b>upgrading</b>: The add-on is being upgraded.</li> <li>● <b>upgradeFailed</b>: Upgrading the add-on failed. In this case, upgrade the add-on again or uninstall the add-on and reinstall it.</li> <li>● <b>deleting</b>: The add-on is being deleted.</li> <li>● <b>deleteFailed</b>: Deleting the add-on failed. In this case, uninstall the add-on again.</li> <li>● <b>deleteSuccess</b>: Deleting the add-on succeeded.</li> <li>● <b>available</b>: Only some instances of the add-on are running. This specifies that some functions of the add-on are available.</li> <li>● <b>rollbacking</b>: The add-on is being rolled back.</li> <li>● <b>rollbackFailed</b>: Rolling back the add-on failed. In this case, roll back the add-on again or uninstall the add-on and reinstall it.</li> <li>● <b>unknown</b>: The add-on chart instance does not exist.</li> </ul>
Reason	String	Cause of the add-on installation failure.
message	String	Installation error details.
targetVersions	Array of strings	Versions to which the current add-on version can be upgraded.
currentVersion	<b>Versions</b> object	Information about the current add-on version.
isRollbackable	Boolean	Whether the add-on version can be rolled back to the source version.

Parameter	Type	Description
previousVersion	String	The add-on version before upgrade or rollback

**Table 4-798** Versions

Parameter	Type	Description
version	String	Add-on version.
input	Object	Add-on installation parameters.
stable	Boolean	Whether the add-on version is a stable release.
translate	Object	Translation information used by the GUI.
supportVersions	Array of <a href="#">SupportVersions</a> objects	Cluster versions that support the add-on.
creationTimestamp	String	Creation time.
updateTimestamp	String	Update time.

**Table 4-799** SupportVersions

Parameter	Type	Description
clusterType	String	Cluster type that supports the add-on.
clusterVersion	Array of strings	Cluster versions that support the add-on. The value is a regular expression.

## Example Requests

None

## Example Responses

**Status code: 200**

OK

```
{
  "kind": "Addon",
  "apiVersion": "v3",
  "metadata": {
    "uid": "24b23108-55c0-11e9-926f-0255ac101a31",
    "name": "gpu-beta",
    "alias": "gpu",
```



```

"creationTimestamp" : "2019-04-03T03:25:34Z",
"updateTimestamp" : "2019-04-03T03:25:34Z"
},
"spec" : {
"clusterID" : "0c0e4a63-5539-11e9-95f7-0255ac10177e",
"version" : "1.0.0",
"addonTemplateName" : "gpu-beta",
"addonTemplateType" : "helm",
"addonTemplateLogo" : "",
"addonTemplateLabels" : [ "Accelerator" ],
"description" : "A device plugin for nvidia.com/gpu resource on nvidia driver",
"values" : {
"basic" : {
"rbac_enabled" : true,
"swr_addr" : "100.125.6.246:20202",
"swr_user" : "hwofficial"
}
}
},
"status" : {
"status" : "installing",
"Reason" : "",
"message" : "",
"targetVersions" : null,
"currentVersion" : {
"version" : "1.0.0",
"input" : {
"basic" : {
"swr_addr" : "100.125.6.246:20202",
"swr_user" : "hwofficial"
}
},
"parameters" : {}
},
"stable" : true,
"translate" : {
"en_US" : {
"addon" : {
"changeLog" : "A device plugin for nvidia.com/gpu resource on nvidia driver",
"description" : "A device plugin for nvidia.com/gpu resource on nvidia driver"
}
},
"zh_CN" : {
"addon" : {
"changeLog" : "",
"description" : ""
}
}
},
"supportVersions" : null,
"creationTimestamp" : "2018-10-23T13:14:55Z",
"updateTimestamp" : "2018-12-07T09:40:24Z"
}
}
}

```

## Status Codes

Status Code	Description
200	OK

## Error Codes

See [Error Codes](#).

## 4.6.7 Listing Add-on Instances

### Function

This API is used to list all add-on instances in the cluster.

### URI

GET /api/v3/addons

**Table 4-800** Query Parameters

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	Cluster ID. For details about how to obtain a cluster ID, see <a href="#">How to Obtain Parameters in the API URI</a> .

### Request Parameters

**Table 4-801** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

### Response Parameters

Status code: 200

**Table 4-802** Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at <b>Addon</b> and cannot be changed.
apiVersion	String	API version. The value is fixed at <b>v3</b> and cannot be changed.

Parameter	Type	Description
items	Array of <a href="#">AddonInstance</a> objects	Add-on instance list.

**Table 4-803** AddonInstance

Parameter	Type	Description
kind	String	API type. The value is fixed at <b>Addon</b> and cannot be changed.
apiVersion	String	API version. The value is fixed at <b>v3</b> and cannot be changed.
metadata	<a href="#">AddonMetadata</a> object	Basic information about the object. Metadata is a collection of attributes.
spec	<a href="#">InstanceSpec</a> object	Detailed description of the add-on instance.
status	<a href="#">AddonInstanceStatus</a> object	Add-on instance status.

**Table 4-804** AddonMetadata

Parameter	Type	Description
uid	String	Unique ID
name	String	Add-on name
alias	String	Add-on alias
labels	Map<String,String>	Add-on labels in key-value pairs. This is a reserved field and does not take effect.
annotations	Map<String,String>	Add-on annotations in the format of key-value pairs. <ul style="list-style-type: none"> <li>For add-on installation, the value is fixed at <b>{"addon.install/type":"install"}</b>.</li> <li>For add-on upgrade, the value is fixed at <b>{"addon.upgrade/type":"upgrade"}</b>.</li> </ul>
updateTimestamp	String	Updated at
creationTimestamp	String	Created at

**Table 4-805** InstanceSpec

Parameter	Type	Description
clusterID	String	Cluster ID.
version	String	Add-on template version, for example, <b>1.0.0</b> .
addonTemplate Name	String	Add-on template name, for example, <b>coredns</b> .
addonTemplate Type	String	Add-on template type.
addonTemplate Logo	String	URL for obtaining the add-on template logo.
addonTemplate Labels	Array of strings	Labels of the add-on template.
description	String	Add-on template description.
values	Map<String, Object>	Add-on template installation parameters (varying depending on the add-on). Set the parameters accordingly.

**Table 4-806** AddonInstanceStatus

Parameter	Type	Description
status	String	<p>Add-on instance status. Options:</p> <ul style="list-style-type: none"> <li>● <b>running</b>: All of the add-on instances are running. This specifies that the add-on runs properly.</li> <li>● <b>abnormal</b>: The add-on instances are abnormal and the add-on cannot be used. You can click the add-on name to view exceptions.</li> <li>● <b>installing</b>: The add-on is being installed.</li> <li>● <b>installFailed</b>: Installing the add-on failed. In this case, uninstall the add-on and then reinstall it.</li> <li>● <b>upgrading</b>: The add-on is being upgraded.</li> <li>● <b>upgradeFailed</b>: Upgrading the add-on failed. In this case, upgrade the add-on again or uninstall the add-on and reinstall it.</li> <li>● <b>deleting</b>: The add-on is being deleted.</li> <li>● <b>deleteFailed</b>: Deleting the add-on failed. In this case, uninstall the add-on again.</li> <li>● <b>deleteSuccess</b>: Deleting the add-on succeeded.</li> <li>● <b>available</b>: Only some instances of the add-on are running. This specifies that some functions of the add-on are available.</li> <li>● <b>rollbacking</b>: The add-on is being rolled back.</li> <li>● <b>rollbackFailed</b>: Rolling back the add-on failed. In this case, roll back the add-on again or uninstall the add-on and reinstall it.</li> <li>● <b>unknown</b>: The add-on chart instance does not exist.</li> </ul>
Reason	String	Cause of the add-on installation failure.
message	String	Installation error details.
targetVersions	Array of strings	Versions to which the current add-on version can be upgraded.
currentVersion	<b>Versions</b> object	Information about the current add-on version.
isRollbackable	Boolean	Whether the add-on version can be rolled back to the source version.

Parameter	Type	Description
previousVersion	String	The add-on version before upgrade or rollback

**Table 4-807** Versions

Parameter	Type	Description
version	String	Add-on version.
input	Object	Add-on installation parameters.
stable	Boolean	Whether the add-on version is a stable release.
translate	Object	Translation information used by the GUI.
supportVersions	Array of <a href="#">SupportVersions</a> objects	Cluster versions that support the add-on.
creationTimestamp	String	Creation time.
updateTimestamp	String	Update time.

**Table 4-808** SupportVersions

Parameter	Type	Description
clusterType	String	Cluster type that supports the add-on.
clusterVersion	Array of strings	Cluster versions that support the add-on. The value is a regular expression.

## Example Requests

None

## Example Responses

**Status code: 200**

ok

```
{
  "kind": "Addon",
  "apiVersion": "v3",
  "items": [ {
    "kind": "Addon",
    "apiVersion": "v3",
    "metadata": {
```

```

"uid" : "8ca259cc-553b-11e9-926f-0255ac101a31",
"name" : "storage-driver",
"alias" : "storage-driver",
"creationTimestamp" : "2019-04-02T11:36:26Z",
"updateTimestamp" : "2019-04-02T11:36:26Z"
},
"spec" : {
"clusterID" : "0c0e4a63-5539-11e9-95f7-0255ac10177e",
"version" : "1.0.10",
"addonTemplateName" : "storage-driver",
"addonTemplateType" : "helm",
"addonTemplateLogo" : "https://192.149.48.66/cce-addon-southchina-aw1hz2u/storage-driverlogo.svg",
"addonTemplateLabels" : [ "Storage" ],
"description" : "A kubernetes FlexVolume Driver used to support cloud storage",
"values" : {
"basic" : {
"addon_version" : "1.0.10",
"euleros_version" : "2.2.5",
"obs_url" : "",
"platform" : "linux-amd64",
"swr_addr" : "100.125.6.246:20202",
"swr_user" : "hwofficial"
},
"flavor" : {
"replicas" : 1
},
},
"parameters" : {}
}
},
"status" : {
"status" : "running",
"Reason" : "Install complete",
"message" : "",
"targetVersions" : null,
"currentVersion" : {
"version" : "1.0.10",
"input" : {
"basic" : {
"euleros_version" : "2.2.5",
"obs_url" : "",
"swr_addr" : "100.125.6.246:20202",
"swr_user" : "hwofficial"
},
"parameters" : {}
},
},
"stable" : true,
"translate" : {
"en_US" : {
"addon" : {
"changeLog" : "The plug-in is upgraded to enhance the storage plug-in function.",
"description" : "A kubernetes FlexVolume Driver used to support cloud storage"
}
},
"zh_CN" : {
"addon" : {
"changeLog" : "",
"description" : ""
}
}
},
"supportVersions" : null,
"creationTimestamp" : "2019-03-29T13:45:37Z",
"updateTimestamp" : "2019-03-29T13:45:37Z"
}
}
}
}
}

```

## Status Codes

Status Code	Description
200	ok

## Error Codes

See [Error Codes](#).

# 4.7 Tag Management

## 4.7.1 Adding Resource Tags to a Specified Cluster in Batches

### Function

This API is used to add resource tags for a specified cluster in batches.

 **NOTE**

- Each cluster supports a maximum of 20 resource tags.
- This API is idempotent. If the to-be-added tag has the same tag key and tag value as an existing tag, the tag will be added. If the to-be-added tag has the same key but different value as an existing tag, the tag will overwrite the existing one.

### URI

POST /api/v3/projects/{project\_id}/clusters/{cluster\_id}/tags/create

**Table 4-809** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .



## Request Parameters

**Table 4-810** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

**Table 4-811** Request body parameters

Parameter	Mandatory	Type	Description
tags	Yes	Array of <a href="#">ResourceTag</a> objects	List of cluster resource tags to be created. Each cluster supports a maximum of 20 resource tags.

**Table 4-812** ResourceTag

Parameter	Mandatory	Type	Description
key	No	String	Key. <ul style="list-style-type: none"> <li>• Cannot be null. Max characters: 128.</li> <li>• Use letters, digits, and spaces in UTF-8 format.</li> <li>• Can contain the following special characters: <code>._:=+@</code></li> <li>• Cannot start with <code>_sys_</code>.</li> </ul>
value	No	String	Value. <ul style="list-style-type: none"> <li>• Can be null but not the default. Max characters: 255.</li> <li>• Use letters, digits, and spaces in UTF-8 format.</li> <li>• Can contain the following special characters: <code>._:/=+@.</code></li> </ul>

## Response Parameters

None

## Example Requests

Adding Resource Tags to a Specified Cluster in Batches

```
POST /api/v3/projects/{project_id}/clusters/{cluster_id}/tags/create
{
  "tags": [ {
    "key": "key1",
    "value": "value1"
  }, {
    "key": "key2",
    "value": "value3"
  } ]
}
```

## Example Responses

None

## Status Codes

Status Code	Description
204	No Content

## Error Codes

See [Error Codes](#).

## 4.7.2 Deleting Resource Tags of a Specified Cluster in Batches

### Function

This API is used to delete resource tags of a specified cluster in batches.

#### NOTE

- This API is idempotent. If the tag key of the to-be-deleted tag does not exist, the tag will be deleted.

### URI

```
POST /api/v3/projects/{project_id}/clusters/{cluster_id}/tags/delete
```

**Table 4-813** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

## Request Parameters

**Table 4-814** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

**Table 4-815** Request body parameters

Parameter	Mandatory	Type	Description
tags	Yes	Array of <a href="#">ResourceDeleteTag</a> objects	List of cluster resource tags to be deleted

**Table 4-816** ResourceDeleteTag

Parameter	Mandatory	Type	Description
key	No	String	<p>The key value of the resource tag</p> <ul style="list-style-type: none"> <li>It cannot be null. It can contain a maximum of 128 characters.</li> <li>It can contain letters, digits, and spaces in UTF-8 format.</li> <li>It can contain the following special characters: <code>./=+-@.</code></li> <li>It cannot start with <code>_sys_</code>.</li> </ul>

## Response Parameters

None

## Example Requests

Deleting Resource Tags of a Specified Cluster in Batches

```
POST /api/v3/projects/{project_id}/clusters/{cluster_id}/tags/delete
{
  "tags": [ {
    "key": "key1"
  }, {
    "key": "key2"
  } ]
}
```

## Example Responses

None

## Status Codes

Status Code	Description
204	No Content

## Error Codes

See [Error Codes](#).

## 4.8 Configuration Management

### 4.8.1 Obtaining the Parameters That Can Be Configured for a Node Pool

#### Function

This API is used to obtain the parameters that can be configured in a specified CCE node pool.

#### URI

GET /api/v3/projects/{project\_id}/clusters/{cluster\_id}/nodepools/{nodepool\_id}/configuration/detail

**Table 4-817** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
nodepool_id	Yes	String	Node pool ID.

#### Request Parameters

**Table 4-818** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

## Response Parameters

None

## Example Requests

None

## Example Responses

**Status code: 200**

The parameters of the specified node pool are obtained.

```
{
  "kubelet" : [ {
    "name" : "kube-api-qps",
    "default" : 300,
    "validAt" : "immediately",
    "empty" : false,
    "schema" : "",
    "type" : "float"
  } ]
}
```

## Status Codes

Status Code	Description
200	The parameters of the specified node pool are obtained.

## Error Codes

See [Error Codes](#).

## 4.8.2 Obtaining the List of Parameters That Can Be Configured for a Cluster

### Function

This API is used to obtain the list of parameters that can be configured for a CCE cluster.

### URI

GET /api/v3/projects/{project\_id}/clusters/{cluster\_id}/configuration/detail

**Table 4-819** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

## Request Parameters

**Table 4-820** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

## Response Parameters

None

## Example Requests

None

## Example Responses

**Status code: 200**

A list of parameters that can be configured for a cluster is obtained successfully.

```
{
  "kube-apiserver" : [ {
    "name" : "default-not-ready-toleration-seconds",
    "default" : 300,
    "validAt" : "immediately",
    "empty" : true,
    "schema" : "kubernetes",
    "type" : "int"
  } ]
}
```

```
}]
}
```

## Status Codes

Status Code	Description
200	A list of parameters that can be configured for a cluster is obtained successfully.

## Error Codes

See [Error Codes](#).

### 4.8.3 Obtaining the Parameters That Can Be Configured for a Node Pool

#### Function

This API is used to obtain the parameters that can be configured for a node pool.

#### URI

GET /api/v3/projects/{project\_id}/clusters/{cluster\_id}/nodepools/{nodepool\_id}/configuration

**Table 4-821** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
nodepool_id	Yes	String	Node pool ID.



## Request Parameters

**Table 4-822** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

## Response Parameters

Status code: 200

**Table 4-823** Response body parameters

Parameter	Type	Description
apiVersion	String	API version
kind	String	API type. The value is fixed at <b>Configuration</b> .
metadata	<a href="#">ConfigurationMetadata</a> object	Configuration metadata
spec	<a href="#">ClusterConfigurationsSpec</a> object	Configuration specifications
status	Object	Configuration status

**Table 4-824** ConfigurationMetadata

Parameter	Type	Description
name	String	Configuration name

Parameter	Type	Description
labels	Map<String,String>	<p>A configuration label in a key-value pair.</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain containing a maximum of 253 characters (such as <b>example.com/my-key</b>) can be prefixed to a key.</li> <li>• <b>Value:</b> The value can be left blank or contain 1 to 63 characters that start and end with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed.</li> </ul> <p>Example: "foo": "bar"</p>

**Table 4-825** ClusterConfigurationsSpec

Parameter	Type	Description
packages	Array of <a href="#">packages</a> objects	Component configuration item details

**Table 4-826** packages

Parameter	Type	Description
name	String	Component name
configurations	Array of <a href="#">ConfigurationItem</a> objects	Component configuration items

**Table 4-827** ConfigurationItem

Parameter	Type	Description
name	String	Component configuration item name.
value	Object	Component configuration item value.

## Example Requests

None

## Example Responses

### Status code: 200

The parameter configurations of the node pool are obtained successfully.

```
{
  "kind": "Configuration",
  "apiVersion": "v3",
  "metadata": {
    "name": "configuration",
    "labels": {
      "nodepool_id": "61de338d-a1f9-11ed-8891-0255ac100036"
    }
  },
  "spec": {
    "packages": [ {
      "name": "kube-apiserver",
      "configurations": [ {
        "name": "event-rate-limit-qps",
        "value": 200
      }, {
        "name": "support-overload",
        "value": false
      } ]
    }, {
      "name": "kube-scheduler",
      "configurations": [ {
        "name": "kube-api-qps",
        "value": 100
      }, {
        "name": "default-scheduler",
        "value": "kube-scheduler"
      } ]
    } ]
  },
  "status": { }
}
```

## Status Codes

Status Code	Description
200	The parameter configurations of the node pool are obtained successfully.

## Error Codes

See [Error Codes](#).

## 4.8.4 Changing the Values of Configuration Parameters of a Node Pool

### Function

This API is used to change the values of parameter configurations of a node pool in the CCE cluster.

## URI

PUT /api/v3/projects/{project\_id}/clusters/{cluster\_id}/nodepools/{nodepool\_id}/configuration

**Table 4-828** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .
nodepool_id	Yes	String	Node pool ID.

## Request Parameters

**Table 4-829** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

**Table 4-830** Request body parameters

Parameter	Mandatory	Type	Description
apiVersion	Yes	String	API version. The value is fixed at <b>v3</b> .
kind	Yes	String	API type. The value is fixed at <b>Configuration</b> .
metadata	Yes	<a href="#">ConfigurationMetadata</a> object	Configuration metadata

Parameter	Mandatory	Type	Description
spec	Yes	<a href="#">ClusterConfigurationsSpec</a> object	Configuration specifications

**Table 4-831** ConfigurationMetadata

Parameter	Mandatory	Type	Description
name	Yes	String	Configuration name
labels	No	Map<String,String>	<p>A configuration label in a key-value pair.</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain containing a maximum of 253 characters (such as <b>example.com/my-key</b>) can be prefixed to a key.</li> <li>• <b>Value:</b> The value can be left blank or contain 1 to 63 characters that start and end with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed.</li> </ul> <p>Example: "foo": "bar"</p>

**Table 4-832** ClusterConfigurationsSpec

Parameter	Mandatory	Type	Description
packages	Yes	Array of <a href="#">packages</a> objects	Component configuration item details

**Table 4-833** packages

Parameter	Mandatory	Type	Description
name	No	String	Component name

Parameter	Mandatory	Type	Description
configurations	No	Array of <b>ConfigurationItem</b> objects	Component configuration items

**Table 4-834** ConfigurationItem

Parameter	Mandatory	Type	Description
name	No	String	Component configuration item name.
value	No	Object	Component configuration item value.

## Response Parameters

Status code: 200

**Table 4-835** Response body parameters

Parameter	Type	Description
apiVersion	String	API version
kind	String	API type. The value is fixed at <b>Configuration</b> .
metadata	<b>ConfigurationMetadata</b> object	Configuration metadata
spec	<b>ClusterConfigurationsSpec</b> object	Configuration specifications
status	Object	Configuration status

**Table 4-836** ConfigurationMetadata

Parameter	Type	Description
name	String	Configuration name

Parameter	Type	Description
labels	Map<String,String>	<p>A configuration label in a key-value pair.</p> <ul style="list-style-type: none"> <li>• <b>Key:</b> Enter 1 to 63 characters starting with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain containing a maximum of 253 characters (such as <b>example.com/my-key</b>) can be prefixed to a key.</li> <li>• <b>Value:</b> The value can be left blank or contain 1 to 63 characters that start and end with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed.</li> </ul> <p>Example: "foo": "bar"</p>

**Table 4-837** ClusterConfigurationsSpec

Parameter	Type	Description
packages	Array of <a href="#">packages</a> objects	Component configuration item details

**Table 4-838** packages

Parameter	Type	Description
name	String	Component name
configurations	Array of <a href="#">ConfigurationItem</a> objects	Component configuration items

**Table 4-839** ConfigurationItem

Parameter	Type	Description
name	String	Component configuration item name.
value	Object	Component configuration item value.

## Example Requests

The following example describes how to update the **system-reserved-mem** and **kube-reserved-mem** parameters in kubelet.

```
/api/v3/projects/{project_id}/clusters/{cluster_id}/nodepools/{nodepool_id}/configuration

{
  "kind": "Configuration",
  "apiVersion": "v3",
  "metadata": {
    "name": "configuration"
  },
  "spec": {
    "packages": [ {
      "name": "kubelet",
      "configurations": [ {
        "name": "system-reserved-mem",
        "value": 600
      }, {
        "name": "kube-reserved-mem",
        "value": 800
      } ]
    } ]
  }
}
```

## Example Responses

**Status code: 200**

The parameter configurations of the node pool are updated successfully.

```
{
  "kind": "Configuration",
  "apiVersion": "v3",
  "metadata": {
    "name": "configuration",
    "labels": {
      "nodepool_id": "61de338d-a1f9-11ed-8891-0255ac100036"
    }
  },
  "spec": {
    "packages": [ {
      "name": "kube-apiserver",
      "configurations": [ {
        "name": "event-rate-limit-qps",
        "value": 200
      }, {
        "name": "support-overload",
        "value": false
      } ]
    }, {
      "name": "kube-scheduler",
      "configurations": [ {
        "name": "kube-api-qps",
        "value": 100
      }, {
        "name": "default-scheduler",
        "value": "kube-scheduler"
      } ]
    } ]
  },
  "status": { }
}
```



## Status Codes

Status Code	Description
200	The parameter configurations of the node pool are updated successfully.

## Error Codes

See [Error Codes](#).

# 4.9 Chart Management

## 4.9.1 Uploading a Chart

### Function

This API is used to upload a chart.

### URI

POST /v2/charts

### Request Parameters

**Table 4-840** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

**Table 4-841** FormData parameters

Parameter	Mandatory	Type	Description
parameters	No	String	Parameter configurations for uploading a chart. The following is an example: <code>{"override":true,"skip_lint":true,"source":"package"}</code> <ul style="list-style-type: none"> <li>skip_lint: whether lint uploaded chart - override: whether override existed chart - visible: update chart visible</li> </ul>
content	Yes	File	Chart package file

## Response Parameters

Status code: 201

**Table 4-842** Response body parameters

Parameter	Type	Description
id	String	Chart ID
name	String	Chart name
values	String	Chart value
translate	String	Chart translation resources
instruction	String	Chart description
version	String	Chart version
description	String	Chart description
source	String	Chart source
icon_url	String	URL to chart icons
public	Boolean	Whether the chart is public
chart_url	String	URL to the chart
create_at	String	Created at
update_at	String	Updated at

## Example Requests

POST /v2/charts

```
{
  "parameters" : "{ \"override\":true, \"skip_lint\":true, \"source\": \"package\" }",
  "content" : "chart-file.tgz"
}
```

## Example Responses

**Status code: 201**

Created

```
{
  "id" : "e99a7e86-afdd-11eb-aca3-0255ac100b0e",
  "name" : "neo4j",
  "values" : "{ \"acceptLicenseAgreement\": \"no\", \"affinity\": {}, \"authEnabled\": true, \"clusterDomain\": \"cluster.local\", \"core\": { \"initContainers\": [], \"numberOfServers\": 3, \"persistentVolume\": { \"enabled\": true, \"mountPath\": \"/data\", \"size\": \"10Gi\" }, \"sidecarContainers\": [], \"defaultDatabase\": \"neo4j\", \"image\": \"neo4j\", \"imagePullPolicy\": \"IfNotPresent\", \"imageTag\": \"4.0.3-enterprise\", \"name\": \"neo4j\", \"nodeSelector\": {}, \"podDisruptionBudget\": {}, \"readReplica\": { \"autoscaling\": { \"enabled\": false, \"maxReplicas\": 3, \"minReplicas\": 1, \"targetAverageUtilization\": 70 }, \"initContainers\": [], \"numberOfServers\": 0, \"resources\": {} }, \"sidecarContainers\": [], \"resources\": {}, \"testImage\": \"markneedham/k8s-kubectl\", \"testImageTag\": \"master\", \"tolerations\": [], \"useAPOC\": \"true\" }",
  "translate" : "",
  "instruction" : "README.md",
  "version" : "3.0.1",
  "description" : "DEPRECATED Neo4j is the world's leading graph database",
  "source" : "",
  "icon_url" : "https://info.neo4j.com/rs/773-GON-065/images/neo4j_logo.png",
  "public" : false,
  "chart_url" : "neo4j-3.0.1.tgz",
  "create_at" : "2021-05-08T08:53:13Z",
  "update_at" : "2021-05-08T08:53:13Z"
}
```

## Status Codes

Status Code	Description
201	Created

## Error Codes

See [Error Codes](#).

## 4.9.2 Obtaining a Chart List

### Function

This API is used to obtain the chart list.

### URI

GET /v2/charts

## Request Parameters

**Table 4-843** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

## Response Parameters

Status code: 200

**Table 4-844** Response body parameters

Parameter	Type	Description
[items]	Array of <a href="#">ChartResp</a> objects	Chart list

**Table 4-845** ChartResp

Parameter	Type	Description
id	String	Chart ID
name	String	Chart name
values	String	Chart value
translate	String	Chart translation resources
instruction	String	Chart description
version	String	Chart version
description	String	Chart description
source	String	Chart source
icon_url	String	URL to chart icons
public	Boolean	Whether the chart is public

Parameter	Type	Description
chart_url	String	URL to the chart
create_at	String	Created at
update_at	String	Updated at

## Example Requests

None

## Example Responses

**Status code: 200**

OK

```
[ {
  "id" : "1abd3bd6-0258-11ec-b8b0-0255ac100b05",
  "name" : "magento-mysql",
  "values" : "{\basic\":{\admin_password\":\*****\",admin_username\":\username\",app_name
\":\magento\",mysql_database\":\magento\",mysql_name\":\mysql\",mysql_password\":\*****
\",mysql_port\":3306,mysql_root_password\":\*****\",mysql_user\":\magento\",storage_class\":\csi-
nas\",storage_mode\":\ReadWriteMany\",storage_size\":\10G\",global\":{\magento_EIP
\":\100.100.100.100\",magento_EPORT\":32080,namespace\":\default\",image\":{\magento_image
\":\example.com/everest/magento:latest\",mysql_image\":\example.com/everest/mysql:5.7.14\"}},
  "translate" : "",
  "instruction" : "",
  "version" : "1.0.0",
  "description" : "chart description",
  "source" : "",
  "icon_url" : "https://example.com/magento-stack-110x117.png",
  "public" : false,
  "chart_url" : "magento-mysql-1.0.0.tgz",
  "create_at" : "2021-08-20T08:00:29Z",
  "update_at" : "2021-08-20T08:00:29Z"
} ]
```

## Status Codes

Status Code	Description
200	OK

## Error Codes

See [Error Codes](#).

## 4.9.3 Obtaining a Release List

### Function

This API is used to obtain a release list.

## URI

GET /cce/cam/v3/clusters/{cluster\_id}/releases

**Table 4-846** Path Parameters

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

**Table 4-847** Query Parameters

Parameter	Mandatory	Type	Description
chart_id	No	String	Chart ID
namespace	No	String	Namespace of the chart

## Request Parameters

**Table 4-848** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

## Response Parameters

Status code: 200

**Table 4-849** Response body parameters

Parameter	Type	Description
[items]	Array of <a href="#">ReleaseResp</a> objects	OK

**Table 4-850** ReleaseResp

Parameter	Type	Description
chart_name	String	Chart name
chart_public	Boolean	Whether the chart is public
chart_version	String	Chart version
cluster_id	String	Cluster ID
cluster_name	String	Current name
create_at	String	Created at
description	String	Release description
name	String	Release name
namespace	String	Namespace to which a chart release belongs
parameters	String	Release parameters
resources	String	Resources required by the release
status	String	Release status. <ul style="list-style-type: none"> <li>● <b>DEPLOYED</b>: The release is normal.</li> <li>● <b>DELETED</b>: The release has been deleted.</li> <li>● <b>FAILED</b>: The release fails to be deployed.</li> <li>● <b>DELETING</b>: The release is being deleted.</li> <li>● <b>PENDING_INSTALL</b>: The release is waiting to be installed.</li> <li>● <b>PENDING_UPGRADE</b>: The release is waiting to be upgraded.</li> <li>● <b>PENDING_ROLLBACK</b>: The release is waiting for rollback.</li> <li>● <b>UNKNOWN</b>: The release status is unknown, indicating that the release is abnormal. You can manually delete the release and reinstall it.</li> </ul>
status_description	String	Release status description

Parameter	Type	Description
update_at	String	Updated at
values	String	Release value
version	Integer	Release version

## Example Requests

None

## Example Responses

**Status code: 200**

OK

```
[ {
  "chart_name": "magento-mysql",
  "chart_public": false,
  "chart_version": "1.0.0",
  "cluster_id": "a870253f-5dc7-11ee-bf71-0255ac100b03",
  "cluster_name": "sfs-turbo-test",
  "create_at": "2023-11-14T20:30:57+08:00",
  "description": "Initial install underway",
  "name": "testwww",
  "namespace": "monitoring",
  "parameters": "",
  "resources": "",
  "status": "PENDING_INSTALL",
  "status_description": "Initial install underway",
  "update_at": "2023-11-14T20:30:57+08:00",
  "values": "{\n  \"basic\": {\n    \"admin_password\": \"*****\",\n    \"admin_username\": \"username\",\n    \"app_name\": \"magento\",\n    \"mysql_database\": \"magento\",\n    \"mysql_name\": \"mysql\",\n    \"mysql_password\": \"*****\",\n    \"mysql_port\": 3306,\n    \"mysql_root_password\": \"*****\",\n    \"mysql_user\": \"magento\",\n    \"storage_class\": \"csi-nas\",\n    \"storage_mode\": \"ReadWriteMany\",\n    \"storage_size\": \"10G\",\n    \"global\": {\n      \"magento_EIP\": \"100.100.100.100\",\n      \"magento_EPORT\": 32080,\n      \"namespace\": \"default\",\n      \"image\": {\n        \"magento_image\": \"example.com/everest/magento:latest\",\n        \"mysql_image\": \"example.com/everest/mysql:5.7.14\"}}}}",
  "version": 1
}
```

## Status Codes

Status Code	Description
200	OK

## Error Codes

See [Error Codes](#).



## 4.9.4 Updating a Chart

### Function

This API is used to update a chart.

### URI

PUT /v2/charts/{chart\_id}

**Table 4-851** Path Parameters

Parameter	Mandatory	Type	Description
chart_id	Yes	String	Chart ID

### Request Parameters

**Table 4-852** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

**Table 4-853** FormData parameters

Parameter	Mandatory	Type	Description
parameters	No	String	Parameter configurations for uploading a chart. The following is an example: {"override":true,"skip_lint":true,"source":"package"} - skip_lint: whether lint uploaded chart - override: whether override existed chart - visible: update chart visible
content	Yes	File	Chart package file

## Response Parameters

Status code: 200

Table 4-854 Response body parameters

Parameter	Type	Description
id	String	Chart ID
name	String	Chart name
values	String	Chart value
translate	String	Chart translation resources
instruction	String	Chart description
version	String	Chart version
description	String	Chart description
source	String	Chart source
icon_url	String	URL to chart icons
public	Boolean	Whether the chart is public
chart_url	String	URL to the chart
create_at	String	Created at
update_at	String	Updated at

## Example Requests

```
PUT /v2/charts/{chart_id}
{
  "parameters": "{ \"override\":true, \"skip_lint\":true, \"source\": \"package\" }",
  "content": "chart-file.tgz"
}
```

## Example Responses

Status code: 200

OK

```
{
  "id": "e99a7e86-afdd-11eb-aca3-0255ac100b0e",
  "name": "neo4j",
  "values": "{ \"acceptLicenseAgreement\": \"no\", \"affinity\": {}, \"authEnabled\": true, \"clusterDomain\": \"cluster.local\", \"core\": { \"initContainers\": [], \"numberOfServers\": 3, \"persistentVolume\": { \"enabled\": true, \"mountPath\": \"/data\", \"size\": \"10Gi\" }, \"sidecarContainers\": [] }, \"defaultDatabase\": \"neo4j\", \"image\": \"neo4j\", \"imagePullPolicy\": \"IfNotPresent\", \"imageTag\": \"4.0.3-enterprise\", \"name\": \"neo4j\", \"nodeSelector\": {}, \"podDisruptionBudget\": {}, \"readReplica\": { \"autoscaling\": { \"enabled\": false, \"maxReplicas\": 3, \"minReplicas\": 1, \"targetAverageUtilization\": 70 }, \"initContainers\": [], \"numberOfServers\": 0, \"resources\": {}, \"sidecarContainers\": [], \"resources\": {}, \"testImage\": \"markhneedham/k8s-kubectl\", \"testImageTag\": \"master\", \"tolerations\": [], \"useAPOC\": \"true\" }, \"translate\": \"\"",
```

```

"instruction" : "README.md",
"version" : "3.0.1",
"description" : "DEPRECATED Neo4j is the world's leading graph database",
"source" : "",
"icon_url" : "https://example.com/images/neo4j_logo.png",
"public" : false,
"chart_url" : "neo4j-3.0.1.tgz",
"create_at" : "2021-05-08T08:53:13Z",
"update_at" : "2021-05-08T08:53:13Z"
}

```

## Status Codes

Status Code	Description
200	OK

## Error Codes

See [Error Codes](#).

## 4.9.5 Creating a Release

### Function

This API is used to create a release.

### URI

POST /cce/cam/v3/clusters/{cluster\_id}/releases

**Table 4-855** Path Parameters

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

### Request Parameters

**Table 4-856** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

**Table 4-857** Request body parameters

Parameter	Mandatory	Type	Description
chart_id	Yes	String	Chart ID
description	No	String	Release description
name	Yes	String	Release name
namespace	Yes	String	Namespace to which a chart release belongs
version	Yes	String	Release version
parameters	No	<a href="#">ReleaseReqBodyParams</a> object	Release parameters
values	Yes	<a href="#">values</a> object	Release value

**Table 4-858** ReleaseReqBodyParams

Parameter	Mandatory	Type	Description
dry_run	No	Boolean	After this function is enabled, only chart parameters are verified, and installation is not performed.
name_template	No	String	Release name template
no_hooks	No	Boolean	Whether to disable hooks during installation
replace	No	Boolean	Whether to replace the release with the same name
recreate	No	Boolean	Whether to rebuild the release

Parameter	Mandatory	Type	Description
reset_values	No	Boolean	Whether to reset values during an update
release_version	No	Integer	Version of the rollback release
include_hooks	No	Boolean	Enable hooks during an update or deletion.

**Table 4-859** values

Parameter	Mandatory	Type	Description
imagePullPolicy	No	String	Image pull policy
imageTag	No	String	Image tag

## Response Parameters

Status code: 201

**Table 4-860** Response body parameters

Parameter	Type	Description
chart_name	String	Chart name
chart_public	Boolean	Whether the chart is public
chart_version	String	Chart version
cluster_id	String	Cluster ID
cluster_name	String	Current name
create_at	String	Created at
description	String	Release description
name	String	Release name
namespace	String	Namespace to which a chart release belongs
parameters	String	Release parameters
resources	String	Resources required by the release

Parameter	Type	Description
status	String	Release status. <ul style="list-style-type: none"> <li>● <b>DEPLOYED</b>: The release is normal.</li> <li>● <b>DELETED</b>: The release has been deleted.</li> <li>● <b>FAILED</b>: The release fails to be deployed.</li> <li>● <b>DELETING</b>: The release is being deleted.</li> <li>● <b>PENDING_INSTALL</b>: The release is waiting to be installed.</li> <li>● <b>PENDING_UPGRADE</b>: The release is waiting to be upgraded.</li> <li>● <b>PENDING_ROLLBACK</b>: The release is waiting for rollback.</li> <li>● <b>UNKNOWN</b>: The release status is unknown, indicating that the release is abnormal. You can manually delete the release and reinstall it.</li> </ul>
status_description	String	Release status description
update_at	String	Updated at
values	String	Release value
version	Integer	Release version

## Example Requests

```
POST /cce/cam/v3/clusters/{cluster_id}/releases
{
  "name": "nino21",
  "namespace": "project01",
  "version": "1.0.0",
  "chart_id": "3c138b72-7ce4-6d76-7c55-604cdb2ce423",
  "values": {
    "imageTag": "v2",
    "imagePullPolicy": "IfNotPresent"
  },
  "parameters": {
    "dry_run": false,
    "no_hooks": false,
    "replace": false,
    "name_template": ""
  }
}
```

## Example Responses

**Status code: 201**

Created

```
{
  "chart_name": "magento-mysql",
```

```

"chart_public" : false,
"chart_version" : "1.0.0",
"cluster_id" : "a870253f-5dc7-11ee-bf71-0255ac100b03",
"cluster_name" : "sfs-turbo-test",
"create_at" : "2023-11-14T20:30:57+08:00",
"description" : "Initial install underway",
"name" : "testwww",
"namespace" : "monitoring",
"parameters" : "",
"resources" : "",
"status" : "PENDING_INSTALL",
"status_description" : "Initial install underway",
"update_at" : "2023-11-14T20:30:57+08:00",
"values" : "{\n  \"basic\":{\n    \"admin_password\": \"*****\",\n    \"admin_username\": \"username\",\n    \"app_name\": \"magento\",\n    \"mysql_database\": \"magento\",\n    \"mysql_name\": \"mysql\",\n    \"mysql_password\": \"*****\",\n    \"mysql_port\": 3306,\n    \"mysql_root_password\": \"*****\",\n    \"mysql_user\": \"magento\",\n    \"storage_class\": \"csi-nas\",\n    \"storage_mode\": \"ReadWriteMany\",\n    \"storage_size\": \"10G\",\n    \"global\":{\n      \"magento_EIP\": \"100.100.100.100\",\n      \"magento_EPORT\": 32080,\n      \"namespace\": \"default\",\n      \"image\": \"magento_image\",\n      \"example.com/everest/magento:latest\",\n      \"mysql_image\": \"example.com/everest/mysql:5.7.14\"}}",
"version" : 1
}

```

## Status Codes

Status Code	Description
201	Created

## Error Codes

See [Error Codes](#).

## 4.9.6 Deleting a Chart

### Function

This API is used to delete a chart.

### URI

DELETE /v2/charts/{chart\_id}

**Table 4-861** Path Parameters

Parameter	Mandatory	Type	Description
chart_id	Yes	String	Chart ID

## Request Parameters

Table 4-862 Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

## Response Parameters

None

## Example Requests

None

## Example Responses

None

## Status Codes

Status Code	Description
200	OK

## Error Codes

See [Error Codes](#).

## 4.9.7 Updating a Release

### Function

This API is used to update a release.

### URI

PUT /cce/cam/v3/clusters/{cluster\_id}/namespace/{namespace}/releases/{name}



**Table 4-863** Path Parameters

Parameter	Mandatory	Type	Description
name	Yes	String	Release name
namespace	Yes	String	Namespace of the chart
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

## Request Parameters

**Table 4-864** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

**Table 4-865** Request body parameters

Parameter	Mandatory	Type	Description
chart_id	Yes	String	Chart ID
action	Yes	String	<b>upgrade</b> for an upgrade and <b>rollback</b> for a rollback
parameters	Yes	<a href="#">ReleaseReqBodyParams</a> object	Release parameters
values	Yes	<a href="#">values</a> object	Release value

**Table 4-866** ReleaseReqBodyParams

Parameter	Mandatory	Type	Description
dry_run	No	Boolean	After this function is enabled, only chart parameters are verified, and installation is not performed.
name_template	No	String	Release name template
no_hooks	No	Boolean	Whether to disable hooks during installation
replace	No	Boolean	Whether to replace the release with the same name
recreate	No	Boolean	Whether to rebuild the release
reset_values	No	Boolean	Whether to reset values during an update
release_version	No	Integer	Version of the rollback release
include_hooks	No	Boolean	Enable hooks during an update or deletion.

**Table 4-867** values

Parameter	Mandatory	Type	Description
imagePullPolicy	No	String	Image pull policy
imageTag	No	String	Image tag

## Response Parameters

**Status code: 200**

**Table 4-868** Response body parameters

Parameter	Type	Description
chart_name	String	Chart name
chart_public	Boolean	Whether the chart is public
chart_version	String	Chart version
cluster_id	String	Cluster ID

Parameter	Type	Description
cluster_name	String	Current name
create_at	String	Created at
description	String	Release description
name	String	Release name
namespace	String	Namespace to which a chart release belongs
parameters	String	Release parameters
resources	String	Resources required by the release
status	String	Release status. <ul style="list-style-type: none"> <li>● <b>DEPLOYED</b>: The release is normal.</li> <li>● <b>DELETED</b>: The release has been deleted.</li> <li>● <b>FAILED</b>: The release fails to be deployed.</li> <li>● <b>DELETING</b>: The release is being deleted.</li> <li>● <b>PENDING_INSTALL</b>: The release is waiting to be installed.</li> <li>● <b>PENDING_UPGRADE</b>: The release is waiting to be upgraded.</li> <li>● <b>PENDING_ROLLBACK</b>: The release is waiting for rollback.</li> <li>● <b>UNKNOWN</b>: The release status is unknown, indicating that the release is abnormal. You can manually delete the release and reinstall it.</li> </ul>
status_description	String	Release status description
update_at	String	Updated at
values	String	Release value
version	Integer	Release version

## Example Requests

```
PUT /cce/cam/v3/clusters/{cluster_id}/namespace/{namespace}/releases/{name}
```

```
{
  "chart_id": "af4b699e-018c-11ec-b8b0-0255ac100b05",
  "action": "upgrade",
  "parameters": {
    "dry_run": false,
    "name_template": "string",
    "no_hooks": false,
    "replace": false,
    "recreate": false,
    "reset_values": false,
    "release_version": 1,
  }
}
```

```

    "include_hooks" : false
  },
  "values" : {
    "imagePullPolicy" : "IfNotPresent",
    "imageTag" : "v2"
  }
}

```

## Example Responses

**Status code: 200**

OK

```

{
  "chart_name" : "magento-mysql",
  "chart_public" : false,
  "chart_version" : "1.0.0",
  "cluster_id" : "a870253f-5dc7-11ee-bf71-0255ac100b03",
  "cluster_name" : "sfs-turbo-test",
  "create_at" : "2023-11-14T20:30:57+08:00",
  "description" : "Initial install underway",
  "name" : "testwww",
  "namespace" : "monitoring",
  "parameters" : "",
  "resources" : "",
  "status" : "PENDING_INSTALL",
  "status_description" : "Initial install underway",
  "update_at" : "2023-11-14T20:30:57+08:00",
  "values" : "{\n\"basic\":{\n\"admin_password\":\n\"*****\",\n\"admin_username\":\n\"username\",\n\"app_name\n\":\n\"magento\",\n\"mysql_database\":\n\"magento\",\n\"mysql_name\":\n\"mysql\",\n\"mysql_password\":\n\"*****\n\",\n\"mysql_port\":3306,\n\"mysql_root_password\":\n\"*****\",\n\"mysql_user\":\n\"magento\",\n\"storage_class\":\n\"csi-nas\",\n\"storage_mode\":\n\"ReadWriteMany\",\n\"storage_size\":\n\"10G\",\n\"global\":{\n\"magento_EIP\n\":\n\"100.100.100.100\",\n\"magento_EPORT\":32080,\n\"namespace\":\n\"default\",\n\"image\":{\n\"magento_image\n\":\n\"example.com/everest/magento:latest\",\n\"mysql_image\":\n\"example.com/everest/mysql:5.7.14\"}}",
  "version" : 1
}

```

## Status Codes

Status Code	Description
200	OK

## Error Codes

See [Error Codes](#).

## 4.9.8 Obtaining a Chart

### Function

This API is used to obtain a chart.

### URI

GET /v2/charts/{chart\_id}

**Table 4-869** Path Parameters

Parameter	Mandatory	Type	Description
chart_id	Yes	String	Chart ID

## Request Parameters

**Table 4-870** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

## Response Parameters

Status code: 200

**Table 4-871** Response body parameters

Parameter	Type	Description
id	String	Chart ID
name	String	Chart name
values	String	Chart value
translate	String	Chart translation resources
instruction	String	Chart description
version	String	Chart version
description	String	Chart description
source	String	Chart source
icon_url	String	URL to chart icons
public	Boolean	Whether the chart is public
chart_url	String	URL to the chart

Parameter	Type	Description
create_at	String	Created at
update_at	String	Updated at

## Example Requests

None

## Example Responses

**Status code: 200**

OK

```
{
  "id" : "e99a7e86-afdd-11eb-aca3-0255ac100b0e",
  "name" : "neo4j",
  "values" : "{\acceptLicenseAgreement\":"no",\affinity\":{\},\authEnabled\":true,\clusterDomain\":"cluster.local",\core\":"{\initContainers\":[\],\numberOfServers\":3,\persistentVolume\":"{\enabled\":true,\mountPath\":"\/data",\size\":"10Gi"},\sidecarContainers\":[\],\defaultDatabase\":"neo4j",\image\":"neo4j",\imagePullPolicy\":"IfNotPresent",\imageTag\":"4.0.3-enterprise",\name\":"neo4j",\nodeSelector\":{\},\podDisruptionBudget\":{\},\readReplica\":"{\autoscaling\":"enabled\":false,\maxReplicas\":3,\minReplicas\":1,\targetAverageUtilization\":70},\initContainers\":[\],\numberOfServers\":0,\resources\":{\},\sidecarContainers\":[\],\resources\":{\},\testImage\":"markneedham/k8s-kubectl",\testImageTag\":"master",\tolerations\":[\],\useAPOC\":"true"}",
  "translate" : "",
  "instruction" : "README.md",
  "version" : "3.0.1",
  "description" : "DEPRECATED Neo4j is the world's leading graph database",
  "source" : "",
  "icon_url" : "https://info.neo4j.com/rs/773-GON-065/images/neo4j_logo.png",
  "public" : false,
  "chart_url" : "neo4j-3.0.1.tgz",
  "create_at" : "2021-05-08T08:53:13Z",
  "update_at" : "2021-05-08T08:53:13Z"
}
```

## Status Codes

Status Code	Description
200	OK

## Error Codes

See [Error Codes](#).

## 4.9.9 Deleting a Release

### Function

This API is used to delete a release.

## URI

DELETE /cce/cam/v3/clusters/{cluster\_id}/namespace/{namespace}/releases/{name}

**Table 4-872** Path Parameters

Parameter	Mandatory	Type	Description
name	Yes	String	Release name
namespace	Yes	String	Namespace of the chart
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

## Request Parameters

**Table 4-873** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

## Response Parameters

None

## Example Requests

None

## Example Responses

None

## Status Codes

Status Code	Description
200	OK

## Error Codes

See [Error Codes](#).

## 4.9.10 Downloading a Chart

### Function

This API is used to download a chart.

### URI

GET /v2/charts/{chart\_id}/archive

**Table 4-874** Path Parameters

Parameter	Mandatory	Type	Description
chart_id	Yes	String	Chart ID

### Request Parameters

**Table 4-875** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

### Response Parameters

Status code: 200



**Table 4-876** Response body parameters

Parameter	Type	Description
-	File	OK

## Example Requests

None

## Example Responses

Status code: 200

OK

```
"chart-file.tgz"
```

## Status Codes

Status Code	Description
200	OK

## Error Codes

See [Error Codes](#).

## 4.9.11 Obtaining a Release

### Function

This API is used to obtain a release.

### URI

GET /cce/cam/v3/clusters/{cluster\_id}/namespace/{namespace}/releases/{name}

**Table 4-877** Path Parameters

Parameter	Mandatory	Type	Description
name	Yes	String	Release name
namespace	Yes	String	Namespace of the chart
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

## Request Parameters

**Table 4-878** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

## Response Parameters

**Status code: 200**

**Table 4-879** Response body parameters

Parameter	Type	Description
chart_name	String	Chart name
chart_public	Boolean	Whether the chart is public
chart_version	String	Chart version
cluster_id	String	Cluster ID
cluster_name	String	Current name
create_at	String	Created at
description	String	Release description
name	String	Release name
namespace	String	Namespace to which a chart release belongs
parameters	String	Release parameters
resources	String	Resources required by the release

Parameter	Type	Description
status	String	Release status. <ul style="list-style-type: none"> <li>● <b>DEPLOYED</b>: The release is normal.</li> <li>● <b>DELETED</b>: The release has been deleted.</li> <li>● <b>FAILED</b>: The release fails to be deployed.</li> <li>● <b>DELETING</b>: The release is being deleted.</li> <li>● <b>PENDING_INSTALL</b>: The release is waiting to be installed.</li> <li>● <b>PENDING_UPGRADE</b>: The release is waiting to be upgraded.</li> <li>● <b>PENDING_ROLLBACK</b>: The release is waiting for rollback.</li> <li>● <b>UNKNOWN</b>: The release status is unknown, indicating that the release is abnormal. You can manually delete the release and reinstall it.</li> </ul>
status_description	String	Release status description
update_at	String	Updated at
values	String	Release value
version	Integer	Release version

### Example Requests

None

### Example Responses

**Status code: 200**

OK

```
{
  "chart_name": "magento-mysql",
  "chart_public": false,
  "chart_version": "1.0.0",
  "cluster_id": "a870253f-5dc7-11ee-bf71-0255ac100b03",
  "cluster_name": "sfs-turbo-test",
  "create_at": "2023-11-14T20:30:57+08:00",
  "description": "Initial install underway",
  "name": "testwww",
  "namespace": "monitoring",
  "parameters": "",
  "resources": "",
  "status": "PENDING_INSTALL",
  "status_description": "Initial install underway",
  "update_at": "2023-11-14T20:30:57+08:00",
  "values": "{\n  \"basic\": {\n    \"admin_password\": \"*****\",\n    \"admin_username\": \"username\",\n    \"app_name\": \"magento\",\n    \"mysql_database\": \"magento\",\n    \"mysql_name\": \"mysql\",\n    \"mysql_password\": \"*****\",\n    \"mysql_port\": 3306,\n    \"mysql_root_password\": \"*****\",\n    \"mysql_user\": \"magento\",\n    \"storage_class\": \"csi-
```

```
nas\,"storage_mode\":"ReadWriteMany\","storage_size\":"10G\","global\":{"magento_EIP
\":"100.100.100.100\","magento_EPORT\":"32080\","namespace\":"default\","image\":{"magento_image
\":"example.com/everest/magento:latest\","mysql_image\":"example.com/everest/mysql:5.7.14\}}",
"version" : 1
}
```

## Status Codes

Status Code	Description
200	OK

## Error Codes

See [Error Codes](#).

## 4.9.12 Obtaining Chart Values

### Function

This API is used to obtain chart values.

### URI

GET /v2/charts/{chart\_id}/values

**Table 4-880** Path Parameters

Parameter	Mandatory	Type	Description
chart_id	Yes	String	Chart ID

## Request Parameters

**Table 4-881** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

## Response Parameters

Status code: 200

Table 4-882 Response body parameters

Parameter	Type	Description
values	Map<String, Object>	Data in <b>values.yaml</b> . The data structure depends on the chart.

## Example Requests

None

## Example Responses

Status code: 200

OK

```
{
  "values": {
    "basic": {
      "admin_password": "*****",
      "admin_username": "username"
    },
    "global": {
      "magento_EIP": "127.0.0.1",
      "magento_EPORT": 32080,
      "namespace": "demo"
    },
    "image": {
      "magento_image": "example.com/demo/magento:latest",
      "mysql_image": "example.com/demo/mysql:5.7.14"
    }
  }
}
```

## Status Codes

Status Code	Description
200	OK

## Error Codes

See [Error Codes](#).

## 4.9.13 Obtaining Historical Records of a Release

### Function

This API is used to obtain historical records of a release.

## URI

GET /cce/cam/v3/clusters/{cluster\_id}/namespace/{namespace}/releases/{name}/history

**Table 4-883** Path Parameters

Parameter	Mandatory	Type	Description
name	Yes	String	Release name
namespace	Yes	String	Namespace of the chart
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

## Request Parameters

**Table 4-884** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

## Response Parameters

Status code: 200

**Table 4-885** Response body parameters

Parameter	Type	Description
[items]	Array of <a href="#">ReleaseResp</a> objects	OK

**Table 4-886** ReleaseResp

Parameter	Type	Description
chart_name	String	Chart name
chart_public	Boolean	Whether the chart is public
chart_version	String	Chart version
cluster_id	String	Cluster ID
cluster_name	String	Current name
create_at	String	Created at
description	String	Release description
name	String	Release name
namespace	String	Namespace to which a chart release belongs
parameters	String	Release parameters
resources	String	Resources required by the release
status	String	Release status. <ul style="list-style-type: none"> <li>● <b>DEPLOYED</b>: The release is normal.</li> <li>● <b>DELETED</b>: The release has been deleted.</li> <li>● <b>FAILED</b>: The release fails to be deployed.</li> <li>● <b>DELETING</b>: The release is being deleted.</li> <li>● <b>PENDING_INSTALL</b>: The release is waiting to be installed.</li> <li>● <b>PENDING_UPGRADE</b>: The release is waiting to be upgraded.</li> <li>● <b>PENDING_ROLLBACK</b>: The release is waiting for rollback.</li> <li>● <b>UNKNOWN</b>: The release status is unknown, indicating that the release is abnormal. You can manually delete the release and reinstall it.</li> </ul>
status_description	String	Release status description
update_at	String	Updated at
values	String	Release value
version	Integer	Release version

## Example Requests

None

## Example Responses

**Status code: 200**

OK

```
[ {
  "chart_name": "magento-mysql",
  "chart_public": false,
  "chart_version": "1.0.0",
  "cluster_id": "a870253f-5dc7-11ee-bf71-0255ac100b03",
  "cluster_name": "sfs-turbo-test",
  "create_at": "2023-11-14T20:30:57+08:00",
  "description": "Initial install underway",
  "name": "testwww",
  "namespace": "monitoring",
  "parameters": "",
  "resources": "",
  "status": "PENDING_INSTALL",
  "status_description": "Initial install underway",
  "update_at": "2023-11-14T20:30:57+08:00",
  "values": "{\n\"basic\":{\n\"admin_password\": \"*****\",\n\"admin_username\": \"username\",\n\"app_name\n\n\": \"magento\",\n\"mysql_database\": \"magento\",\n\"mysql_name\": \"mysql\",\n\"mysql_password\": \"*****\n\n\",\n\"mysql_port\": 3306,\n\"mysql_root_password\": \"*****\",\n\"mysql_user\": \"magento\",\n\"storage_class\": \"csi-nas\",\n\"storage_mode\": \"ReadWriteMany\",\n\"storage_size\": \"10G\"},\n\"global\":{\n\"magento_EIP\n\n\": \"100.100.100.100\",\n\"magento_EPORT\": 32080,\n\"namespace\": \"default\",\n\"image\":{\n\"magento_image\n\n\": \"example.com/everest/magento:latest\",\n\"mysql_image\": \"example.com/everest/mysql:5.7.14\"}}",
  "version": 1
} ]
```

## Status Codes

Status Code	Description
200	OK

## Error Codes

See [Error Codes](#).

### 4.9.14 Obtaining the Quota of a User Chart

#### Function

This API is used to obtain the quota of a user chart.

#### URI

GET /v2/charts/{project\_id}/quotas



**Table 4-887** Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see <a href="#">How to Obtain Parameters in the API URI</a> .

## Request Parameters

**Table 4-888** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format).
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a User Token</a> .

## Response Parameters

Status code: 200

**Table 4-889** Response body parameters

Parameter	Type	Description
quotas	<a href="#">quotas</a> object	Chart quota

**Table 4-890** quotas

Parameter	Type	Description
resources	Array of <a href="#">resources</a> objects	Resources

**Table 4-891** resources

Parameter	Type	Description
type	String	Resource type
quota	Integer	Resource quota
used	Integer	Used resources

## Example Requests

None

## Example Responses

**Status code: 200**

OK

```
{
  "quotas": {
    "resources": [ {
      "type": "Charts",
      "quota": 200,
      "used": 2
    } ]
  }
}
```

## Status Codes

Status Code	Description
200	OK

## Error Codes

See [Error Codes](#).

# 4.10 Add-on Instance Parameters

## 4.10.1 CoreDNS

### Add-on Overview

CoreDNS is a DNS server that uses chain plug-ins to provide domain name resolution services for Kubernetes clusters. It is the recommended DNS server solution by the Kubernetes community.

## Add-on Parameters

**Table 4-892** Parameters

Parameter	Mandatory	Type	Description
basic	No	object	Basic configuration parameters, which do not need to be specified
flavor	Yes	<a href="#">Table 4-893</a> object	Flavor parameters
custom	Yes	<a href="#">Table 4-894</a> object	Custom parameters

**Table 4-893** Configuration of flavor

Parameter	Mandatory	Type	Description
replicas	Yes	int	Number of pods. The default value is 2.
resources	Yes	Array of <a href="#">resources</a> object	Container resource (CPU and memory) quotas

**Table 4-894** Configuration of custom

Parameter	Mandatory	Type	Description
servers	No	object	Configuration of servers
stub_domains	No	Map<String>[]string	Stub domain configuration
multiAZEnable	No	bool	Whether to enable the forcible mode of multi-AZ deployment for the deployment component. The default value is <b>false</b> . Deployment pods of the add-on will be forcibly scheduled to nodes in different AZs. If there are fewer AZs than pods, the extra pods will fail to run. If both <b>multiAZEnable</b> and <b>multiAZBalance</b> are set to true, the settings of <b>multiAZBalance</b> take effect, which means, the equivalent mode of multi-AZ deployment is used.

Parameter	Mandatory	Type	Description
multiAZBalance	No	bool	Whether to enable the equivalent mode of multi-AZ deployment for the deployment component. The default value is <b>false</b> . Deployment pods of the add-on are evenly scheduled to the nodes in the cluster in each AZ. If a new AZ is added, it is recommended that you increase add-on pods for cross-AZ HA deployment. With the equivalent multi-AZ deployment, the difference between the number of add-on pods in different AZs will be less than or equal to 1. If resources in one of the AZs are insufficient, pods cannot be scheduled to that AZ.
tolerations	No	Array of <a href="#">Table 4-896</a>	Toleration configuration

**Table 4-895** Data structure of the resources field

Parameter	Mandatory	Type	Description
limitsCpu	Yes	String	CPU size limit (unit: m)
limitsMem	Yes	String	Memory size limit (unit: Mi)
name	Yes	String	Add-on name. The value is fixed at <b>coredns</b> .
requestsCpu	Yes	String	Requested CPU size (unit: m)
requestsMem	Yes	String	Requested memory size (unit: Mi)

**Table 4-896** Taints and tolerations

Parameter	Mandatory	Type	Description
key	No	String	Taint key
effect	No	String	Taint policy
operator	No	String	Operator
tolerationSeconds	No	Int	Toleration time window

**Table 4-897** Configuration of server

Parameter	Mandatory	Type	Description
port	No	Int	Domain port number. The default value is <b>5353</b> .
zones	No	Array of <a href="#">Table 4-899</a>	Domain configuration
plugins	No	Array of <a href="#">Table 4-898</a>	Configuration of plugin

**Table 4-898** Configuration of plugin

Parameter	Mandatory	Type	Description
name	Yes	String	Plugin name
configBlock	No	String	Configuration of plugin
parameters	No	String/Int	Extended parameters of the plugin

**Table 4-899** Configuration of zone

Parameter	Mandatory	Type	Description
zone	Yes	String	Domain to be listened on. The default value is a period (.).

## Example Request

```
{
  "kind": "Addon",
  "apiVersion": "v3",
  "metadata": {
    "annotations": {
      "addon.install/type": "install"
    }
  },
  "spec": {
    "clusterID": "2292498e-*****-0255ac1001ba",
    "version": "1.29.2",
    "addonTemplateName": "coredns",
    "values": {
      "basic": {
        "basickey": "val"
      },
      "flavor": {
        "replicas": 2,
        "resources": [{
          "limitsCpu": "2000m",
```

```

        "limitsMem": "2000Mi",
        "name": "coredns",
        "requestsCpu": "2000m",
        "requestsMem": "2000Mi"
    }
}],
"custom": {
    "multiAZBalance": false,
    "multiAZEnabled": false,
    "node_match_expressions": [],
    "servers": [{
        "plugins": [{
            "name": "bind",
            "parameters": "{$POD_IP}"
        }],
        {
            "configBlock": "servfail 5s",
            "name": "cache",
            "parameters": 30
        },
        {
            "name": "errors"
        },
        {
            "name": "health",
            "parameters": "{$POD_IP}:8080"
        },
        {
            "name": "ready",
            "parameters": "{$POD_IP}:8081"
        },
        {
            "configBlock": "pods insecure\nfallthrough in-addr.arpa ip6.arpa",
            "name": "kubernetes",
            "parameters": "cluster.local in-addr.arpa ip6.arpa"
        },
        {
            "name": "loadbalance",
            "parameters": "round_robin"
        },
        {
            "name": "prometheus",
            "parameters": "{$POD_IP}:9153"
        },
        {
            "configBlock": "policy random",
            "name": "forward",
            "parameters": ". /etc/resolv.conf"
        },
        {
            "name": "reload"
        }
    ]},
    "port": 5353,
    "zones": [{
        "zone": "."
    }
}],
"stub_domains": {
},
"tolerations": [{
    "key": "node.kubernetes.io/not-ready",
    "operator": "Exists",
    "effect": "NoExecute",
    "tolerationSeconds": 60
},
{
    "key": "node.kubernetes.io/unreachable",
    "operator": "Exists",

```

```

    "effect": "NoExecute",
    "tolerationSeconds": 60
  }}
}
}
}
}

```

## 4.10.2 CCE Container Storage (Everest)

### Add-on Overview

The CCE Container Storage (Everest) add-on uses the Container Storage Interface (CSI) to link Kubernetes clusters with cloud storage services.

### Add-on Parameters

**Table 4-900** Parameters

Parameter	Mandatory	Type	Description
basic	No	object	Basic configuration parameters, which do not need to be specified
flavor	Yes	<a href="#">Table 4-901</a> object	Flavor parameters
custom	Yes	<a href="#">Table 4-902</a> object	Custom parameters

**Table 4-901** Configuration of flavor

Parameter	Mandatory	Type	Description
description	No	String	Add-on description
name	Yes	String	Add-on specification name
replicas	Yes	String	Number of pods. The default value is <b>1</b> .
resources	Yes	<a href="#">resources</a> object	Container resource (CPU and memory) quotas

**Table 4-902** Configuration of custom

Parameter	Mandatory	Type	Description
default_vpc_id	Yes	String	VPC ID

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	Cluster ID
cluster_name	Yes	String	Cluster name
project_id	Yes	String	Project ID
disable_auto_mount_secret	No	bool	Whether to allow the default AK/SK to be used during OBS volume mounting Default value: <b>false</b>
over_subscription	No	String	Local PV overcommitment ratio Default value: <b>80</b>
csi_attach_worker_threads	No	String	The number of concurrent workers for detaching volumes Default value: <b>60</b>
volume_attaching_flow_ctrl	No	String	Attaching flow control data Default value: <b>0</b> .
number_of_reserved_disks	No	String	Disk attaching capability reserved for non-container scenarios Default value: <b>6</b>
flow_control	No	Map<String>string	Flow control parameter Default value: <b>{}</b>
enable_node_attacher	No	bool	Whether to start agent attacher Default value: <b>false</b>
multiAZEnabled	No	bool	Whether to enable the forcible mode of multi-AZ deployment for the deployment component. The default value is <b>false</b> . Deployment pods of the add-on will be forcibly scheduled to nodes in different AZs. If there are fewer AZs than pods, the extra pods will fail to run. If both <b>multiAZEnable</b> and <b>multiAZBalance</b> are set to true, the settings of <b>multiAZBalance</b> take effect, which means, the equivalent mode of multi-AZ deployment is used.



Parameter	Mandatory	Type	Description
multiAZBalance	No	bool	Whether to enable the equivalent mode of multi-AZ deployment for the deployment component. The default value is <b>false</b> . Deployment pods of the add-on are evenly scheduled to the nodes in the cluster in each AZ. If a new AZ is added, it is recommended that you increase add-on pods for cross-AZ HA deployment. With the equivalent multi-AZ deployment, the difference between the number of add-on pods in different AZs will be less than or equal to 1. If resources in one of the AZs are insufficient, pods cannot be scheduled to that AZ.
tolerations	No	Array of <a href="#">Table 4-904</a>	Toleration configuration
node_match_expressions	No	Array of <a href="#">Table 4-905</a>	Add-on pod affinity configuration

**Table 4-903** Data structure of the resources field

Parameter	Mandatory	Type	Description
limitsCpu	Yes	String	CPU size limit (unit: m)
limitsMem	Yes	String	Memory size limit (unit: Mi)
name	Yes	String	Add-on name. The value is fixed at <b>everest-csi-controller</b> or <b>everest-csi-driver</b> .
requestsCpu	Yes	String	Requested CPU size (unit: m)
requestsMem	Yes	String	Requested memory size (unit: Mi)

**Table 4-904** Taints and tolerations

Parameter	Mandatory	Type	Description
key	No	String	Taint key
effect	No	String	Taint policy

Parameter	Mandatory	Type	Description
operator	No	String	Operator
tolerationSeconds	No	Int	Toleration time window

**Table 4-905** nodeMatchExpression node affinity

Parameter	Mandatory	Type	Description
key	No	String	Taint key
values	No	List<String >	Node affinity name
operator	No	String	Operator

## Example Request

```
{
  "kind": "Addon",
  "apiVersion": "v3",
  "metadata": { "annotations": { "addon.install/type": "install" } },
  "spec": {
    "clusterID": "cea63ae5-df9b-11ee-9f27-0255ac1001b2",
    "version": "2.4.4",
    "addonTemplateName": "everest",
    "values": {
      "basic": {
        "bms_url": "bms.***.com",
        "driver_init_image_version": "2.4.4",
        "ecsEndpoint": "https://ecs.***.com",
        "everest_image_version": "2.4.4",
        "evs_url": "evs.***.com",
        "iam_url": "iam.***.com",
        "ims_url": "ims.***.com",
        "obs_url": "obs.***.com",
        "platform": "linux-amd64",
        "sfs30_url": "obs.***.com",
        "sfs_turbo_url": "sfs-turbo.***.com",
        "sfs_url": "sfs.***.com",
        "supportHcs": false,
        "swr_addr": "swr.***.com",
        "swr_user": "hwofficial",
        "rbac_enabled": true,
        "cluster_version": "v1.23"
      },
      "flavor": {
        "description": "High available",
        "is_default": true,
        "name": "HA50",
        "recommend_cluster_flavor_types": ["small"],
        "replicas": 2,
        "resources": [
          {
            "limitsCpu": "250m",
            "limitsMem": "600Mi",
            "name": "everest-csi-controller",
            "requestsCpu": "250m",
            "requestsMem": "600Mi"
          }
        ]
      }
    }
  }
}
```

```
{
  "limitsCpu": "300m",
  "limitsMem": "300Mi",
  "name": "everest-csi-driver",
  "requestsCpu": "300m",
  "requestsMem": "300Mi"
},
"category": ["CCE", "Turbo"]
},
"custom": {
  "annotations": {},
  "cluster_id": "",
  "cluster_name": "",
  "csi_attacher_detach_worker_threads": "60",
  "csi_attacher_worker_threads": "60",
  "default_vpc_id": "",
  "disable_auto_mount_secret": false,
  "enable_node_attacher": false,
  "flow_control": {},
  "multiAZBalance": false,
  "multiAZEnabled": false,
  "node_match_expressions": [],
  "number_of_reserved_disks": "6",
  "over_subscription": "80",
  "project_id": "",
  "tolerations": [
    {
      "key": "node.kubernetes.io/not-ready",
      "operator": "Exists",
      "effect": "NoExecute",
      "tolerationSeconds": 60
    },
    {
      "key": "node.kubernetes.io/unreachable",
      "operator": "Exists",
      "effect": "NoExecute",
      "tolerationSeconds": 60
    }
  ],
  "volume_attaching_flow_ctrl": "0"
}
}
```

### 4.10.3 CCE Node Problem Detector

#### Add-on Overview

CCE Node Problem Detector (node-problem-detector, NPD) is an add-on that monitors abnormal events of cluster nodes and connects to a third-party monitoring platform. It is a daemon running on each node. It collects node issues from different daemons and reports them to the API server. It can run as a DaemonSet or a daemon.

## Add-on Parameters

**Table 4-906** Parameters

Parameter	Mandatory	Type	Description
basic	No	object	Basic configuration parameters, which do not need to be specified
flavor	Yes	<a href="#">Table 4-907</a> object	Flavor parameters
custom	Yes	<a href="#">Table 4-908</a> object	Custom parameters

**Table 4-907** Configuration of flavor

Parameter	Mandatory	Type	Description
description	No	String	Add-on description
name	Yes	String	Add-on specification name. The value is fixed at <b>Single-instance</b> .
replicas	Yes	String	Number of pods. The default value is <b>1</b> .
resources	Yes	<a href="#">resources</a> object	Container resource (CPU and memory) quotas

**Table 4-908** Configuration of custom

Parameter	Mandatory	Type	Description
feature_gate	No	String	Feature gate, which is used to enable the beta features
multiAZBalance	No	Bool	Multi AZ deployment
multiAZEnabled	No	Bool	Whether to deploy the add-on pods in multiple AZs. The default value is <b>false</b> . If this parameter is set to <b>true</b> , cross-AZ deployment is forcibly performed. If this parameter is set to <b>false</b> , cross-AZ deployment is preferred.

Parameter	Mandatory	Type	Description
npc	Yes	object <a href="#">Table 4-910</a>	node-problem-controller configuration
tolerations	No	List<Object> <a href="#">Table 4-912</a>	Tolerations of the add-on
node_match_expressions	No	List<Object> <a href="#">Table 4-912</a>	Node affinity configuration of the add-on

**Table 4-909** Data structure of the resources field

Parameter	Mandatory	Type	Description
limitsCpu	Yes	String	CPU size limit (unit: m)
limitsMem	Yes	String	Memory size limit (unit: Mi)
name	Yes	String	Add-on name. The value is fixed at <b>custom-resources</b> .
requestsCpu	Yes	String	Requested CPU size (unit: m)
requestsMem	Yes	String	Requested memory size (unit: Mi)

**Table 4-910** Data structure of the npc field

Parameter	Mandatory	Type	Description
maxTaintedNode	Yes	String or Int	The maximum number of nodes that NPC can add taints to when a single fault occurs on multiple nodes for minimizing impact. The value can be in int or percentage format.

**Table 4-911** Taints and tolerations

Parameter	Mandatory	Type	Description
key	No	String	Taint key
effect	No	String	Taint policy

Parameter	Mandatory	Type	Description
operator	No	String	Operator
tolerationSeconds	No	Int	Toleration time window

**Table 4-912** nodeMatchExpression node affinity

Parameter	Mandatory	Type	Description
key	No	String	Taint key
values	No	List<String>	Node affinity name
operator	No	String	Operator

## Example Request

```
{
  "kind": "Addon",
  "apiVersion": "v3",
  "metadata": {
    "annotations": {
      "addon.install/type": "install"
    }
  },
  "spec": {
    "clusterID": "b78fb690-b82c-11ee-83cf-0255ac100b0f",
    "version": "1.18.48",
    "addonTemplateName": "npd",
    "values": {
      "basic": {
        "image_version": "1.18.48",
        "swr_addr": "****",
        "swr_user": "****",
        "rbac_enabled": true,
        "cluster_version": "v1.23"
      }
    },
    "flavor": {
      "description": "custom resources",
      "name": "custom-resources",
      "replicas": 2,
      "resources": [
        {
          "limitsCpu": "100m",
          "limitsMem": "300Mi",
          "name": "node-problem-controller",
          "requestsCpu": "30m",
          "requestsMem": "100Mi"
        },
        {
          "limitsCpu": "100m",
          "limitsMem": "300Mi",
          "name": "node-problem-detector",
          "requestsCpu": "30m",
          "requestsMem": "100Mi"
        }
      ]
    },
    "category": [
```



**Table 4-914** Configuration of flavor

Parameter	Mandatory	Type	Description
replicas	Yes	String	Number of pods. The default value is <b>1</b> .
resources	Yes	Array of <b>resources</b> object	Container resource (CPU and memory) quotas

**Table 4-915** Configuration of custom

Parameter	Mandatory	Type	Description
serviceType	Yes	String	External access type. The value can be <b>NodePort</b> or <b>ELB</b> and defaults to <b>NodePort</b> .
port	No	int	Port number of the kubernetes-dashboard Service. The default value is <b>8443</b> .
loadBalancerIP	No	String	IP address of the ELB when external access type is ELB.
elbClass	No	String	ELB type when external access type is ELB. The value can be <b>union</b> (shared load balancer) or <b>performance</b> (dedicated load balancer) and defaults to <b>union</b> .
elbID	No	String	ID of the ELB when external access type is ELB.
certUploaded	No	bool	Whether to use a custom certificate. The default value is <b>true</b> .
cert	No	String	Cert of a custom certificate
key	No	String	Key of a custom certificate

**Table 4-916** Data structure of the resources field

Parameter	Mandatory	Type	Description
limitsCpu	Yes	String	CPU size limit (unit: m)
limitsMem	Yes	String	Memory size limit (unit: Mi)
name	Yes	String	Add-on name. The value is fixed at <b>dashboard</b> .



Parameter	Mandatory	Type	Description
requestsCpu	Yes	String	Requested CPU size (unit: m)
requestsMem	Yes	String	Requested memory size (unit: Mi)

## Example Request

```
{
  "kind": "Addon",
  "apiVersion": "v3",
  "metadata": {
    "annotations": {
      "addon.install/type": "install"
    }
  },
  "spec": {
    "clusterID": "2292498e-d169-**-9151-***1001ba",
    "version": "2.2.27",
    "addonTemplateName": "dashboard",
    "values": {
      "basic": {
        "basickey": "val"
      },
      "flavor": {

        "replicas": 1,
        "resources": [{
          "limitsCpu": "100m",
          "limitsMem": "512Mi",
          "name": "dashboard",
          "requestsCpu": "50m",
          "requestsMem": "256Mi"
        }]
      }
    },
    "custom": {
      "cert": "****",
      "certUploaded": true,
      "elbClass": "union",
      "elbID": 0,
      "key": "****",
      "loadBalancerIP": "",
      "port": 8443,
      "serviceType": "NodePort",
      "cluster_id": "2292498e-d169-**-9151-***1001ba",
      "tenant_id": "*****"
    }
  }
}
```

## 4.10.5 CCE Cluster Autoscaler

### Add-on Overview

The CCE Cluster Autoscaler (autoscaler) add-on enables node pools to be scaled in or out.

## Add-on Parameters

**Table 4-917** Parameters

Parameter	Mandatory	Type	Description
basic	No	object	Basic configuration parameters, which do not need to be specified
flavor	Yes	<a href="#">Table 4-918</a> object	Flavor parameters
custom	Yes	<a href="#">Table 4-919</a> object	Custom parameters

**Table 4-918** flavor

Parameter	Mandatory	Type	Description
replicas	Yes	String	Number of pods. The default value is <b>2</b> .
resources	Yes	Array of <a href="#">resources</a> object	Container resource (CPU and memory) quotas

**Table 4-919** custom

Parameter	Mandatory	Type	Description
cluster_id	Yes	string	Cluster ID.
tenant_id	Yes	string	ID of the project where the current cluster is in.
scaleDownEnabled	No	bool	Whether to enable auto scale-in. The default value is <b>false</b> .
scaleDownDelayAfterAdd	No	int	Cooldown period (in minutes) for starting scale-in evaluation again after auto scale-out was triggered in a cluster. The default value is <b>10</b> .
scaleDownDelayAfterDelete	No	int	Cooldown period (in minutes) for starting scale-in evaluation again after auto scale-in was triggered in a cluster. The default value is <b>10</b> .

Parameter	Mandatory	Type	Description
scaleDownDelayAfterFailure	No	int	Cooldown period (in minutes) for starting scale-in evaluation again after auto scale-in triggered by a cluster failed. The default value is <b>10</b> .
maxEmptyBulkDeleteFlag	No	int	Number of idle nodes that can be concurrently scaled in. The default value is <b>10</b> .
unremovableNodeRecheckTimeout	No	int	Interval (in minutes) for starting the checks again after a node is determined not to be scaled in. The default value is <b>5</b> .
scaleDownUtilizationThreshold	No	double	CPU and memory usage thresholds for determining whether a node can be scaled in. The default value is <b>0.5</b> .
maxNodesTotal	No	int	Maximum number of nodes that can be added to a cluster. The default value is <b>1000</b> .
coresTotal	No	int	Maximum number of CPU cores that can be added to a cluster. The default value is <b>32000</b> .
memoryTotal	No	int	Memory upper limit (in Gi) for cluster scale-out. The default value is <b>128000</b> .
scaleUpUtilizationEnabled	No	bool	Whether to enable custom scaling. The default value is <b>true</b> .
scaleUpUnscheduledPodEnabled	No	bool	Whether to enable automatic scale-out for unscheduled pods. The default value is <b>true</b> .
ignoreDaemonSetsUtilization	No	bool	Whether to ignore DaemonSets' resource usage when thresholds determine if scale-in is performed. The default value is <b>false</b> .
skipNodesWithCustomControllerPods	No	bool	Scale-in is not performed on the node which has containers created by a third-party controller running on it. The default value is <b>true</b> .
logLevel	No	int	Log level. The default value is <b>4</b> .

Parameter	Mandatory	Type	Description
multiAZEnable	No	bool	Whether to enable the forcible mode of multi-AZ deployment for the deployment component. The default value is <b>false</b> . Deployment pods of the add-on will be forcibly scheduled to nodes in different AZs. If there are fewer AZs than pods, the extra pods will fail to run. If both <b>multiAZEnable</b> and <b>multiAZBalance</b> are set to true, the settings of <b>multiAZBalance</b> take effect, which means, the equivalent mode of multi-AZ deployment is used.
multiAZBalance	No	bool	Whether to enable the equivalent mode of multi-AZ deployment for the deployment component. The default value is <b>false</b> . Deployment pods of the add-on are evenly scheduled to the nodes in the cluster in each AZ. If a new AZ is added, it is recommended that you increase add-on pods for cross-AZ HA deployment. With the equivalent multi-AZ deployment, the difference between the number of add-on pods in different AZs will be less than or equal to 1. If resources in one of the AZs are insufficient, pods cannot be scheduled to that AZ.
tolerations	No	Array of <a href="#">Table 4-921</a>	Toleration configuration
node_match_expressions	No	Array of <a href="#">Table 4-922</a>	Add-on pod affinity configuration

**Table 4-920** Data structure of the resources field

Parameter	Mandatory	Type	Description
limitsCpu	Yes	String	CPU size limit (unit: m)
limitsMem	Yes	String	Memory size limit (unit: Mi)
name	Yes	String	Add-on name. The value is fixed at <b>autoscaler</b> .

Parameter	Mandatory	Type	Description
requestsCpu	Yes	String	Requested CPU size (unit: m)
requestsMem	Yes	String	Requested memory size (unit: Mi)

**Table 4-921** Taints and tolerations

Parameter	Mandatory	Type	Description
key	No	String	Taint key
effect	No	String	Taint policy
operator	No	String	Operator
tolerationSeconds	No	Int	Toleration time window

**Table 4-922** nodeMatchExpression node affinity

Parameter	Mandatory	Type	Description
key	No	String	Taint key
values	No	List<String>	Node affinity name
operator	No	String	Operator

## Example Request

```
{
  "kind": "Addon",
  "apiVersion": "v3",
  "metadata": {
    "annotations": {
      "addon.install/type": "install"
    }
  },
  "spec": {
    "clusterID": "2292498e-*****-0255ac1001ba",
    "version": "1.23.116",
    "addonTemplateName": "autoscaler",
    "values": {
      "basic": {
        "basickey": "val"
      },
      "flavor": {
        "replicas": 2,
        "resources": [{
          "limitsCpu": "1000m",
          "limitsMem": "1000Mi",
          "name": "autoscaler",
          "requestsCpu": "1000m",

```

```

        "requestsMem": "1000Mi"
    }
  },
  "custom": {
    "cluster_id": "2292498e-*****-0255ac1001ba",
    "coresTotal": 32000,
    "ignoreDaemonSetsUtilization": false,
    "logLevel": 4,
    "maxEmptyBulkDeleteFlag": 10,
    "maxNodeProvisionTime": 15,
    "maxNodesTotal": 1000,
    "memoryTotal": 128000,
    "multiAZBalance": false,
    "multiAZEnabled": false,
    "node_match_expressions": [],
    "scaleDownDelayAfterAdd": 10,
    "scaleDownDelayAfterDelete": 10,
    "scaleDownDelayAfterFailure": 3,
    "scaleDownEnabled": false,
    "scaleDownUnneededTime": 10,
    "scaleDownUtilizationThreshold": 0.5,
    "scaleUpUnscheduledPodEnabled": true,
    "scaleUpUtilizationEnabled": true,
    "skipNodesWithCustomControllerPods": true,
    "tenant_id": "*****",
    "tolerations": [{
      "key": "node.kubernetes.io/not-ready",
      "operator": "Exists",
      "effect": "NoExecute",
      "tolerationSeconds": 60
    },
    {
      "key": "node.kubernetes.io/unreachable",
      "operator": "Exists",
      "effect": "NoExecute",
      "tolerationSeconds": 60
    }
  ],
    "unremovableNodeRecheckTimeout": 5
  }
}
}
}

```

## 4.10.6 NGINX Ingress Controller

### Add-on Overview

The NGINX Ingress Controller add-on can automatically modify configurations when there are changes to pods in Services. This add-on uses Nginx's excellent stability, performance, and concurrent processing ability to effectively manage containerized applications at the application layer.

### Add-on Parameters

**Table 4-923** Parameters

Parameter	Mandatory	Type	Description
basic	No	object	Basic configuration parameters, which do not need to be specified

Parameter	Mandatory	Type	Description
flavor	Yes	<a href="#">Table 4-924</a> object	Flavor parameters
custom	Yes	<a href="#">Table 4-925</a> object	Custom parameters

**Table 4-924** flavor

Parameter	Mandatory	Type	Description
replicas	Yes	String	Number of pods. The default value is <b>1</b> .
resources	Yes	<a href="#">resources</a> object	Container resource (CPU and memory) quotas

**Table 4-925** custom

Parameter	Mandatory	Type	Description
ingressClass	Yes	String	Controller name. The default value is <b>nginx</b> .
namespace	Yes	String	Namespace where the add-on is in. The default value is <b>kube-system</b> .
service	Yes	<a href="#">Table 4-930</a>	Configuration of a Service that provides external access
config	No	Map<String>String	Nginx configuration parameters. For details, see <a href="#">ConfigMaps</a> .
admissionWebhooks	No	<a href="#">Table 4-929</a>	Configuration of ingress admission verification.
metrics	No	<a href="#">Table 4-931</a>	Monitoring metric configuration.
defaultBackendService	No	String	Default 404 service, which is in the format of <i>&lt;namespace&gt;/&lt;service_name&gt;</i> .
extraArgs	No	<a href="#">Table 4-932</a>	Extended parameter configuration.

Parameter	Mandatory	Type	Description
multiAZEnable	No	bool	Whether to enable the forcible mode of multi-AZ deployment for the deployment component. The default value is <b>false</b> . Deployment pods of the add-on will be forcibly scheduled to nodes in different AZs. If there are fewer AZs than pods, the extra pods will fail to run. If both <b>multiAZEnable</b> and <b>multiAZBalance</b> are set to true, the settings of <b>multiAZBalance</b> take effect, which means, the equivalent mode of multi-AZ deployment is used.
multiAZBalance	No	bool	Whether to enable the equivalent mode of multi-AZ deployment for the deployment component. The default value is <b>false</b> . Deployment pods of the add-on are evenly scheduled to the nodes in the cluster in each AZ. If a new AZ is added, it is recommended that you increase add-on pods for cross-AZ HA deployment. With the equivalent multi-AZ deployment, the difference between the number of add-on pods in different AZs will be less than or equal to 1. If resources in one of the AZs are insufficient, pods cannot be scheduled to that AZ.
tolerations	No	Array of <a href="#">Table 4-927</a>	Toleration configuration
node_match_expressions	No	Array of <a href="#">Table 4-928</a>	Add-on pod affinity configuration

**Table 4-926** Data structure of the resources field

Parameter	Mandatory	Type	Description
limitsCpu	Yes	String	CPU size limit (unit: m)
limitsMem	Yes	String	Memory size limit (unit: Mi)
name	Yes	String	Add-on name. The value is fixed at <b>nginx-ingress</b> .
requestsCpu	Yes	String	Requested CPU size (unit: m)



Parameter	Mandatory	Type	Description
requestsMemory	Yes	String	Requested memory size (unit: Mi)

**Table 4-927** Taints and tolerations

Parameter	Mandatory	Type	Description
key	No	String	Taint key
effect	No	String	Taint policy
operator	No	String	Operator
tolerationSeconds	No	Int	Toleration time window

**Table 4-928** nodeMatchExpression node affinity

Parameter	Mandatory	Type	Description
key	No	String	Taint key
values	No	List<String>	Node affinity name
operator	No	String	Operator

**Table 4-929** admissionWebhook

Parameter	Mandatory	Type	Description
enable	No	bool	Whether to enable ingress resource admission verification. The default value is <b>true</b> .

**Table 4-930** service

Parameter	Mandatory	Type	Description
annotations	No	Map<String>String	Annotations of a Service that provides external access. This parameter is only for configurations of ELB-related options, such as <b>kubernetes.io/elb.class</b> , <b>kubernetes.io/elb.id</b> , and <b>kubernetes.io/elb.pass-through</b> .

Parameter	Mandatory	Type	Description
loadBalancerIP	No	String	Load balancer IP address used by the Service that is interconnected with the load balancer.

**Table 4-931** metrics

Parameter	Mandatory	Type	Description
enable	No	bool	Whether to monitor metrics. The default value is <b>true</b> .
excludeSocketMetrics	No	String	Shielded monitoring metrics. The default value is <b>"nginx_ingress_controller_success,nginx_ingress_controller_header_duration_seconds,nginx_ingress_controller_ingress_upstream_latency_seconds"</b> .

**Table 4-932** extraArg extended parameter

Parameter	Mandatory	Type	Description
default-ssl-certificate	No	String	Default certificate configuration. For details, see <a href="#">Default SSL Certificate</a> .

## Example Request

```
{
  "kind": "Addon",
  "apiVersion": "v3",
  "metadata": {
    "annotations": {
      "addon.install/type": "install"
    }
  },
  "spec": {
    "clusterID": "2292498e-*****ac1001ba",
    "version": "2.2.52",
    "addonTemplateName": "nginx-ingress",
    "values": {
      "basic": {
        "basickey": "val"
      },
      "flavor": {
        "replicas": 2,
        "resources": [{
          "limitsCpu": "8000m",
          "limitsMem": "4000Mi",
          "name": "nginx-ingress",
          "requestsCpu": "8000m",
          "requestsMem": "4000Mi"
        }]
      }
    }
  }
}
```

```

    },
    "custom": {
      "config": {
        "keep-alive-requests": "100"
      },
      "defaultBackend": {
        "enabled": true
      },
      "defaultBackendService": "",
      "extraArgs": {
        "default-ssl-certificate": ""
      },
      "ingressClass": "nginx",
      "multiAZBalance": false,
      "multiAZEnabled": false,
      "namespace": "kube-system",
      "node_match_expressions": [],
      "service": {
        "annotations": {
          "kubernetes.io/elb.class": "performance",
          "kubernetes.io/elb.id": "8d6bd485-d8ac-4693-815d-9d54d79b0666"
        },
        "loadBalancerIP": ""
      },
      "tolerations": [{
        "key": "node.kubernetes.io/not-ready",
        "operator": "Exists",
        "effect": "NoExecute",
        "tolerationSeconds": 60
      },
      {
        "key": "node.kubernetes.io/unreachable",
        "operator": "Exists",
        "effect": "NoExecute",
        "tolerationSeconds": 60
      }
    ]
  }
}

```

## 4.10.7 Kubernetes Metrics Server

### Add-on Overview

From version 1.8 onwards, Kubernetes provides resource usage metrics, such as the container CPU and memory usage, through the Metrics API. These metrics can be directly accessed by users (for example, by using the **kubectl top** command) or used by controllers (for example, Horizontal Pod Autoscaler) in a cluster for decision-making. The specific component is metrics-server, which is used to substitute for Heapster for providing the similar functions. Heapster has been gradually abandoned since v1.11.

### Add-on Parameters

**Table 4-933** Parameters

Parameter	Mandatory	Type	Description
basic	No	object	Basic configuration parameters, which do not need to be specified

Parameter	Mandatory	Type	Description
flavor	Yes	<a href="#">Table 4-934</a> object	Flavor parameters
custom	Yes	<a href="#">Table 4-935</a> object	Custom parameters

**Table 4-934** flavor

Parameter	Mandatory	Type	Description
description	No	String	Add-on description
name	Yes	String	Add-on specification name. The value is fixed at <b>Single-instance</b> .
replicas	Yes	String	Number of pods. The default value is <b>1</b> .
resources	Yes	<a href="#">resources</a> object	Container resource (CPU and memory) quotas

**Table 4-935** custom

Parameter	Mandatory	Type	Description
multiAZBalance	No	Bool	Whether the equivalent mode of multi-AZ deployment is enabled. If this function is enabled, the equivalent mode is used.
multiAZEnabled	No	Bool	Whether to deploy the add-on pods in multiple AZs. The default value is <b>false</b> . If this parameter is set to <b>true</b> , cross-AZ deployment is forcibly performed. If this parameter is set to <b>false</b> , cross-AZ deployment is preferred.
tolerations	No	List<Object> <a href="#">Table 4-938</a>	Tolerations of the add-on
node_match_expressions	No	List<Object> <a href="#">Table 4-938</a>	Node affinity configuration of the add-on

**Table 4-936** Data structure of the resources field

Parameter	Mandatory	Type	Description
limitsCpu	Yes	String	CPU size limit (unit: m)
limitsMem	Yes	String	Memory size limit (unit: Mi)
name	Yes	String	Add-on name. The value is fixed at <b>metrics-server</b> .
requestsCpu	Yes	String	Requested CPU size (unit: m)
requestsMem	Yes	String	Requested memory size (unit: Mi)

**Table 4-937** Taints and tolerations

Parameter	Mandatory	Type	Description
key	No	String	Taint key
effect	No	String	Taint policy
operator	No	String	Operator
tolerationSeconds	No	Int	Toleration time window

**Table 4-938** nodeMatchExpression node affinity

Parameter	Mandatory	Type	Description
key	No	String	Taint key
values	No	List<String>	Node affinity name
operator	No	String	Operator

## Example Request

```
{
  "kind": "Addon",
  "apiVersion": "v3",
  "metadata": {
    "annotations": {
      "addon.install/type": "install"
    }
  },
  "spec": {
    "clusterID": "b78fb690-b82c-11ee-83cf-0255ac100b0f",
    "version": "1.3.39",
  }
}
```

```
"addonTemplateName": "metrics-server",
"values": {
  "basic": {
    "image_version": "v0.6.2",
    "swr_addr": "****",
    "swr_user": "****",
    "rbac_enabled": true,
    "cluster_version": "v1.23"
  },
  "flavor": {
    "description": "Has only one instance",
    "name": "Single",
    "replicas": 1,
    "resources": [
      {
        "limitsCpu": "1000m",
        "limitsMem": "1000Mi",
        "name": "metrics-server",
        "requestsCpu": "100m",
        "requestsMem": "300Mi"
      }
    ],
    "category": [
      "CCE",
      "Turbo"
    ]
  },
  "custom": {
    "annotations": {},
    "multiAZBalance": false,
    "multiAZEnabled": false,
    "node_match_expressions": [],
    "tolerations": [
      {
        "key": "node.kubernetes.io/not-ready",
        "operator": "Exists",
        "effect": "NoExecute",
        "tolerationSeconds": 60
      },
      {
        "key": "node.kubernetes.io/unreachable",
        "operator": "Exists",
        "effect": "NoExecute",
        "tolerationSeconds": 60
      }
    ]
  }
}
}
```

## 4.10.8 CCE Advanced HPA

### Add-on Overview

CCE Advanced HPA (cce-hpa-controller) is an in-house add-on, which can be used to flexibly scale in or out Deployments based on metrics such as CPU usage and memory usage.

## Add-on Parameters

**Table 4-939** Parameters

Parameter	Mandatory	Type	Description
basic	No	object	Basic configuration parameters, which do not need to be specified
flavor	Yes	<a href="#">Table 4-940</a> object	Flavor parameters
custom	Yes	<a href="#">Table 4-941</a> object	Custom parameters

**Table 4-940** flavor

Parameter	Mandatory	Type	Description
replicas	Yes	String	Number of pods. The default value is <b>2</b> .
resources	Yes	<a href="#">resources</a> object	Container resource (CPU and memory) quotas.

**Table 4-941** custom

Parameter	Mandatory	Type	Description
multiAZEnable	No	bool	Whether to enable the forcible mode of multi-AZ deployment for the deployment component. The default value is <b>false</b> . Deployment pods of the add-on will be forcibly scheduled to nodes in different AZs. If there are fewer AZs than pods, the extra pods will fail to run. If both <b>multiAZEnable</b> and <b>multiAZBalance</b> are set to true, the settings of <b>multiAZBalance</b> take effect, which means, the equivalent mode of multi-AZ deployment is used.

Parameter	Mandatory	Type	Description
multiAZBalance	No	bool	Whether to enable the equivalent mode of multi-AZ deployment for the deployment component. The default value is <b>false</b> . Deployment pods of the add-on are evenly scheduled to the nodes in the cluster in each AZ. If a new AZ is added, it is recommended that you increase add-on pods for cross-AZ HA deployment. With the equivalent multi-AZ deployment, the difference between the number of add-on pods in different AZs will be less than or equal to 1. If resources in one of the AZs are insufficient, pods cannot be scheduled to that AZ.
tolerations	No	Array of <a href="#">Table 4-943</a>	Toleration configuration
node_match_expressions	No	Array of <a href="#">Table 4-944</a>	Add-on pod affinity configuration

**Table 4-942** Data structure of the resources field

Parameter	Mandatory	Type	Description
limitsCpu	Yes	String	CPU size limit (unit: m)
limitsMem	Yes	String	Memory size limit (unit: Mi)
name	Yes	String	Add-on name. The value is fixed at <b>customedhpa-controller</b> .
requestsCpu	Yes	String	Requested CPU size (unit: m)
requestsMem	Yes	String	Requested memory size (unit: Mi)

**Table 4-943** Taints and tolerations

Parameter	Mandatory	Type	Description
key	No	String	Taint key
effect	No	String	Taint policy
operator	No	String	Operator



Parameter	Mandatory	Type	Description
tolerationSeconds	No	Int	Toleration time window

**Table 4-944** nodeMatchExpression node affinity

Parameter	Mandatory	Type	Description
key	No	String	Taint key
values	No	List<String >	Node affinity name
operator	No	String	Operator

## Example Request

```
{
  "kind": "Addon",
  "apiVersion": "v3",
  "metadata": {
    "annotations": {
      "addon.install/type": "install"
    }
  },
  "spec": {
    "clusterID": "2292498e-*****-0255ac1001ba",
    "version": "1.4.2",
    "addonTemplateName": "cce-hpa-controller",
    "values": {
      "basic": {
        "basickey": "val"
      },
      "flavor": {
        "replicas": 1,
        "resources": [{
          "limitsCpu": "100m",
          "limitsMem": "300Mi",
          "name": "customedhpa-controller",
          "requestsCpu": "100m",
          "requestsMem": "300Mi"
        }]
      },
      "custom": {
        "multiAZBalance": false,
        "multiAZEnabled": false,
        "node_match_expressions": [],
        "tolerations": [{
          "key": "node.kubernetes.io/not-ready",
          "operator": "Exists",
          "effect": "NoExecute",
          "tolerationSeconds": 60
        },
        {
          "key": "node.kubernetes.io/unreachable",
          "operator": "Exists",
          "effect": "NoExecute",
          "tolerationSeconds": 60
        }
      ]
    }
  }
}
```

```
}  
}
```

## 4.10.9 CCE AI Suite (NVIDIA GPU)

### Add-on Overview

CCE AI Suite (NVIDIA GPU) is a device management add-on that supports GPUs in containers. To use GPU nodes in a cluster, this add-on must be installed.

### Add-on Parameters

**Table 4-945** Parameters

Parameter	Mandatory	Type	Description
basic	Yes	object	Basic add-on configuration parameters
custom	Yes	<a href="#">Table 4-947</a> object	Custom parameters

**Table 4-946** Configuration of basic

Parameter	Mandatory	Type	Description
cluster_version	No	String	CCE cluster version
device_version	Yes	String	Add-on version
driver_version	Yes	String	Image tag of an add-on pod where a driver is installed. Generally, the value is the same as that of <b>device_version</b> .
obs_url	Yes	String	When a GPU driver is downloaded from the default driver address, the value is the GPU driver address.
swr_addr	Yes	String	Image repository address

Parameter	Mandatory	Type	Description
swr_user	Yes	String	Tenant path of an image repository

**Table 4-947** Configuration of custom

Parameter	Mandatory	Type	Description
compatible_with_legacy_api	No	Bool	API compatibility switch Default value: <b>false</b> <b>true</b> : The add-on supports the GPU native mode and xGPU virtualization.
component_schedulername	Yes	String	Name of the scheduler used by the add-on. Default value: <b>default-scheduler</b>
disable_mount_path_v1	No	Bool	Default value: <b>false</b> <b>true</b> : <b>/opt/cloud/cce/nvidia</b> is not mounted to the <b>/usr/lib/nvidia</b> directory of a GPU container.
disable_nvidia_gsp	No	Bool	Default value: <b>true</b> <b>true</b> : The GPU GSP firmware is disabled.
driver_mount_paths	No	String	Driver file directory that needs to be automatically mounted to a GPU container Default value: <b>"bin,lib64"</b>
enable_fault_isolation	No	Bool	Default value: <b>true</b> <b>true</b> : The add-on detects hardware faults or driver issues of a GPU and then sets the GPU to be unavailable.
enable_health_monitoring	No	Bool	Default value: <b>true</b> <b>true</b> : The add-on detects hardware faults or driver issues of a GPU.
enable_metrics_monitoring	No	Bool	Default value: <b>true</b> <b>true</b> : The add-on collects GPU metrics and reports these metrics to Prometheus.
enable_simple_lib64_mount	No	Bool	Default value: <b>true</b> <b>true</b> : Only the <b>libxxx.so.x</b> file is mounted to a container.

Parameter	Mandatory	Type	Description
enable_xgpu	No	Bool	Default value: <b>false</b> Whether to enable xGPU virtualization.
gpu_driver_config	No	Map	Configurations of the GPU driver for a single node pool Default value: <b>{}</b>
health_check_xids_v2	No	String	GPU error range for the add-on health checks Default value: <b>"74,79"</b>
inject_ld_library_path	No	String	Value of the <b>LD_LIBRARY_PATH</b> environment variable automatically injected by the add-on to a GPU container Default value: <b>""</b>
lib64_container_paths	No	String	Mount path of NVIDIA lib64 in a GPU container Default value: <b>"/usr/lib64,/usr/lib/x86_64-linux-gnu"</b>
metrics_delete_interval	No	int	Timeout threshold for deleting a metric when the metric cannot be obtained. The unit is millisecond. Default value: <b>30000</b>
metrics_monitor_interval	No	int	Interval for obtaining metrics, in milliseconds. Default value: <b>15000</b>
nvidia_driver_download_url	Yes	String	Path for downloading the NVIDIA driver Default value: <b>""</b>

### Example Request

```
{
  "kind": "Addon",
  "apiVersion": "v3",
  "metadata": {
    "name": "gpu-beta",
  },
  "spec": {
    "clusterID": "80c9e306-***-***-***-0255ac100043",
    "version": "2.0.69",
    "addonTemplateName": "gpu-beta",
    "values": {
      "basic": {
        "cluster_version": "v1.27",
        "device_version": "2.0.69",

```

```

"driver_version": "2.0.69",
"obs_url": "****",
"region": "****",
"swr_addr": "****",
"swr_user": "****"
},
"custom": {
"compatible_with_legacy_api": true,
"component_schedulername": "kube-scheduler",
"disable_mount_path_v1": false,
"disable_nvidia_gsp": true,
"driver_mount_paths": "bin,lib64",
"enable_fault_isolation": true,
"enable_health_monitoring": true,
"enable_metrics_monitoring": true,
"enable_simple_lib64_mount": true,
"enable_xgpu": true,
"gpu_driver_config": {},
"health_check_xids_v2": "74,79",
"inject_ld_library_path": "",
"lib64_container_paths": "/usr/lib64,/usr/lib/x86_64-linux-gnu",
"metrics_delete_interval": 30000,
"metrics_monitor_interval": 15000,
"nvidia_driver_download_url": ""
},
}
}
}

```

## 4.10.10 Volcano Scheduler

### Add-on Overview

Volcano is a batch scheduling platform based on Kubernetes. It provides a series of features required by machine learning, deep learning, bioinformatics, genomics, and other big data applications, as a powerful supplement to Kubernetes capabilities.

### Add-on Parameters

**Table 4-948** Parameters

Parameter	Mandatory	Type	Description
basic	No	<a href="#">Table 4-949</a> object	Basic configuration parameters, which do not need to be specified
flavor	Yes	<a href="#">Table 4-950</a> object	Flavor parameters
custom	Yes	<a href="#">Table 4-951</a> object	Custom parameters

**Table 4-949** Configuration of basic

Parameter	Mandatory	Type	Description
swr_addr	Yes	String	Add-on download address, which does not need to be specified
swr_user	Yes	String	User who can download the add-on. This parameter does not need to be specified.
platform	Yes	String	Add-on platform, which does not need to be specified
escEndpoint	Yes	String	ECS address, which does not need to be specified

**Table 4-950** Configuration of flavor

Parameter	Mandatory	Type	Description
description	No	String	Add-on description
name	Yes	String	Add-on specification name. The value is fixed at <b>Single-instance</b> .
replicas	Yes	String	Number of pods. The default value is <b>1</b> .
resources	Yes	<b>resources</b> object	Container resource (CPU and memory) quotas

**Table 4-951** Configuration of custom

Parameter	Mandatory	Type	Description
multiAZEnabled	No	Bool	Whether multi-AZ deployment is enabled <b>true</b> : The Volcano Scheduler pods are deployed based on anti-affinity.
node_match_expressions	No	<b>Table 4-954</b>	Expression for matching the Volcano Scheduler pods to nodes
tolerations	No	<b>Table 4-953</b>	The format is the same as that of Kubernetes tolerations. It is used to add taints to Volcano Scheduler pods.
oversubscription_ratio	No	int	Node resource overcommitment ratio in the Volcano scheduling environment

Parameter	Mandatory	Type	Description
deschedule r_enable	No	Bool	Whether to support rescheduling
enable_wor kload_bala ncher	No	Bool	Whether load balancers are supported
default_sch eduler_con f	Yes	yaml	The format is the same as that of the YAML for Volcano.
deschedule rPolicy	No	yaml	The format is the same as that of the YAML for Volcano descheduling configuration.

**Table 4-952** Data structure of the resources field

Parameter	Mandatory	Type	Description
limitsCpu	Yes	String	CPU size limit (unit: m) Default value: <b>250m</b>
limitsMem	Yes	String	Memory size limit (unit: Mi) Default value: <b>512Mi</b>
name	Yes	String	Add-on name. The value is fixed at <b>virtual-kubelet</b> .
requestsCp u	Yes	String	Requested CPU size (unit: m) Default value: <b>250m</b>
requestsMe m	Yes	String	Requested memory size (unit: Mi) Default value: <b>512Mi</b>

**Table 4-953** Taints and tolerations

Parameter	Mandatory	Type	Description
key	No	String	Taint key
effect	No	String	Taint policy
operator	No	String	Operator
tolerationSeconds	No	Int	Toleration time window

**Table 4-954** nodeMatchExpression node affinity

Parameter	Mandatory	Type	Description
key	No	String	Taint key
values	No	List<String >	Node affinity name
operator	No	String	Operator

### Example Request

```
{
  "kind": "Addon",
  "apiVersion": "v3",
  "metadata": {
    "annotations": {
      "addon.install/type": "install"
    }
  },
  "spec": {
    "clusterID": "2292498e-d169-*****-0255ac1001ba",
    "version": "1.12.14",
    "addonTemplateName": "volcano",
    "values": {
      "basic": {
        "swr_addr": "x.x.x.x:20202",
        "swr_user": "hwofficial",
        "platform": "linux-amd64",
        "ecsEndpoint": "x.x.x.x"
      },
      "flavor": {
        "resources": [{
          "limitsCpu": "500m",
          "limitsMem": "500Mi",
          "name": "volcano-scheduler",
          "requestsCpu": "2000m",
          "requestsMem": "2000Mi"
        },
        {
          "limitsCpu": "500m",
          "limitsMem": "500Mi",
          "name": "volcano-controller",
          "requestsCpu": "2000m",
          "requestsMem": "2000Mi"
        },
        {
          "limitsCpu": "200m",
          "limitsMem": "500Mi",
          "name": "volcano-admission",
          "requestsCpu": "500m",
          "requestsMem": "500Mi"
        },
        {
          "limitsCpu": "100m",
          "limitsMem": "150Mi",
          "name": "volcano-agent",
          "requestsCpu": "200m",
          "requestsMem": "200Mi"
        },
        {
          "limitsCpu": "50m",
          "limitsMem": "50Mi",
          "name": "volcano-exporter",
          "requestsCpu": "100m",

```



```

    "requestsMem": "100Mi"
  },
  {
    "limitsCpu": "500m",
    "limitsMem": "256Mi",
    "name": "volcano-descheduler",
    "requestsCpu": "1000m",
    "requestsMem": "512Mi"
  },
  {
    "limitsCpu": "300m",
    "limitsMem": "300Mi",
    "name": "volcano-recommender",
    "requestsCpu": "500m",
    "requestsMem": "500Mi"
  },
  {
    "limitsCpu": "200m",
    "limitsMem": "200Mi",
    "name": "volcano-recommender-prometheus-adapter",
    "requestsCpu": "300m",
    "requestsMem": "300Mi"
  }
}
"custom": {
  "default_scheduler_conf": {
    "actions": "allocate,backfill,preempt",
    "tiers": [
      {
        "plugins": [
          {
            "name": "priority"
          },
          {
            "name": "conformance"
          },
          {
            "name": "gang",
            "enablePreemptable": "false",
            "enableJobStarving": "false",
          }
        ]
      },
      {
        "plugins": [
          {
            "name": "predicates"
          },
          {
            "name": "nodeorder"
          },
          {
            "name": "drf",
            "enablePreemptable": "false",
          }
        ]
      },
      {
        "plugins": [
          {
            "name": "cce-gpu-topology-predicate"
          },
          {
            "name": "cce-gpu-topology-priority"
          },
          {
            "name": "xgpu"
          }
        ]
      },
      {
        "plugins": [
          {
            "name": "nodelocalvolume"
          },
          {
            "name": "nodeemptydirvolume"
          }
        ]
      }
    ]
  }
}

```



## Example Request

```
{
  "kind": "Addon",
  "apiVersion": "v3",
  "metadata": {
    "annotations": {
      "addon.install/type": "install"
    }
  },
  "spec": {
    "clusterID": "2292498e-d169-*****-0255ac1001ba",
    "version": "1.1.1",
    "addonTemplateName": "dew-provider",
    "values": {
      "basic": {
        "basickey": "val"
      },
      "flavor": {
        "resources": [{
          "limitsCpu": "100m",
          "limitsMem": "100Mi",
          "name": "dew-provider",
          "requestsCpu": "100m",
          "requestsMem": "100Mi"
        }]
      },
      "custom": {
        "rotation_poll_interval": "2m"
      }
    }
  }
}
```

## 4.10.12 CCE Network Metrics Exporter

### Add-on Overview

CCE Network Metrics Exporter (dolphin) is an add-on for monitoring and managing container network traffic. It collects traffic statistics of containers that do not use the host network in CCE Turbo clusters and performs node-wide container connectivity checks.

### Add-on Parameters

**Table 4-957** Parameters

Parameter	Mandatory	Type	Description
basic	No	object	Basic configuration parameters, which do not need to be specified
flavor	Yes	<a href="#">Table 4-958</a> object	Flavor parameters
custom	Yes	<a href="#">Table 4-959</a> object	Custom parameters

**Table 4-958** Configuration of flavor

Parameter	Mandatory	Type	Description
name	Yes	String	Add-on specification name. The value is fixed at <b>default</b> .
resources	Yes	<b>resources</b> object	Container resource (CPU and memory) quotas

**Table 4-959** Configuration of custom

Parameter	Mandatory	Type	Description
annotations	No	Map<String>String	Custom annotation

**Table 4-960** Data structure of the resources field

Parameter	Mandatory	Type	Description
limitsCpu	Yes	String	CPU size limit (unit: m) Default value: <b>500m</b>
limitsMem	Yes	String	Memory size limit (unit: Mi) Default value: <b>512Mi</b>
name	Yes	String	Add-on name. The value is fixed at <b>dolphin</b> .
requestsCpu	Yes	String	Requested CPU size (unit: m) Default value: <b>500m</b>
requestsMem	Yes	String	Requested memory size (unit: Mi) Default value: <b>512Mi</b>

## Example Request

```
{
  "kind": "Addon",
  "apiVersion": "v3",
  "metadata": {
    "name": "dolphin",
    "alias": "CCE Network Metrics Exporter",
    "addon.install/type": "install"
  },
  "spec": {
    "clusterID": "****",
    "version": "1.4.5",
    "addonTemplateName": "dolphin",
    "values": {
      "basic": {
        "cluster_version": "v1.28",
```

```

    "image_version": "1.4.5",
    "platform": "linux-amd64",
    "rbac_enabled": true,
    "swr_addr": "****",
    "swr_user": "****"
  },
  "custom": {
    "annotations": {}
  },
  "flavor": {
    "name": "default",
    "resources": [
      {
        "limitsCpu": "500m",
        "limitsMem": "512Mi",
        "name": "dolphin",
        "requestsCpu": "500m",
        "requestsMem": "512Mi"
      }
    ]
  }
}

```

## 4.10.13 NodeLocal DNSCache

### Add-on Overview

NodeLocal DNSCache is an add-on developed based on the community NodeLocal DNSCache. This add-on functions as a DaemonSet to run the DNS cache proxy on cluster nodes to improve cluster DNS performance.

### Add-on Parameters

**Table 4-961** Parameters

Parameter	Mandatory	Type	Description
basic	No	object	Basic configuration parameters, which do not need to be specified
flavor	Yes	<a href="#">Table 4-962</a> object	Flavor parameters
custom	Yes	<a href="#">Table 4-963</a> object	Custom parameters

**Table 4-962** Configuration of flavor

Parameter	Mandatory	Type	Description
replicas	No	String	Number of admission-controller component instances in the add-on. The default value is <b>2</b> .

Parameter	Mandatory	Type	Description
resources	No	Array <b>resources</b> object	Container resource (CPU and memory) quotas

**Table 4-963** Configuration of custom

Parameter	Mandatory	Type	Description
enable_dnsconfig_admission	No	bool	Enables DNSConfig automatic injection. The default value is <b>true</b> . After this function is enabled, a DNSConfig admission controller will be created. The controller intercepts pod creation requests in the namespace labeled with <b>node-localdns-injection=enabled</b> based on Admission Webhook, and automatically configures <b>Pod dnsConfig</b> that uses the DNS cache. If this function is disabled or the pod belongs to a non-target namespace, you must manually configure DNSConfig for the pod.
enable_namespace_admission	No	bool	Adds <b>node-local-dns-injection=enabled</b> to a created namespace. The default value is <b>true</b> . After this tag is added to a namespace, the system identifies the namespace creation request and automatically adds the tag. The target of these operations does not include the built-in namespaces (such as <b>kube-system</b> ).
multiAZEnable	No	bool	Whether to enable the forcible mode of multi-AZ deployment for the deployment component. The default value is <b>false</b> . Deployment pods of the add-on will be forcibly scheduled to nodes in different AZs. If there are fewer AZs than pods, the extra pods will fail to run. If both <b>multiAZEnable</b> and <b>multiAZBalance</b> are set to true, the settings of <b>multiAZBalance</b> take effect, which means, the equivalent mode of multi-AZ deployment is used.

Parameter	Mandatory	Type	Description
multiAZBalance	No	bool	Whether to enable the equivalent mode of multi-AZ deployment for the deployment component. The default value is <b>false</b> . Deployment pods of the add-on are evenly scheduled to the nodes in the cluster in each AZ. If a new AZ is added, it is recommended that you increase add-on pods for cross-AZ HA deployment. With the equivalent multi-AZ deployment, the difference between the number of add-on pods in different AZs will be less than or equal to 1. If resources in one of the AZs are insufficient, pods cannot be scheduled to that AZ.
tolerations	No	Array of <a href="#">Table 4-965</a>	Tolerations of the admission-controller component

**Table 4-964** Data structure of the resources field

Parameter	Mandatory	Type	Description
limitsCpu	Yes	String	CPU size limit (unit: m)
limitsMem	Yes	String	Memory size limit (unit: Mi)
name	Yes	String	Add-on name. The value is fixed at <b>node-local-dns-admission-controller</b> or <b>node-local-dns-cache</b> .
requestsCpu	Yes	String	Requested CPU size (unit: m)
requestsMem	Yes	String	Requested memory size (unit: Mi)

**Table 4-965** Taints and tolerations

Parameter	Mandatory	Type	Description
key	No	String	Taint key
effect	No	String	Taint policy
operator	No	String	Operator
tolerationSeconds	No	Int	Toleration time window

## Example Request

```
{
  "kind": "Addon",
  "apiVersion": "v3",
  "metadata": {
    "annotations": {
      "addon.install/type": "install"
    }
  },
  "spec": {
    "clusterID": "2292498e-d169-*****-0255ac1001ba",
    "version": "1.6.2",
    "addonTemplateName": "node-local-dns",
    "values": {
      "basic": {
        "basickey": "val"
      },
      "flavor": {
        "replicas": 2,
        "resources": [{
          "limitsCpu": "250m",
          "limitsMem": "512Mi",
          "name": "node-local-dns-admission-controller",
          "requestsCpu": "250m",
          "requestsMem": "512Mi"
        },
        {
          "limitsCpu": "500m",
          "limitsMem": "512Mi",
          "name": "node-local-dns-cache",
          "requestsCpu": "25m",
          "requestsMem": "5Mi"
        }
      ]
    },
    "custom": {
      "enable_dnsconfig_admission": true,
      "enable_namespace_admission": true,
      "multiAZBalance": false,
      "multiAZEnabled": false,
      "node_match_expressions": [],
      "tolerations": [{
        "key": "node.kubernetes.io/not-ready",
        "operator": "Exists",
        "effect": "NoExecute",
        "tolerationSeconds": 60
      },
      {
        "key": "node.kubernetes.io/unreachable",
        "operator": "Exists",
        "effect": "NoExecute",
        "tolerationSeconds": 60
      }
    ]
  }
}
```



## 4.10.14 Cloud Native Cluster Monitoring

### Add-on Overview

The Cluster Native Cluster Monitoring add-on (kube-prometheus-stack) uses Prometheus-operator and Prometheus to provide easy-to-use, end-to-end Kubernetes cluster monitoring.

This add-on enables monitoring data to connect to the monitoring center, so that you can view monitoring data and configure alarms on the console.

### Add-on Parameters

**Table 4-966** Parameters

Parameter	Mandatory	Type	Description
basic	No	object	Basic configuration parameters, which do not need to be specified
flavor	Yes	<a href="#">Table 4-967</a> object	Flavor parameters
custom	Yes	<a href="#">Table 4-968</a> object	Custom parameters

**Table 4-967** Configuration of flavor

Parameter	Mandatory	Type	Description
description	No	String	Add-on description
name	Yes	String	Add-on specification name
resources	Yes	<a href="#">resources</a> object	Container resource (CPU and memory) quotas

**Table 4-968** Configuration of custom

Parameter	Mandatory	Type	Description
aom_enable	No	Boolean	Whether to interconnect with AOM
aom_instance_id	No	String	AOM instance ID. This parameter is mandatory when AOM is interconnected with the cluster.

Parameter	Mandatory	Type	Description
aom_project_id	No	String	ID of the project to which the AOM instance belongs. This parameter is mandatory when AOM is interconnected with the cluster.
aom_auth_type	No	String	Authentication type for interconnecting with AOM. This parameter is mandatory when AOM is interconnected with the cluster. The value is fixed at <b>Bearer</b> .
aom_app_key	No	String	app_key for interconnecting with AOM. This parameter is mandatory when AOM is interconnected with the cluster.
aom_app_secret	No	String	app_secret for interconnecting with AOM. This parameter is mandatory when AOM is interconnected with the cluster.
deploy_mode	Yes	String	Prometheus add-on model. The value can be <b>agent</b> or <b>server</b> . <b>agent</b> is recommended.
enablethird	No	Boolean	Whether to report metrics to the third-party monitoring system.
url_third	No	String	URL for the third-party monitoring system to report metrics.
basic_auth_username_third	No	String	Account for interconnecting with a third-party monitoring system (in the BasicAuth authentication mode)
basic_auth_password_third	No	String	Password for interconnecting with a third-party monitoring system (in the BasicAuth authentication mode)
bearer_token	No	String	Token for interconnecting with a third-party monitoring system (in the BearerToken authentication mode)
cluster	Yes	String	Cluster name
clusterId	Yes	String	Cluster ID
enable_custom_metrics	No	Boolean	Whether to enable custom metric collection. The default value is <b>false</b> .
highAvailability	No	Boolean	Whether HA is enabled. The default value is <b>false</b> .

Parameter	Mandatory	Type	Description
scrapeInterval	Yes	String	Default metric collection period. The default value is <b>15s</b> .
shards	No	Integer	Number of collected fragments. This parameter is valid only in the agent mode. The default value is <b>1</b> .

**Table 4-969** Data structure of the resources field

Parameter	Mandatory	Type	Description
limitsCpu	Yes	String	CPU size limit (unit: m)
limitsMem	Yes	String	Memory size limit (unit: Mi)
name	Yes	String	Workload name
requestsCpu	Yes	String	Requested CPU size (unit: m)
requestsMem	Yes	String	Requested memory size (unit: Mi)

## Example Request

```
{
  "kind": "Addon",
  "apiVersion": "v3",
  "metadata": {
    "annotations": {
      "addon.install/type": "install"
    }
  },
  "spec": {
    "clusterID": "b889764e-c100-11ee-80cf-0255ac100b0f",
    "version": "3.9.5",
    "addonTemplateName": "cie-collector",
    "values": {
      "basic": {
        "aom_endpoint": "https://**",
        "aom_url": "https://**",
        "region_id": "**",
        "swr_addr": "**",
        "swr_user": "**",
        "rbac_enabled": true,
        "cluster_version": "v1.27"
      },
      "flavor": {
        "deploy_mode": "agent",
        "description": "Default flavor for agent mode.",
        "name": "Agent-Default",
        "resources": [
          {
            "limitsCpu": "500m",
            "limitsMem": "500Mi",
            "name": "prometheusOperator",
            "requestsCpu": "100m",

```

```
    "requestsMem": "100Mi"
  },
  {
    "limitsCpu": "4",
    "limitsMem": "8Gi",
    "name": "prometheus",
    "requestsCpu": "500m",
    "requestsMem": "500Mi"
  },
  {
    "limitsCpu": "500m",
    "limitsMem": "500Mi",
    "name": "kubeStateMetrics",
    "requestsCpu": "200m",
    "requestsMem": "200Mi"
  },
  {
    "limitsCpu": "500m",
    "limitsMem": "1Gi",
    "name": "nodeExporter",
    "requestsCpu": "200m",
    "requestsMem": "100Mi"
  }
],
"category": [
  "CCE",
  "Turbo"
],
"custom": {
  "aom_app_key": "*****",
  "aom_app_secret": "*****",
  "aom_auth_type": "Bearer",
  "aom_enable": true,
  "aom_instance_id": "a2bf5a6f-2c64-4a7f-a369-78dccd9cb6fd",
  "aom_project_id": "b6315dd3d0ff4be5b31a963256794989",
  "basic_auth_password_third": "",
  "basic_auth_username_third": "",
  "bearer_token": "",
  "cluster": "test-cluster",
  "clusterId": "b889764e-c100-11ee-80cf-0255ac100b0f",
  "deploy_mode": "agent",
  "enable_custom_metrics": true,
  "enablethird": false,
  "highAvailability": false,
  "projectId": "b6315dd3d0ff4be5b31a963256794989",
  "region": "****",
  "scrapeInterval": "15s",
  "shards": 1,
  "url_third": ""
}
}
}
```

## 4.10.15 Cloud Native Logging

### Add-on Overview

The Cloud Native Logging add-on (log-agent) is developed based on Fluent Bit and OpenTelemetry for collecting logs and Kubernetes events. This add-on supports CRD-based log collection policies. It collects and forwards standard output logs, container file logs, node logs, and Kubernetes event logs in a cluster based on configured policies. It also reports Kubernetes events to AOM for configuring event alarms. By default, all abnormal events and some normal events are reported.

## Add-on Parameters

**Table 4-970** Parameters

Parameter	Mandatory	Type	Description
basic	No	object	Basic configuration parameters, which do not need to be specified
flavor	Yes	<a href="#">Table 4-971</a> object	Flavor parameters
custom	Yes	<a href="#">Table 4-972</a> object	Custom parameters

**Table 4-971** Configuration of flavor

Parameter	Mandatory	Type	Description
description	No	String	Add-on description
name	Yes	String	Add-on specification name. The value is fixed at <b>Single-instance</b> .
replicas	Yes	String	Number of otel-collector component instances. The default value is <b>1</b> .
resources	Yes	<b>resources</b> object	Container resource (CPU and memory) quotas

**Table 4-972** Configuration of custom

Parameter	Mandatory	Type	Description
caCert	Yes	String	Specifies a client certificate, which is encrypted using Base64. This parameter is optional during installation. It is generated by the CCE add-on center and used for two-way authentication between components and webhook requests.  Domain name to be issued: <b>*.monitoring.svc</b> .

Parameter	Mandatory	Type	Description
serverCert	Yes	String	Specifies a server certificate, which is encrypted using Base64. This parameter is optional during installation. It is generated by CCE add-ons and used for two-way authentication between components and webhook requests. Domain name to be issued: <b>*.monitoring.svc</b> .
serverKey	Yes	String	Specifies the private key of a server, which is encrypted using Base64. This parameter is optional during installation. It is generated by CCE add-ons and used for two-way authentication between components and webhook requests. Domain name to be issued: <b>*.monitoring.svc</b> .
accessKey	No	String	Specifies the access key ID of a user, which is used to request AOM and LTS APIs. If this parameter is left blank, a temporary AK/SK is used.
secretKey	No	String	Specifies the access key of a user, which is used to request AOM and LTS APIs. If this parameter is left blank, a temporary AK/SK is used.
createDefaultStdout	No	Bool	Specifies whether to create the default policy for collecting container standard output logs and reporting them to LTS. This parameter is valid only after it is specified during the add-on installation. The default value is <b>false</b> .
createDefaultEvent	No	Bool	Specifies whether to create the default policy for collecting Kubernetes events and reporting them to LTS. This parameter is valid only after it is specified during the add-on installation. The default value is <b>false</b> .
multiAZEnabled	No	Bool	Specifies whether to deploy the add-on pods in multiple AZs. The default value is <b>false</b> . If this parameter is set to <b>true</b> , cross-AZ deployment is forcibly performed. If this parameter is set to <b>false</b> , cross-AZ deployment is preferred.

Parameter	Mandatory	Type	Description
cluster_category	No	String	The value is fixed at <b>CCE</b> .
ltsAccessEndpoint	No	String	Specifies the address for reporting logs to LTS. If this parameter is left blank, the address specified in <b>basic</b> is used.
ltsEndpoint	No	String	Specifies the address of the LTS APIs. If this parameter is left blank, the address specified in <b>basic</b> is used.
aomEndpoint	No	String	Specifies the address for reporting events to AOM. If this parameter is left blank, the address specified in <b>basic</b> is used.
projectID	Yes	String	Specifies the ID of the project to which the current CCE cluster belongs. This parameter is optional during installation.
clusterID	Yes	String	Specifies the ID of the current CCE cluster. This parameter is optional during installation.
clusterName	Yes	String	Specifies the name of the current CCE cluster.

**Table 4-973** Data structure of the resources field

Parameter	Mandatory	Type	Description
limitsCpu	Yes	String	CPU size limit (unit: m)
limitsMem	Yes	String	Memory size limit (unit: Mi)
name	Yes	String	Component name. Options: <b>fluent-bit</b> : component for collecting node logs <b>cop-logs</b> : generates soft links for collected files on each node. <b>log-operator</b> : generates fluent-bit and otel-collector configurations. <b>otel-collector</b> : reports collected logs and events to AOM and LTS.
requestsCpu	Yes	String	Requested CPU size (unit: m)

Parameter	Mandatory	Type	Description
requestsMem	Yes	String	Requested memory size (unit: Mi)

## Example Request

```
{
  "kind": "Addon",
  "apiVersion": "v3",
  "metadata": {
    "annotations": {
      "addon.install/type": "install"
    }
  },
  "spec": {
    "clusterID": "ccbe7bdf-4**9-3**b-b**4-0*****78",
    "version": "1.4.4",
    "addonTemplateName": "log-agent",
    "values": {
      "basic": {
        "aomEndpoint": "https://**",
        "iam_url": "**",
        "ltsAccessEndpoint": "https://**:8102",
        "ltsEndpoint": "https://**",
        "region": "**",
        "swr_addr": "**",
        "swr_user": "**",
        "rbac_enabled": true,
        "cluster_version": "v1.25"
      }
    },
    "flavor": {
      "description": "Recommended when the number of logs per second does not exceed 5000.",
      "name": "Low",
      "replicas": 1,
      "resources": [
        {
          "limitsCpu": "500m",
          "limitsMem": "500Mi",
          "name": "fluent-bit",
          "requestsCpu": "100m",
          "requestsMem": "200Mi"
        },
        {
          "limitsCpu": 1,
          "limitsMem": "500Mi",
          "name": "cop-logs",
          "requestsCpu": "100m",
          "requestsMem": "100Mi"
        },
        {
          "limitsCpu": "500m",
          "limitsMem": "500Mi",
          "name": "log-operator",
          "requestsCpu": "100m",
          "requestsMem": "100Mi"
        },
        {
          "limitsCpu": 1,
          "limitsMem": "2Gi",
          "name": "otel-collector",
          "requestsCpu": "200m",
          "requestsMem": "1Gi"
        }
      ]
    },
    "category": [
```



```
    "CCE",
    "Turbo"
  ]
},
"custom": {
  "accessKey": "",
  "aomEndpoint": "https://****",
  "aomPrivateEndpointIP": "",
  "caCert": "",
  "clusterID": "",
  "clusterName": "clusterName",
  "cluster_category": "CCE",
  "createAudit": false,
  "createDefaultEvent": false,
  "createDefaultEventToAOM": true,
  "createDefaultStdout": false,
  "createKubeApiserver": false,
  "createKubeControllerManager": false,
  "createKubeScheduler": false,
  "ltsAccessEndpoint": "https://***:8102",
  "ltsAuditStreamID": "",
  "ltsEndpoint": "https://****",
  "ltsEventStreamID": "",
  "ltsGroupID": "",
  "ltsKubeApiserverStreamID": "",
  "ltsKubeControllerManagerStreamID": "",
  "ltsKubeSchedulerStreamID": "",
  "ltsLogReportDomain": "",
  "ltsPrivateEndpointIP": "",
  "ltsStdoutStreamID": "",
  "multiAZEnabled": false,
  "paasakskEnable": true,
  "projectID": "",
  "secretKey": "",
  "securityToken": "",
  "serverCert": "",
  "serverKey": ""
}
}
}
```

# 5 Kubernetes APIs

## Description

Kubernetes APIs are resource-based (RESTful) programming interfaces provided through HTTP. It supports query, creation, update, and deletion of various cluster resources using standard HTTP request methods (POST, PUT, PATCH, DELETE, and GET).

CCE allows you to use native [Kubernetes APIs](#) in the following ways:

- [Calling Kubernetes APIs Through the Cluster API Server](#). It is suitable for API calls on scale thanks to its direct connection to the API Server. This is a recommended option.
- [Calling Kubernetes APIs Through API Gateway](#). It applies to small-scale API calls. API gateway flow control may be triggered when APIs are called on scale.

## Calling Kubernetes APIs Through the Cluster API Server

You can use the API server of a Kubernetes cluster to call Kubernetes-native APIs.

### Step 1 Obtain the cluster certificate and API server.

- Method 1: Obtain the certificate by calling the API for [obtaining the cluster certificate](#), save the returned information to the `kubeconfig.json` file, and extract the certificate, private key, and API server information. The commands are as follows:

```
# Obtain the CA certificate of the cluster and save it as ca.crt.  
cat ./kubeconfig.json |grep certificate-authority-data | awk -F "" '{print $4}' | base64 -d > ./ca.crt  
# Obtain the client certificate and save it as client.crt.  
cat ./kubeconfig.json |grep client-certificate-data | awk -F "" '{print $4}' | base64 -d > ./client.crt  
# Obtain the client private key and save it as client.key.  
cat ./kubeconfig.json |grep client-key-data | awk -F "" '{print $4}' | base64 -d > ./client.key  
# Obtain the API server.  
cat ./kubeconfig.json |grep server | awk -F "" '{print $4}'
```

- Method 2: Obtain the API server IP address (private or public network address) on the [Overview](#) page of the CCE console and download the certificate (`ca.crt`, `client.crt`, and `client.key` files).

### Step 2 Call Kubernetes-native APIs using the cluster certificate.

For example, run the `curl` command to call an API to view the pod information. The following is an example:

```
curl --cacert ./ca.crt --cert ./client.crt --key ./client.key https://192.168.0.198:5443/api/v1/namespaces/default/pods/
```

Specifically:

- `./ca.crt`, `./client.crt`, and `./client.key` specify that the certificate files in the current path are used. Replace them with the actual file paths.
- **192.168.0.198:5443** is the IP address of the cluster API server.
- `/api/v1/namespaces/default/pods/` specifies the URI of the cluster API for viewing pod information in the **default** namespace. For more cluster APIs, see [Kubernetes API](#).

----End

## Calling Kubernetes APIs Through API Gateway

You can call Kubernetes-native APIs through API Gateway using the URL in the format of **https://{clusterid}.Endpoint/uri**. In the URL, `{clusterid}` indicates the cluster ID, and `uri` indicates the resource path, that is, the path for API access.

**Table 5-1** URL parameters

Parameter	Description
{clusterid}	Cluster ID. After a cluster is created, call the <a href="#">API for obtaining a cluster in a specified project</a> to obtain the cluster ID.
Endpoint	Entry (URL) for a web service, which can be obtained from <a href="#">Endpoints</a> .
uri	Access path of an API for performing an operation. Obtain the value from the URI of the API. For details, see <a href="#">Kubernetes API</a> .

**Step 1** Obtain the token of the region where the cluster is located. For details about how to obtain the token, see [Obtaining a Token](#).

**Step 2** Obtain the cluster ID using either of the following methods:

- Method 1: Obtain the cluster UID by calling the API for obtaining cluster information.
- Method 2: Obtain the cluster ID on the **Overview** page of the CCE console.

**Step 3** Determine the requested URL based on the URL format **https://{clusterid}.Endpoint/uri**.

- **{clusterid}**: Obtain the value by using [Step 2](#).
- **Endpoint**: Obtain the endpoint from [Regions and Endpoints](#).  
For example, the endpoint of CCE in the **my-kualalumpur-1** region is **cce.my-kualalumpur-1.myhuaweicloud.com**.
- **uri**: Set this parameter based on the API to be called. For example, if you want to create a Deployment, the request method is POST and the API URI is `/apis/apps/v1/namespaces/{namespace}/deployments`, where

*{namespace}* indicates the cluster namespace name. In this example, the value is **default**.

For more APIs, see [Kubernetes APIs](#).

Combine the preceding parameters following the URL format **https://{clusterid}.Endpoint/uri**.

The following is an example of the URL for calling the API to view information about all pods:

```
https://07da5****.cce.my-kualalumpur-1.myhuaweicloud.com/apis/apps/v1/namespaces/default/deployments
```

- Step 4** Use the request method specified by the API and set the request header parameters. If parameters in the body need to be added, add the structure corresponding to the API by referring to [Kubernetes APIs](#).

Example curl command to call the API for creating a Deployment using POST and adding the corresponding body:

In this example, the **nginx.json** file is used to create a Deployment named **nginx**. The Deployment uses the **nginx:latest** image and contains two pods. Each pod occupies 100m CPUs and 200 MiB of memory.

```
curl --location --request POST 'https://07da5****.cce.my-kualalumpur-1.myhuaweicloud.com/apis/apps/v1/namespaces/default/deployments' \
--header 'Content-Type: application/json' \
--header 'X-Auth-Token: MIIWVW****' \
--data @nginx.json
```

Header parameters contained in the request are as follows:

**Table 5-2** Request header parameters

Parameter	Mandatory	Type	Description
Content-Type	Yes	String	Message body type (format), for example, application/json.
X-Auth-Token	Yes	String	Requests for calling an API can be authenticated using either a token or AK/SK. If token-based authentication is used, this parameter is mandatory and must be set to a user token. For details, see <a href="#">Obtaining a Token</a> .

**nginx.json** is located in the current directory and contains the following content:

```
{
  "apiVersion": "apps/v1",
  "kind": "Deployment",
  "metadata": {
    "name": "nginx"
  },
  "spec": {
    "replicas": 2,
    "selector": {
```

```

    "matchLabels": {
      "app": "nginx"
    }
  },
  "template": {
    "metadata": {
      "labels": {
        "app": "nginx"
      }
    },
    "spec": {
      "containers": [
        {
          "image": "nginx:latest",
          "name": "container-0",
          "resources": {
            "limits": {
              "cpu": "100m",
              "memory": "200Mi"
            },
            "requests": {
              "cpu": "100m",
              "memory": "200Mi"
            }
          }
        }
      ],
      "imagePullSecrets": [
        {
          "name": "default-secret"
        }
      ]
    }
  }
}

```

----End

### Related Documents

- [Accessing a Cluster Using Kubernetes APIs](#)
- [Kubernetes official SDKs](#) (including Go, Python, and Java)

Language	Client Library	Sample Program
C	<a href="https://github.com/kubernetes-client/c">github.com/kubernetes-client/c</a>	<a href="#">Browse</a>
dotnet	<a href="https://github.com/kubernetes-client/csharp">github.com/kubernetes-client/csharp</a>	<a href="#">Browse</a>
Go	<a href="https://github.com/kubernetes/client-go">github.com/kubernetes/client-go/</a>	<a href="#">Browse</a>
Haskell	<a href="https://github.com/kubernetes-client/haskell">github.com/kubernetes-client/haskell</a>	<a href="#">Browse</a>
Java	<a href="https://github.com/kubernetes-client/java">github.com/kubernetes-client/java</a>	<a href="#">Browse</a>
JavaScript	<a href="https://github.com/kubernetes-client/javascript">github.com/kubernetes-client/javascript</a>	<a href="#">Browse</a>
Perl	<a href="https://github.com/kubernetes-client/perl/">github.com/kubernetes-client/perl/</a>	<a href="#">Browse</a>

Language	Client Library	Sample Program
Python	<a href="https://github.com/kubernetes-client/python/">github.com/kubernetes-client/python/</a>	<a href="#">Browse</a>
Ruby	<a href="https://github.com/kubernetes-client/ruby/">github.com/kubernetes-client/ruby/</a>	<a href="#">Browse</a>

# 6 Permissions and Supported Actions

---

You can use Identity and Access Management (IAM) for fine-grained permissions management of your CCE clusters. If your account does not need individual IAM users, you can skip this section.

New IAM users do not have any permissions assigned by default. You need to first add them to one or more groups and attach policies or roles to these groups. Users inherit permissions from the groups to which they are added and can perform specific operations on cloud services based on the assigned permissions.

You can grant users permissions by using roles and policies. Roles are provided by IAM to define service-based permissions that match users' job responsibilities. Policies define API-based permissions for operations on specific resources under certain conditions, allowing for more fine-grained, secure access control of cloud resources.

## NOTE

If you want to allow or deny the access to an API, fine-grained authorization is a good choice.

An account has all of the permissions required to call all APIs, but IAM users must have the required permissions specifically assigned. The required permissions are determined by the actions supported by the API. Only users with the permissions allowing for those actions can call the API successfully. For example, if an IAM user wants to query ECSs using an API, the user must have been granted permissions that allow the **ecs:servers:list** action.

## Supported Actions

CCE provides system-defined policies that can be directly used in IAM. You can also create custom policies to supplement system-defined policies for more refined access control. Operations supported by policies are specific to APIs. The following are common concepts related to policies:

- Permissions: statements in a policy that allow or deny certain operations.
- APIs: REST APIs that can be called by a user who has been granted specific permissions.
- Actions: specific operations that are allowed or denied in a custom policy.

- **Dependencies:** actions which a specific action depends on. When allowing an action for a user, you also need to allow any existing action dependencies for that user.
- **IAM projects/Enterprise projects:** the authorization scope of a custom policy. A custom policy can be applied to IAM projects or enterprise projects or both. Policies that contain actions for IAM and enterprise projects can be assigned to user groups and be applied in IAM and Enterprise Management. Policies that contain actions only for IAM projects can be assigned to user groups and be applied in IAM. They cannot be applied in Enterprise Management.

 **NOTE**

The check mark (√) and cross symbol (x) respectively indicate that an action takes effect or does not take effect for the corresponding type of projects.

CCE supports the following actions in custom policies.

**Table 6-1** Cluster management actions

Permission	API	Action	IAM Project	Enterprise Project
Obtaining clusters in a project	GET /api/v3/projects/{project_id}/clusters	cce:cluster:list	√	√
Obtaining a cluster	GET /api/v3/projects/{project_id}/clusters/{cluster_id}	cce:cluster:get	√	√
Creating a cluster	POST /api/v3/projects/{project_id}/clusters	cce:cluster:create	√	√
Updating a cluster	PUT /api/v3/projects/{project_id}/clusters/{cluster_id}	cce:cluster:update	√	√
Deleting a cluster	DELETE /api/v3/projects/{project_id}/clusters/{cluster_id}	cce:cluster:delete	√	√
Upgrading a cluster	POST /api/v2/projects/:projectid/clusters/:clusterid/upgrade	cce:cluster:upgrade	√	√
Waking up a cluster	POST /api/v3/projects/{project_id}/clusters/{cluster_id}/operation/awake	cce:cluster:start	√	√
Hibernating a cluster	POST /api/v3/projects/{project_id}/clusters/{cluster_id}/operation/hibernate	cce:cluster:stop	√	√



Permission	API	Action	IAM Project	Enterprise Project
Changing the specifications of a cluster	POST /api/v2/projects/{project_id}/clusters/:clusterid/resize	cce:cluster:resize	√	√
Obtaining the certificate of a cluster	POST /api/v3/projects/{project_id}/clusters/{cluster_id}/clustercert	cce:cluster:get	√	√

**Table 6-2** Node management actions

Permission	API	Action	IAM Project	Enterprise Project
Obtaining all nodes in a cluster	GET /api/v3/projects/{project_id}/clusters/{cluster_id}/nodes	cce:node:list	√	√
Obtaining a node	GET /api/v3/projects/{project_id}/clusters/{cluster_id}/nodes/{node_id}	cce:node:get	√	√
Creating a node	POST /api/v3/projects/{project_id}/clusters/{cluster_id}/nodes	cce:node:create	√	√ <b>NOTE</b> If you use enterprise project authorization to create a node, you need to add the global permission of <b>evs:quota:get</b> .
Updating a node	PUT /api/v3/projects/{project_id}/clusters/{cluster_id}/nodes/{node_id}	cce:node:update	√	√
Deleting a node	DELETE /api/v3/projects/{project_id}/clusters/{cluster_id}/nodes/{node_id}	cce:node:delete	√	√

**Table 6-3** Job management actions

Permission	API	Action	IAM Project	Enterprise Project
Obtaining information about a job	GET /api/v3/projects/{project_id}/jobs/{job_id}	cce:job:get	√	√
Listing all jobs	GET /api/v2/projects/{project_id}/jobs	cce:job:list	√	√
Deleting one or all jobs	DELETE /api/v2/projects/{project_id}/jobs DELETE /api/v2/projects/{project_id}/jobs/{job_id}	cce:job:delete	√	√

**Table 6-4** Node pool management actions

Permission	API	Action	IAM Project	Enterprise Project
Obtaining all node pools in a cluster	GET /api/v3/projects/{project_id}/clusters/{cluster_id}/nodepools	cce:nodepool:list	√	√
Obtaining a node pool	GET /api/v3/projects/{project_id}/clusters/{cluster_id}/nodepools/{nodepool_id}	cce:nodepool:get	√	√
Creating a node pool	POST /api/v3/projects/{project_id}/clusters/{cluster_id}/nodepools	cce:nodepool:create	√	√
Updating a node pool	PUT /api/v3/projects/{project_id}/clusters/{cluster_id}/nodepools/{nodepool_id}	cce:nodepool:update	√	√
Deleting a node pool	DELETE /api/v3/projects/{project_id}/clusters/{cluster_id}/nodepools/{nodepool_id}	cce:nodepool:delete	√	√

**Table 6-5** Chart management actions

Permission	API	Action	IAM Project	Enterprise Project
Updating a chart	PUT /v2/charts/{id}	cce:chart:update	√	x
Uploading a chart	POST /v2/charts	cce:chart:upload	√	x
Listing all charts	GET /v2/charts	cce:chart:list	√	x
Obtaining information about a chart	GET /v2/charts/{id}	cce:chart:get	√	x
Deleting a chart	DELETE /v2/charts/{id}	cce:chart:delete	√	x

**Table 6-6** Release management actions

Permission	API	Action	IAM Project	Enterprise Project
Updating a release	PUT /v2/releases/{name}	cce:release:update	√	√
Listing all releases	GET /v2/releases	cce:release:list	√	√
Creating a release	POST /v2/releases	cce:release:create	√	√
Obtaining information about a release	GET /v2/releases/{name}	cce:release:get	√	√
Deleting a release	DELETE /v2/releases/{name}	cce:release:delete	√	√

**Table 6-7** Storage management actions

Permission	API	Action	IAM Project	Enterprise Project
Creating a PersistentVolumeClaim	POST /api/v1/namespaces/{namespace}/cloudpersistentvolumeclaims	cce:storage:create	√	√
Deleting a PersistentVolumeClaim	DELETE /api/v1/namespaces/{namespace}/cloudpersistentvolumeclaims/{name}	cce:storage:delete	√	√
Listing all volumes	GET /storage/api/v1/namespaces/{namespace}/listvolumes	cce:storage:list	√	√

**Table 6-8** Add-on management actions

Permission	API	Action	IAM Project	Enterprise Project
Creating an add-on instance	POST /api/v3/addons	cce:addonInstance:create	√	√
Obtaining an add-on instance	GET /api/v3/addons/{id}?cluster_id={cluster_id}	cce:addonInstance:get	√	√
Listing all add-on instances	GET /api/v3/addons?cluster_id={cluster_id}	cce:addonInstance:list	√	√
Deleting an add-on instance	DELETE /api/v3/addons/{id}?cluster_id={cluster_id}	cce:addonInstance:delete	√	√
Updating an add-on instance	PUT /api/v3/addons/{id}	cce:addonInstance:update	√	√

**Table 6-9** Quota management actions

Permission	API	Action	IAM Project	Enterprise Project
Obtaining quota details	GET /api/v3/projects/{project_id}/quotas	cce:quota:get	√	√

# 7 Appendix

## 7.1 Status Code

[Table 7-1](#) describes the status codes.

**Table 7-1** Status code

Stat us Cod e	Code	Description
100	Continue	The server has received the initial part of the request and the client should continue to send the remaining part. It is issued on a provisional basis while request processing continues. It alerts the client to wait for a final response.
101	Switching Protocols	The requester has asked the server to switch protocols and the server has agreed to do so. The target protocol must be more advanced than the source protocol. For example, the current HTTP protocol is switched to a later version of HTTP.
201	Created	The request has been fulfilled, resulting in the creation of a new resource.
202	Accepted	The request has been accepted for processing, but the processing has not been completed.
203	Non-Authoritative Information	The server successfully processed the request, but is returning information that may be from another source.

Stat us Cod e	Code	Description
204	NoContent	The server has successfully processed the request, but does not return any content. The status code is returned in response to an HTTP OPTIONS request.
205	Reset Content	The server successfully processed the request, but is not returning any content.
206	Partial Content	The server has successfully processed a part of the GET request.
300	Multiple Choices	There are multiple options for the requested resource. For example, this code could be used to present a list of resource characteristics and addresses from which the client such as a browser may choose.
301	Moved Permanently	This and all future requests should be permanently directed to the given URI indicated in this response.
302	Found	The requested resource was temporarily moved.
303	See Other	The response to the request can be found under a different URI, and should be retrieved using a GET or POST method.
304	Not Modified	The requested resource has not been modified. In such case, there is no need to retransmit the resource since the client still has a previously-downloaded copy.
305	Use Proxy	The requested resource is available only through a proxy.
306	Unused	This HTTP status code is no longer used.
400	BadRequest	The request is invalid. The client should modify the request instead of re-initiating it.
401	Unauthorized	The authorization information provided by the client is incorrect or invalid.
402	Payment Required	This status code is reserved for future use.

Stat us Cod e	Code	Description
403	Forbidden	The server has received the request and understood it, but the server is refusing to respond to it. The client should modify the request instead of re-initiating it.
404	NotFound	The requested resource cannot be found. The client should modify the request instead of re-initiating it.
405	MethodNotAllowed	A request method is not supported for the requested resource. The client should modify the request instead of re-initiating it.
406	Not Acceptable	The server cannot fulfill the request based on the content characteristics of the request.
407	Proxy Authentication Required	This code is similar to 401, but indicates that the client must first authenticate itself with the proxy.
408	Request Time-out	The server timed out waiting for the request. The client may re-initiate the request without modifications at any later time.
409	Conflict	The request cannot be processed due to a conflict. This status code indicates that the resource that the client attempts to create already exists, or the request fails to be processed because of the update of the conflict request.
410	Gone	The requested resource cannot be found. The status code indicates that the requested resource has been deleted permanently.
411	Length Required	The server refused to process the request because the request does not specify the length of its content.
412	Precondition Failed	The server does not meet one of the preconditions that the requester puts on the request.



Stat us Cod e	Code	Description
413	Request Entity Too Large	The server refuses to process a request because the request entity is too large. The server may disable the connection to prevent the client from sending requests consecutively. If the server temporarily cannot process the request, the response will contain a Retry-After header field.
414	Request-URI Too Large	The URI provided was too long for the server to process.
415	Unsupported Media Type	The server does not support the media type in the request.
416	Requested range not satisfiable	The requested range is invalid.
417	Expectation Failed	The server fails to meet the requirements of the Expect request-header field.
422	UnprocessableEntity	The request is well-formed but is unable to be processed due to semantic errors.
429	TooManyRequests	The client sends excessive requests to the server within a given time (exceeding the limit on the access frequency of the client), or the server receives excessive requests within a given time (beyond its processing capability). In this case, the client should repeat requests after the time specified in the Retry-After header of the response expires.
500	InternalServerError	The server is able to receive the request but it could not understand the request.
501	Not Implemented	The server does not support the requested function.
502	Bad Gateway	The server acting as a gateway or proxy receives an invalid response from a remote server.
503	ServiceUnavailable	The requested service is invalid. The client should modify the request instead of re-initiating it.
504	ServerTimeout	The server could not return a timely response. The response will reach the client only if the request carries a timeout parameter.

Status Code	Code	Description
505	HTTP Version not supported	The server does not support the HTTP protocol version used in the request.

## 7.2 Error Codes

If an error occurs in API calling, no result is returned. Identify the cause based on the error code of each API. If an error occurs in API calling, HTTP status code 4xx or 5xx is returned. The response body contains the specific error code and information. If you fail to locate the cause of an error, contact customer service and provide the error code, so that we can help you solve the problem as soon as possible.

### Format of an Error Response Body

If an error occurs during API calling, an error code and a message will be displayed. The following shows an error response body.

```
{
  "errorMessage": "The format of message is error",
  "errorCode": "CCE.01400001"
}
```

In the preceding information, **errorCode** is an error code, and **errorMessage** describes the error.

### Error Code Description

Status Code	Error Code	Error Message	Description	Troubleshooting
400	CCE.01400001	Invalid request.	Invalid request body.	Modify the request body based on the returned message and the CCE API documentation, or contact technical support.
400	CCE.01400002	Subnet not found in the VPC.	No subnet is found in the VPC.	Check whether the subnet in the request body is in the corresponding VPC.

Status Code	Error Code	Error Message	Description	Troubleshooting
400	CCE.01400003	IPv6 not supported for the subnet.	The subnet does not support IPv6.	Use a subnet that supports IPv6.
400	CCE.01400004	No available flavors for master nodes.	There is no available flavor for the master node.	Change to another available cluster flavor or contact technical support.
400	CCE.01400005	Container network CIDR blocks conflict.	The container CIDR block conflicts with the existing one.	Check the container CIDR block based on the returned message.
400	CCE.01400006	Content type not supported.	The content type is invalid.	Use the supported content type by referring to the CCE API documentation.
400	CCE.01400007	Insufficient cluster quota.	Insufficient cluster quota.	Submit a service ticket to increase the cluster quota.
400	CCE.01400008	Insufficient server quota	Insufficient ECS quota.	Submit a service ticket to increase the ECS quota.
400	CCE.01400009	Insufficient CPU quota.	Insufficient ECS CPU quota.	Submit a service ticket to increase the ECS CPU quota.
400	CCE.01400010	Insufficient memory quota.	Insufficient ECS memory quota.	Submit a service ticket to increase the ECS memory quota.
400	CCE.01400011	Insufficient security group quota.	Insufficient security group quota.	Submit a service ticket to increase the security group quota.
400	CCE.01400012	Insufficient EIP quota.	Insufficient EIP quota.	Submit a service ticket to increase the EIP quota.

Status Code	Error Code	Error Message	Description	Troubleshooting
400	CCE.01400013	Insufficient volume quota.	Insufficient disk quota.	Submit a service ticket to increase the disk quota based on the returned message.
400	CCE.01400014	Excessive nodes in the cluster.	The number of nodes exceeds the cluster scale.	Submit a service ticket to change the cluster flavor.
400	CCE.01400015	Version not supported.	The cluster version is not supported.	Create a cluster of the supported version based on the returned message.
400	CCE.01400016	Current cluster type does not support this node flavor.	The current cluster type does not support this node flavor.	Use the correct node flavor based on the returned message.
400	CCE.01400017	No available container CIDR block found.	No available container CIDR block is found.	Use a correct container CIDR block based on the returned message.
400	CCE.01400018	This type of OS cannot be created in this CCE version.	The current CCE version does not support this type of OS.	Use a supported OS based on the returned message.
400	CCE.01400019	Insufficient resource tenant quota.	The quota of the resource tenant is insufficient.	Rectify the fault based on the returned message or contact technical support.
400	CCE.01400020	Insufficient VPC quota.	Insufficient VPC quota.	Rectify the fault based on the returned message or contact technical support.
400	CCE.01400021	No available flavors for nodes.	There is no available flavor for the node.	Change to another available node flavor or contact technical support.

Status Code	Error Code	Error Message	Description	Troubleshooting
400	CCE.01400022	No available node volumes for nodes.	There is no available EVS disk flavor for the node.	Change to another available EVS disk flavor or contact technical support.
400	CCE.01400023	operation conflict	Nodes cannot be created during cluster scale-out.	Try again later.
400	CCE.01400024	operation conflict	The cluster cannot be deleted during node creation.	Try again later.
400	CCE.01400025	Unsupported flavor with insufficient sub-ENI quota	The sub-ENI quota is insufficient. The VM flavor is not supported by CCE Turbo clusters.	Select a VM flavor whose sub-ENI quota is not 0.
400	CCE.01400033	Snapshot task already exists.	The cluster backup task already exists.	Wait until the cluster backup task is complete and try again.
400	CCE.02400001	Invalid request.	Invalid request body.	Modify the request body based on the returned message and the CCE API documentation, or contact technical support.
400	CCE.03400001	Invalid request.	Invalid request body.	Modify the request body based on the returned message and the CCE API documentation, or contact technical support.

Status Code	Error Code	Error Message	Description	Troubleshooting
400	CCE.03400002	Missing access key.	The access key is missing.	Ensure that the version of the installed or upgraded storage add-on is correct, or contact technical support.
401	CCE.01401001	Authorization failed.	Authentication failed.	Rectify the fault based on the returned message or contact technical support.
401	CCE.02401001	Authorization failed.	Authentication failed.	Rectify the fault based on the returned message or contact technical support.
401	CCE.03401001	Authorization failed.	Authentication failed.	Rectify the fault based on the returned message or contact technical support.
403	CCE.01403001	Forbidden.	Access denied.	Rectify the fault based on the returned message or contact technical support.
403	CCE.01403002	Current cluster status does not allow node pool to be deleted.	A node pool in the cluster that is in the current state cannot be deleted.	Wait until the cluster runs properly and try again.
403	CCE.01403003	Current node pool status does not allow node pool to be deleted.	The node pool in the current state cannot be deleted.	Wait until the node pool runs properly and try again.
403	CCE.01403005	Node pool cannot be deleted when it is scaling.	The node pool is being scaled out and cannot be deleted.	Wait until the node pool scale-out completes and try again.

Status Code	Error Code	Error Message	Description	Troubleshooting
403	CCE.01403006	Node pool cannot be deleted when exists installing or deleting nodes.	There are nodes that are being installed or deleted and the node pool cannot be deleted.	Wait until the nodes in the node pool are installed or deleted and try again.
403	CCE.02403001	Forbidden.	Access denied.	Rectify the fault based on the returned message or contact technical support.
403	CCE.03403001	Forbidden.	Access denied.	Rectify the fault based on the returned message or contact technical support.
404	CCE.01404001	Resource not found.	Resources not found.	Check whether the resource to be accessed has been deleted.
404	CCE.02404001	Resource not found.	Resources not found.	Check whether the resource to be accessed has been deleted.
404	CCE.03404001	Resource not found.	Resources not found.	Check whether the resource to be accessed has been deleted.
409	CCE.01409001	The resource already exists.	The resource already exists.	Delete the resource and try again.
409	CCE.01409002	Resource updated with out-of-date version.	An out-of-date version is used to update the target resource.	Ensure that the target resource version is the latest one or contact technical support.
409	CCE.02409001	The resource already exists.	The resource already exists.	Delete the resource and try again.

Status Code	Error Code	Error Message	Description	Troubleshooting
409	CCE.03409001	Addon instance has installed.	The add-on instance has been installed.	Delete the add-on instance and install it again.
429	CCE.01429002	Resource locked by other requests.	The resource is locked by another request.	Rectify the fault based on the returned message or contact technical support.
429	CCE.01429003	The concurrency limit of tasks has been reached.	The number of concurrent tasks has reached the upper limit.	Reduce the frequency of sending requests or contact technical support.
429	CCE.02429001	The throttling threshold has been reached.	The maximum number of requests has been reached.	Reduce the frequency of sending requests or contact technical support.
500	CCE.01500001	Internal error.	Internal error.	Rectify the fault based on the returned message or contact technical support.
500	CCE.02500001	Internal error.	Internal error.	Rectify the fault based on the returned message or contact technical support.
500	CCE.03500001	Internal error.	Internal error.	Rectify the fault based on the returned message or contact technical support.

## 7.3 Obtaining a Project ID

### Scenarios

A project ID is required for some URLs when APIs are called. Therefore, obtain a project ID in advance. Two methods are available:

- [Call an API](#)



- [Use the console](#)

## Obtaining the Project ID by Calling an API

You can obtain the project ID by calling the IAM API used to query project information based on the specified criteria.

The API for obtaining a project ID is **GET `https://{Endpoint}/v3/projects`**. `{Endpoint}` indicates the endpoint of IAM, which can be obtained from [Endpoints](#). For details about API authentication, see [Authentication](#).

The following is an example response. The value of `id` is the project ID.

```
{
  "projects": [
    {
      "domain_id": "65382450e8f64ac0870cd180d14e684b",
      "is_domain": false,
      "parent_id": "65382450e8f64ac0870cd180d14e684b",
      "name": "project_name",
      "description": "",
      "links": {
        "next": null,
        "previous": null,
        "self": "https://www.example.com/v3/projects/a4a5d4098fb4474fa22cd05f897d6b99"
      },
      "id": "a4a5d4098fb4474fa22cd05f897d6b99",
      "enabled": true
    }
  ],
  "links": {
    "next": null,
    "previous": null,
    "self": "https://www.example.com/v3/projects"
  }
}
```

## Obtaining a Project ID from the Console

To obtain a project ID from the console, perform the following operations:

1. Log in to the management console.
2. Hover over the username and choose **My Credentials** from the drop-down list.

On the **Projects** tab page, view project IDs.

## 7.4 Obtaining an Account ID

An account ID (domain-id) is required for some URLs when an API is called. To obtain an account ID, perform the following operations:

1. Sign up and log in to the management console.
2. Click the username in the upper right corner and choose **My Credentials** from the drop-down list.

On the **API Credentials** page, view the account ID.

## 7.5 Specifying Add-ons to Be Installed During Cluster Creation

When creating a cluster, you can add a key-value pair to **annotations** of the **metadata** field in the request body to install an add-on in the cluster. The key is **cluster.install.addons/install**, and the value is a JSON array of **AddonTemplate**.

**Table 7-2** Value data structure

Parameter	Mandatory	Type	Description
Value	Yes	JSON array of AddonTemplate string	Add-on to be installed in the cluster. If this parameter is not specified, the CoreDNS and CCE Container Storage (Everest) add-ons are installed in the cluster by default. For details about the parameter values, see <a href="#">Table 7-3</a> .

**Table 7-3** Data structure of the AddonTemplate field

Parameter	Mandatory	Type	Description
addonTemplateName	Yes	String	Add-on name. Options: <ul style="list-style-type: none"> <li>• <b>coredns</b>: Install CoreDNS.</li> <li>• <b>everest</b>: Install CCE Container Storage (Everest).</li> <li>• <b>node-local-dns</b>: Install NodeLocal DNSCache.</li> <li>• <b>volcano</b>: Install Volcano Scheduler.</li> <li>• <b>npd</b>: Install CCE Node Problem Detector.</li> </ul>
version	No	String	Add-on version. To view the add-on version, log in to the CCE console, choose <b>Add-ons</b> in the navigation pane, click the name of the target add-on, and check the change history. If this parameter is left blank, the latest version is used by default.

Parameter	Mandatory	Type	Description
values	No	Json Map	<ul style="list-style-type: none"> <li>CoreDNS: For details about the parameters required for installing the add-on, see <a href="#">CoreDNS</a>.</li> </ul> <p><b>NOTE</b> You do not need to specify this parameter when installing the CCE Container Storage (Everest), NodeLocal DNSCache, Volcano Scheduler, or CCE Node Problem Detector add-on.</p>

### Example Request

Create a cluster using a VPC network and install CoreDNS and CCE Container Storage (Everest) in it.

```
{
  "kind": "Cluster",
  "apiVersion": "v3",
  "metadata": {
    "name": "test",
    "annotations": {
      "cluster.install.addons.external/install": "[{\"addonTemplateName\":\"icagent\",\"extendParam\":{\"logSwitch\":\"false\",\"tDSEnable\":\"false\"}}]",
      "cluster.install.addons/install": "[{\"addonTemplateName\":\"coredns\",\"values\":{\"flavor\":\"is_default\":false,\"name\":\"2500\",\"recommend_cluster_flavor_types\":[\"small\"],\"replicas\":\"2\",\"resources\":{\"limitsCpu\":\"500m\",\"limitsMem\":\"512Mi\",\"name\":\"coredns\",\"replicas\":\"2\",\"requestsCpu\":\"500m\",\"requestsMem\":\"512Mi\"}},\"size\":\"small\",\"category\":[\"CCE\",\"Turbo\"]}],\n{\"addonTemplateName\":\"everest\"}]"
    }
  },
  "spec": {
    "category": "CCE",
    "flavor": "cce.s1.small",
    "version": "v1.29",
    "type": "VirtualMachine",
    "hostNetwork": {
      "vpc": "*****",
      "subnet": "*****"
    },
    "containerNetwork": {
      "mode": "vpc-router",
      "cidrs": [
        {
          "cidr": "10.0.0.0/16"
        }
      ]
    },
    "ipv6enable": false,
    "description": "",
    "billingMode": 0,
    "kubeProxyMode": "iptables",
    "extendParam": {
      "alpha.cce/fixPoolMask": "25",
      "enterpriseProjectId": "0"
    },
    "authentication": {
      "mode": "rbac"
    },
    "configurationsOverride": [
      {

```

```
    "name": "kube-apiserver",
    "configurations": [
      {
        "name": "support-overload",
        "value": true
      }
    ]
  },
  "deletionProtection": false,
  "serviceNetwork": {
    "IPv4CIDR": "10.247.0.0/16"
  }
}
```

## 7.6 How to Obtain Parameters in the API URI

### Obtaining a Project ID (`project_id`)

`project_id` indicates the project ID, which can be obtained from the console or APIs. For details, see [Obtaining a Project ID](#).

### Obtaining a Cluster ID (`cluster_id`)

**Step 1** Log in to the CCE console. In the navigation pane, choose **Clusters**.

**Step 2** Click the name of the created cluster. The cluster details page is displayed. Obtain the cluster ID.

----End

### Obtaining a Node ID (`node_id`)

**Step 1** Log in to the CCE console. In the navigation pane, choose **Clusters**.

**Step 2** Click the name of the created cluster to access the cluster console. In the navigation pane, choose **Nodes**. On the **Nodes** tab, move the cursor to the node name and view the node ID.

----End

### Obtaining a Node Pool ID (`nodepool_id`)

**Step 1** Log in to the CCE console. In the navigation pane, choose **Clusters**.

**Step 2** Click the name of the created cluster to access the cluster console. In the navigation pane, choose **Nodes**. On the **Node Pools** tab, move the cursor to the node pool name and view the node pool ID.

----End


### Obtaining a Job ID (`job_id`)

**Step 1** Log in to the CCE console. In the navigation pane, choose **Clusters**. The following uses cluster management as an example to describe how to obtain the `job_id` of the cluster that is being created.

**Step 2** Obtain the job ID.

1. If you are using Google Chrome, press **F12**. On the pane displayed on the right, click the **Network** tab.
2. Click **Operation Records** on the CCE console to view details about cluster operation records.
3. Enter **jobs** in the **Filter** text box to filter out the jobs. Select a job from the list on the left and click **Preview**. The **uid** field indicates the job UID.

**Figure 7-1** Obtaining the job ID



----End

## 7.7 Creating a VPC and Subnet

### Context

To provide a secure and isolated network environment for CCE, create a VPC before creating a cluster.

If you have already created a VPC, you do not need to create it again.

### Creating a VPC

- Step 1** On the management console, click **Service List**, and choose **Network > Virtual Private Cloud** to launch the VPC console.
- Step 2** On the VPC console, click **Create VPC** to create a VPC.
- Step 3** The created VPC is displayed in the list. Click its name and obtain the VPC ID, which will be required in [cluster creation](#).

----End

### Creating a Subnet

- Step 1** On the management console, click **Service List**, and choose **Network > Virtual Private Cloud** to launch the VPC console.
- Step 2** On the **VPC** console, choose **Virtual Private Cloud > Subnets** in the navigation pane, and click **Create Subnet** in the upper right corner.
- Step 3** Create a subnet as prompted and click its name to obtain the network ID, which will be required in [cluster creation](#).

----End

## 7.8 Creating a Key Pair

### Context

Create a key pair before you create a container cluster. Key pairs are used for user identity authentication upon login to a worker node.

If you have already created a key pair, you do not need to create it again.

### Procedure

- Step 1** On the management console, click **Service List**, and choose **Compute > Elastic Cloud Server** to go to the ECS console.
- Step 2** In the navigation pane, choose **Key Pair**.
- Step 3** Click **Create Key Pair** and create a key pair as prompted.
- Step 4** After the creation is complete, a key file in the **.pem** format is generated and automatically saved to the default directory on your local computer.

----End

## 7.9 Node Flavor Description

### NOTE

Different regions support different node flavors, and node flavors may be changed or sold out. You are advised to log in to the CCE console and check whether the required node flavors are supported on the page for creating nodes.

- **CCE standard cluster**

CCE clusters support only 2 vCPUs and 4 GiB or higher specifications. You are advised to query node specifications on the console. For details about node specifications, see ECS Specifications.

You need to enter the specific flavor name, for example, **c6.large.2**.

## 7.10 Adding a Salt in the password Field When Creating a Node

When a node is created through the API, you need to add a salt to the **password** field to safeguard the password. The procedure is as follows:

 NOTE

The salt must be set based on the password complexity requirements:

- A string of 8–26 characters.
- Contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters !@\$%^\_-=+[{]};,./?
- Cannot contain the username or the username spelled backwards.
- Cannot contain the username, the username spelled backwards, or more than two consecutive characters in the username (for Windows EC2s).

## Python

**To salt a password in the Python 3.7.7 environment, perform the following steps:**

```
pip install passlib
python -c "import base64; from passlib.hash import sha512_crypt; salted_password =
base64.b64encode(sha512_crypt.hash('*****', salt='salt', rounds=5000).encode()).decode();
print(salted_password)"
```

 NOTE

The **python crypt** package has compatibility issues in macOS. If the package cannot be executed, run it in Linux.

## Java

**To salt a password in the Java environment, perform the following steps:**

1. Obtain a random number as the salt.

```
private static String getCharAndNumr(int length) {
    String val = "";
    Random random = new SecureRandom();
    for (int i = 0; i < length; i++) {
        // Indicates whether to output letters or digits.
        String charOrNum = random.nextInt(2) % 2 == 0 ? "char" : "num";
        // Character string
        if ("char".equalsIgnoreCase(charOrNum)) {
            // Indicates whether an upper-case or lower-case letter is obtained.
            int choice = random.nextInt(2) % 2 == 0 ? 65 : 97;
            val += (char) (choice + random.nextInt(26));
        } else if ("num".equalsIgnoreCase(charOrNum)) { // Digit
            val += String.valueOf(random.nextInt(10));
        }
    }
    return val;
}
```

2. Generate a salt.

```
private static String generateSalt() {
    String salt;
    try {
        salt = "$6$" + getCharAndNumr(16);
    } catch (Exception e) {
        salt = defaultSalt;
    }

    return salt;
}
```

3. Generate a ciphertext password based on the salt.

```
public static String getSaltPassword(String password) {
    if (StringUtil.isBlank(password)) {
        throw new BizException("password is empty");
    }
}
```

```
    }  
  
    String salt = generateSalt();  
  
    Crypt crypt = new Crypt();  
    return crypt.crypt(password, salt);  
}
```

4. Encode the value of the **password** field using Base64.

```
(Base64.getEncoder().encodeToString(AddSaltPasswordUtil.getSaltPassword(cceNodeCreateVo.getPassword().getBytes()))
```

5. A complete example is as follows:

```
import java.util.Base64;  
import java.util.Random;  
import java.security.SecureRandom;  
  
import org.apache.commons.codec.digest.Crypt;  
import org.apache.commons.lang.StringUtils;  
  
public class PassWord {  
  
    static String defaultSalt = null;  
  
    public static void main(String[] args) throws Exception {  
        System.out.println(Base64.getEncoder().encodeToString(PassWord.getSaltPassword("Custom  
password").getBytes()));  
    }  
  
    // Generate a ciphertext password based on the salt.  
    public static String getSaltPassword(String password) throws Exception {  
        if(StringUtils.isBlank(password)) {  
            throw new Exception("password is empty");  
        }  
        String salt = generateSalt();  
        return Crypt.crypt(password, salt);  
    }  
  
    //Generate a salt.  
    private static String generateSalt() {  
        String salt;  
        try {  
            salt = "$6$" + getCharAndNumr(16);  
        }catch (Exception e){  
            salt = defaultSalt;  
        }  
        return salt;  
    }  
  
    // Obtain a random number as the salt.  
    private static String getCharAndNumr(int length) {  
        String val = "";  
        Random random = new SecureRandom();  
        for (int i = 0; i < length; i++) {  
            // Indicates whether to output letters or digits.  
            String charOrNum = random.nextInt(2) % 2 == 0 ? "char" : "num";  
            // Character string  
            if ("char".equalsIgnoreCase(charOrNum)) {  
                // Indicates whether an upper-case or lower-case letter is obtained.  
                int choice = random.nextInt(2) % 2 == 0 ? 65 : 97;  
                val += (char) (choice + random.nextInt(26));  
            } else if ("num".equalsIgnoreCase(charOrNum)) { // Digit  
                val += String.valueOf(random.nextInt(10));  
            }  
        }  
        return val;  
    }  
}
```



## Go

You can use either of the following methods to salt passwords for the Go language:

- <https://github.com/amoghe/go-crypt>
- <https://github.com/GehirnInc/crypt>

## 7.11 Maximum Number of Pods That Can Be Created on a Node

### Calculation of the Maximum Number of Pods on a Node

The maximum number of pods that can be created on a node is calculated based on the cluster type:

Network Model	Maximum Number of Pods on a Node	Recommended Configuration
Tunnel network	<a href="#">Maximum number of pods on a node</a>	None
VPC network	The smaller value between the <a href="#">maximum number of pods on a node</a> and the <a href="#">number of container IP addresses that can be allocated on a node</a>	To ensure new pods run smoothly on a node, make sure that the number of pods on the node does not exceed the number of container IP addresses that can be assigned to it. If there are not enough container IP addresses available on the node, the new pods may not function properly.

### Number of Allocatable Container IP Addresses on a Node

When creating a cluster in the VPC network model, specify the number of container IP addresses that can be allocated on each node using `alpha.cce/fixPoolMask` based on the rules for managing container IP addresses.

The maximum number of pods that can be created on a node is determined by the number of container IP addresses that can be allocated to it. In a [container network](#), each pod needs its own IP address. If there are not enough pre-allocated container IP addresses on the node, pods cannot be created. If `hostNetwork: true` is configured in the YAML file, pods will use the [host network](#) instead of the allocatable container IP addresses. For details, see [Pod IP Address Allocation Differences Between the Container Network and Host Network](#).

By default, a node occupies three container IP addresses (network address, gateway address, and broadcast address). Therefore, the number of container IP

addresses that can be allocated to a node equals the number of selected container IP addresses minus 3.

## Maximum Number of Pods on a Node

When creating a node, you can configure the maximum number of pods (maxPods) that can be created on the node. This parameter is a configuration item of kubelet and determines the maximum number of pods that can be created by kubelet.

### NOTICE

For nodes in the default node pool (**DefaultPool**), the maximum number of pods cannot be changed after the nodes are created.

After a node in a custom node pool is created, you can modify the **max-pods** parameter in the node pool configuration to change the maximum number of pods on the node.

**Table 7-4** lists the default maximum number of pods on a node based on node specifications.

**Table 7-4** Default maximum number of pods on a node

Memory	Max. Pods
4 GB	20
8 GB	40
16 GB	60
32 GB	80
64 GB or above	110

## Pod IP Address Allocation Differences Between the Container Network and Host Network

When creating a pod, you can select the container network or host network for the pod.

- Container network (default): **Each pod is assigned an IP address by the cluster networking add-ons, which occupies the IP addresses of the container network.**
- Host network: Pods with **hostNetwork: true** configured directly use the network of the host. After the configuration, the pods use the ports on the host and their IP address is identical to that of the host, **without relying on the IP addresses of the container network.** When using a host network, avoid conflicts between the pod ports and service ports on the host. Use a host network only if a particular application requires access to a specific port on the host.

## 7.12 Node OS

### Node OSs and Cluster Types

Table 7-5 Mapping between node OSs and cluster types

OS	Cluster Version	CCE Standard Cluster	
		VPC Network	Tunnel Network
Ubuntu 22.04	v1.30	√	×
	v1.29	√	×
	v1.28	√	×
	v1.27	√	×
	v1.25	√	×
CentOS Linux release 7.6	v1.30	√	√
	v1.29	√	√
	v1.28	√	√
	v1.27	√	√
	v1.25	√	√
	v1.23	√	√
	v1.21	√	√

## 7.13 Space Allocation of a Data Disk

This section describes how to allocate data disk space to nodes so that you can configure the data disk space accordingly.

### Allocating Default Data Disk Space

 NOTE

- In clusters of a version earlier than v1.23.18-r0, v1.25.13-r0, v1.27.10-r0, v1.28.8-r0, or v1.29.4-r0, a default data disk will be added to the node its container runtime and kubelet components. You can customize the space allocation of the default data disk.
- In clusters of v1.23.18-r0, v1.25.13-r0, v1.27.10-r0, v1.28.8-r0, v1.29.4-r0, or later, a default data disk will be added to the node for its container runtime and kubelet components only if **System Component Storage** is set to **Data Disk**. You can customize the space allocation of the default data disk.

When creating a node, you can customize Data Disk Space Allocation in the expanded area of **Data Disk**.

- **Space Allocation for Container Engines**
  - Specified disk space: CCE divides the data disk space for two parts by default. One part is used to store the Docker/containerd working directories, container image data, and image metadata. The other is reserved for kubelet and emptyDir volumes. The available container engine space affects image pulls and container startup and running.
    - Container engine and container image space (90% by default): stores the container runtime working directories, container image data, and image metadata.
    - kubelet and emptyDir space (10% by default): stores pod configuration files, secrets, and mounted storage such as emptyDir volumes.

#### NOTE

If the sum of the container engine and container image space and the kubelet and emptyDir space is less than 100%, the remaining space will be allocated for user data. You can mount the storage volume to a service path. Do not leave the path empty or set it to a key OS path such as the root directory.

- **Space Allocation for Pods:** indicates the basesize of a pod. You can set an upper limit for the disk space occupied by each workload pod (including the space occupied by container images). This setting prevents the pods from taking all the disk space available, which may cause service exceptions. It is recommended that the value is less than or equal to 80% of the container engine space. This parameter is related to the node OS and container storage rootfs and is not supported in some scenarios. For details, see [Mapping Between OS and Container Storage Rootfs](#).
- Write Mode
  - **Linear:** A linear logical volume integrates one or more physical volumes. Data is written to the next physical volume when the previous one is used up.
  - **Striped:** available only if there are at least two data disks. A striped logical volume stripes data into blocks of the same size and stores them in multiple physical volumes in sequence. This allows data to be concurrently read and written. A storage pool consisting of striped volumes cannot be scaled-out.

## Space Allocation for Container Engines

For a node using a non-shared data disk (100 GiB for example), the division of the disk space varies depending on the container storage Rootfs type **Device Mapper** or **OverlayFS**. For details about the container storage Rootfs corresponding to different OSs, see [Mapping Between OS and Container Storage Rootfs](#).

- **Rootfs (Device Mapper)**


By default, the container engine and image space, occupying 90% of the data disk, can be divided into the following two parts:

- The **/var/lib/docker** directory is used as the Docker working directory and occupies 20% of the container engine and container image space by default. (Space size of the **/var/lib/docker** directory = **Data disk space x 90% x 20%**)

- The thin pool is used to store container image data, image metadata, and container data, and occupies 80% of the container engine and container image space by default. (Thin pool space = **Data disk space x 90% x 80%**)

The thin pool is dynamically mounted. You can view it by running the **lsblk** command on a node, but not the **df -h** command.


**Figure 7-2** Space allocation for container engines of Device Mapper



- **Rootfs (OverlayFS)**

No separate thin pool. The entire container engine and container image space (90% of the data disk by default) are in the **/var/lib/docker** directory.

**Figure 7-3** Space allocation for container engines of OverlayFS



## Space Allocation for Pods

The customized pod container space (basesize) is related to the node OS and container storage Rootfs. For details about the container storage Rootfs, see [Mapping Between OS and Container Storage Rootfs](#).

- Device Mapper supports custom pod basesize. The default value is 10 GiB.
- In OverlayFS mode, the pod container space is not limited by default.

When configuring **basesize**, consider the maximum number of pods allowed on one node. The container engine space should be greater than the total disk space used by containers. Formula: **Container engine space and container image space (90% by default) > Number of containers x basesize**. Otherwise, the container engine space allocated to the node may be insufficient and the container cannot be started.

For nodes that support **basesize**, when Device Mapper is used, although you can limit the size of the **/home** directory of a single container (to 10 GiB by default), all containers on the node still share the thin pool of the node for storage. They are not completely isolated. When the sum of the thin pool space used by certain containers reaches the upper limit, other containers cannot run properly.

In addition, after a file is deleted in the **/home** directory of the container, the thin pool space occupied by the file is not released immediately. Therefore, even if

**basesize** is set to 10 GiB, the thin pool space occupied by files keeps increasing until 10 GiB when files are created in the container. The space released after file deletion will be reused but after a while. If **the number of containers on the node multiplied by basesize** is greater than the thin pool space size of the node, there is a possibility that the thin pool space has been used up.

## Mapping Between OS and Container Storage Rootfs

Table 7-6 Node OSs and container engines in CCE clusters

OS	Container Storage Rootfs	Custom Basesize
CentOS 7.x	Clusters of v1.19.16 and earlier use Device Mapper. Clusters of v1.19.16 and later use OverlayFS.	Supported when Rootfs is set to Device Mapper and the runtime is Docker. The default value is 10 GiB. If Rootfs is set to OverlayFS, the basesize cannot be customized.
Ubuntu 22.04	OverlayFS	Not supported

## Garbage Collection Policies for Container Images

When the container engine space is insufficient, image garbage collection is triggered.

The policy for garbage collecting images takes two factors into consideration: **HighThresholdPercent** and **LowThresholdPercent**. Disk usage exceeding the high threshold (default: 80%) will trigger garbage collection. The garbage collection will delete least recently used images until the low threshold (default: 70%) is met.

## Recommended Configuration for the Container Engine Space

- The container engine space should be greater than the total disk space used by containers. Formula: **Container engine space > Number of containers x basesize**
- You are advised to create and delete files of containerized services in local storage volumes (such as emptyDir and hostPath volumes) or cloud storage directories mounted to the containers. In this way, the thin pool space is not occupied. emptyDir volumes occupy the kubelet space. Therefore, properly plan the size of the kubelet space.
- You can deploy services on nodes that use the OverlayFS (for details, see [Mapping Between OS and Container Storage Rootfs](#)) so that the disk space occupied by files created or deleted in containers can be released immediately.

## 7.14 Attaching Disks to a Node

### Challenges

In disk planning and striped logical disk creation, it is difficult for users to flexibly attach and partition disks when creating a node.

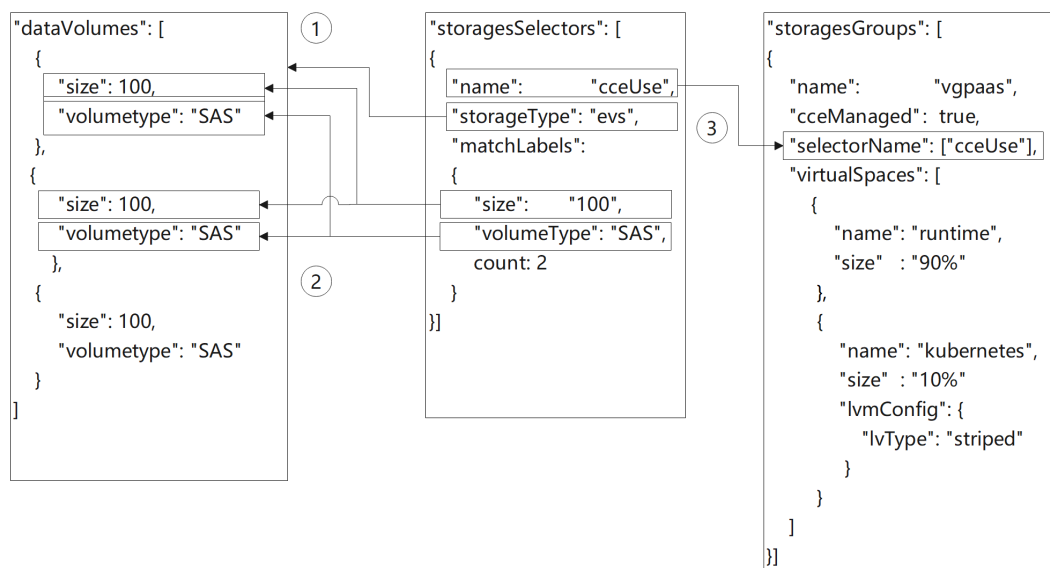
During node creation, the **storage** field selects a data disk based on parameters, such as the disk size and disk type, to prevent failures in node creation, resetting, migration, and management caused by drive letter matching failures.

### Solution

This section provides details about the **storage** field used in node creation so that you can implement complex disk selection and partitioning by calling the node creation API.

The **storage** field consists of **storageSelectors** and **storageGroups**. The **storageSelectors** field is responsible for disk selection, and the **storageGroups** field is responsible for disk processing.

The basic logic for field matching is as follows:

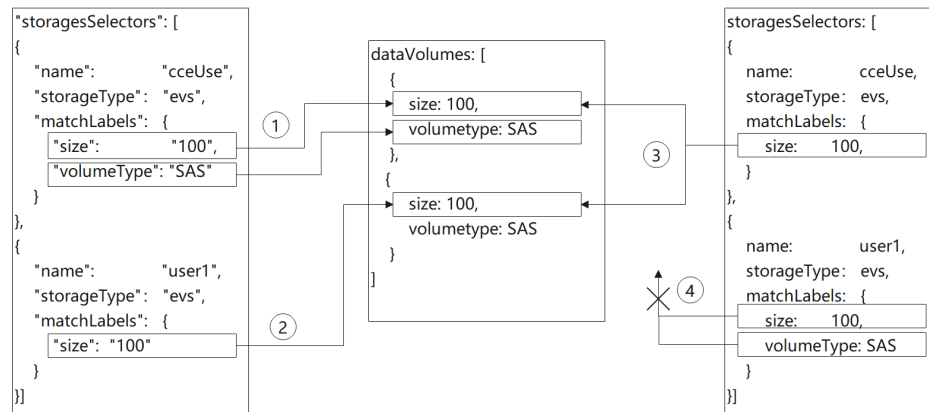


1. **storageSelectors** selects an EVS disk or a local disk based on the value of **storageType**.
  - a. Local disks do not support exact match. All local disks will be selected as data disks.

If some local disks need to be reserved, occupy the disks in the pre-installation script. The script is similar to the following:

```
# prepare
vgName=vg-test
storageDevice=/dev/vdb
# vgcreate
vgcreate ${vgName} ${storageDevice}
```

- b. EVS disks match the disks created in **dataVolumes** based on the settings of **matchLabels**.
- 2. Policies have priorities to match **matchLabels**. The **matchLabels** policy nearest to **storageSelectors** has a higher priority and the disk nearest to **dataVolumes** will be preferentially matches. As **matchLabels** uses loose matching, you are advised to place the **matchLabels** policy with a small matching range on the top. For example:

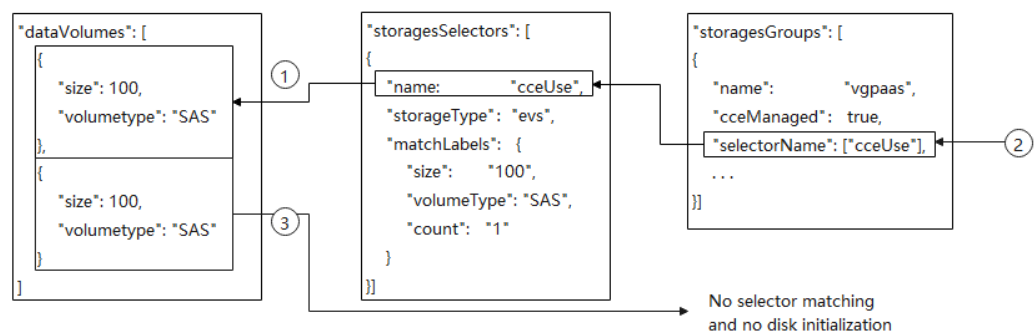


- a. In step 1, the first disk in **dataVolumes** is matched by the EVS disk whose size is 100 GiB and storage class is SAS. In step 2, the second disk in **dataVolumes** is matched by the EVS disk whose size is 100 GiB because the first disk has been selected.
- b. In step 3, two disks in **dataVolumes** can be matched because **volumeType** or **count** is not specified in **matchLabels**. In this case, no disk is available for matching in step 4.
- 3. **storageGroups** associates with **storageSelectors** based on **selectorName**. Finally, two 100 GiB disks are selected. The CCE backend groups the two PVs into a volume group (VG) and divides the VG into two logical volumes (LVs) in the ratio of 9:1. 10% of Kubernetes LVs are partitioned in striped mode. 90% runtime LVs are partitioned in linear mode by default because runtimeConfig is not configured.

## Creating a Raw Disk

During node creation on the CCE console, click **Add Data Disk**. Then, click **Expand** next to the newly added data disk and select **Default**. The created disk is a raw disk.

The following figure shows the API calling logic.





1. The **cceUse** selector matches a 100 GiB data disk.
2. The selected disk is managed by CCE and used as a data disk.
3. The other 100 GiB data disk created in **dataVolumes** is not selected by any selector and is managed by storageGroups. Therefore, this EVS disk will be attached to the node as a raw disk and will not be initialized.

After the node is created, log in to the node and check whether a 100 GiB disk has been attached but not initialized.

```
[root@test-83790 ~]# lsblk -n
sda                                     8:0    0   50G  0 disk
├─sda1                                  8:1    0   50G  0 part /
└─sdb                                    8:16   0  100G  0 disk
   └─vgpaas-dockersys                    253:0  0   18G  0 lvm  /var/lib/docker
      └─vgpaas-thinpool_tmeta             253:1  0    3G  0 lvm
         └─vgpaas-thinpool                253:3  0   67G  0 lvm
            ├─docker-253:0-786433-7cb37dc21202bfe2fc78dd1d33b70571e7e1982e56a4118f6facdc630cbc8b38 253:5  0   10G  0 dm  /var/lib/docker/devicemapper/
            ├─docker-253:0-786433-e17cd8670b9f423eaff34b92bd82a2e620118227c26da2e41eda7894361c9942 253:6  0   10G  0 dm  /var/lib/docker/devicemapper/
            ├─docker-253:0-786433-0dedb47e75eed3f635ce2d47c584587ae622c70dcb0eafeade9e14693a3146a0 253:7  0   10G  0 dm  /var/lib/docker/devicemapper/
            ├─docker-253:0-786433-93ed7e6e14313d13ecfa1152937b153fe599c48cfdaf9ecd43c1c36cae89a38a 253:8  0   10G  0 dm  /var/lib/docker/devicemapper/
            ├─docker-253:0-786433-e706be08bf5c6249850a09e080cf43d9a7be499eae33aa8feb06c027d26fa1e9 253:9  0   10G  0 dm  /var/lib/docker/devicemapper/
            └─docker-253:0-786433-5ecc4420da9a58fb66108db599a8267af3e8856da86b9c3d7fb82090a8781ae8 253:10 0   10G  0 dm  /var/lib/docker/devicemapper/
   └─vgpaas-thinpool_tdata                253:2  0   67G  0 lvm
      └─vgpaas-thinpool                  253:3  0   67G  0 lvm
         ├─docker-253:0-786433-7cb37dc21202bfe2fc78dd1d33b70571e7e1982e56a4118f6facdc630cbc8b38 253:5  0   10G  0 dm  /var/lib/docker/devicemapper/
         ├─docker-253:0-786433-e17cd8670b9f423eaff34b92bd82a2e620118227c26da2e41eda7894361c9942 253:6  0   10G  0 dm  /var/lib/docker/devicemapper/
         ├─docker-253:0-786433-0dedb47e75eed3f635ce2d47c584587ae622c70dcb0eafeade9e14693a3146a0 253:7  0   10G  0 dm  /var/lib/docker/devicemapper/
         ├─docker-253:0-786433-93ed7e6e14313d13ecfa1152937b153fe599c48cfdaf9ecd43c1c36cae89a38a 253:8  0   10G  0 dm  /var/lib/docker/devicemapper/
         └─docker-253:0-786433-e706be08bf5c6249850a09e080cf43d9a7be499eae33aa8feb06c027d26fa1e9 253:9  0   10G  0 dm  /var/lib/docker/devicemapper/
   └─vgpaas-kubernetes                    253:4  0   10G  0 lvm  /mnt/paas/kubernetes/kubelet
sdc                                     8:32   0  100G  0 disk
```

The following is an API example:

```
{
  "kind": "Node",
  "apiVersion": "v3",
  "metadata": {
    "name": "test-83790"
  },
  "spec": {
    "flavor": "c3.large.2",
    "az": "",
    "os": "EulerOS 2.9",
    "dataVolumes": [
      {
        "size": 100,
        "volumetype": "SAS"
      },
      {
        "size": 100,
        "volumetype": "SAS"
      }
    ],
    "billingMode": 0,
    "extendParam": {
      "maxPods": 110
    },
    "nodeNicSpec": {
      "primaryNic": {
        "subnetId": "ca964acf-8468-4735-8229-97940ef6c881"
      }
    },
    "rootVolume": {
      "size": 50,
      "volumetype": "SAS"
    },
    "runtime": {
      "name": "docker"
    },
    "login": {
      "userPassword": {
        "username": "root",
        "password": "*****"
      }
    }
  }
}
```

```

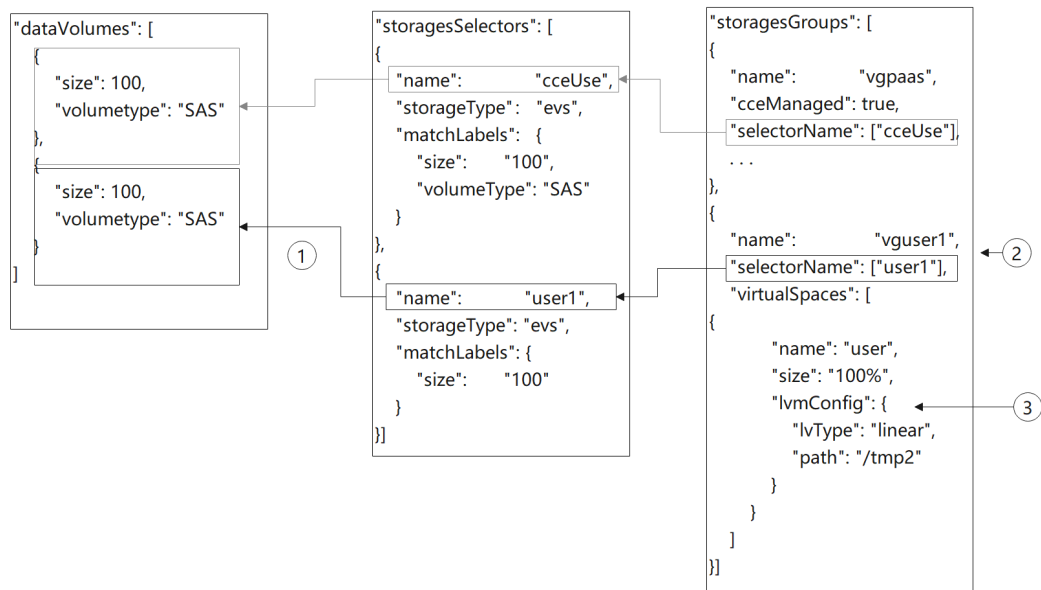
    }
  },
  "storage": {
    "storageSelectors": [
      {
        "name": "cceUse",
        "storageType": "evs",
        "matchLabels": {
          "size": "100",
          "volumeType": "SAS",
          "count": "1"
        }
      }
    ],
    "storageGroups": [
      {
        "name": "vgpaas",
        "selectorNames": [
          "cceUse"
        ],
        "cceManaged": true,
        "virtualSpaces": [
          {
            "name": "runtime",
            "size": "90%"
          },
          {
            "name": "kubernetes",
            "size": "10%"
          }
        ]
      }
    ]
  },
  "count": 1
}

```

## Attaching a Disk to a Specified Path

During node creation on the CCE console, click **Add Data Disk**. Then, click **Expand** next to the newly added data disk, select **Mount Disk**, and set the mount path. In this case, CCE initializes and attaches the disk by default.

The following figure shows the API calling logic.



1. The **user1** selector selects a 100 GiB data disk.
2. Create a VG named **vguser1** using LVM.
3. Strip all the space of **vguser1** into an LV named **user** and format the disk in ext4 format. Finally, attach the disk to the **/tmp2** directory.

After the node is created, log in to the node and check whether a 100 GiB disk has been attached and managed by LVM.

```
[root@test-37106 ~]# lsblk -n
sda                                8:0    0 50G 0 disk
├─sda1                              8:1    0 50G 0 part /
└─sdb                                8:16   0 100G 0 disk
   └─vgpaas-dockersys                253:0  0 16G 0 lvm  /var/lib/docker
      └─vgpaas-thinpool_tmeta         253:1  0 3G  0 lvm
         └─vgpaas-thinpool            253:3  0 59G 0 lvm
            ├─docker-253:0-917505-3a36be80c1a49db5da9639d222f19ce5983489080a36efdda1f17fa2d0bb7da9 253:6  0 10G 0 dm  /var/lib/docker/devicemapper/
            ├─docker-253:0-917505-46a876d16929a54d4f5ea97da81c3603c79cd5630be1c1010b476387a5d3c086 253:7  0 10G 0 dm  /var/lib/docker/devicemapper/
            ├─docker-253:0-917505-93081c85109968299fdca13a077e82252e725a6e37cae7299841db482656b815 253:8  0 10G 0 dm  /var/lib/docker/devicemapper/
            ├─docker-253:0-917505-513c5bda896de61ac85d917366da4ea4d78ab9f87cd4caae9e465badc0003c62 253:9  0 10G 0 dm  /var/lib/docker/devicemapper/
            ├─docker-253:0-917505-a6ac0d3ae8bffb57a92e6812079e503db49942619d5bbc69bb9516b31e15e67 253:10 0 10G 0 dm  /var/lib/docker/devicemapper/
            └─docker-253:0-917505-f9dfa31cdc3eb514a797c98311372ac8497d9a99581acdfefff0114bdf8e525 253:11 0 10G 0 dm  /var/lib/docker/devicemapper/
         └─vgpaas-thinpool_tdata       253:2  0 59G 0 lvm
            └─vgpaas-thinpool         253:3  0 59G 0 lvm
               ├─docker-253:0-917505-3a36be80c1a49db5da9639d222f19ce5983489080a36efdda1f17fa2d0bb7da9 253:6  0 10G 0 dm  /var/lib/docker/devicemapper/
               ├─docker-253:0-917505-46a876d16929a54d4f5ea97da81c3603c79cd5630be1c1010b476387a5d3c086 253:7  0 10G 0 dm  /var/lib/docker/devicemapper/
               ├─docker-253:0-917505-93081c85109968299fdca13a077e82252e725a6e37cae7299841db482656b815 253:8  0 10G 0 dm  /var/lib/docker/devicemapper/
               ├─docker-253:0-917505-513c5bda896de61ac85d917366da4ea4d78ab9f87cd4caae9e465badc0003c62 253:9  0 10G 0 dm  /var/lib/docker/devicemapper/
               ├─docker-253:0-917505-a6ac0d3ae8bffb57a92e6812079e503db49942619d5bbc69bb9516b31e15e67 253:10 0 10G 0 dm  /var/lib/docker/devicemapper/
               └─docker-253:0-917505-f9dfa31cdc3eb514a797c98311372ac8497d9a99581acdfefff0114bdf8e525 253:11 0 10G 0 dm  /var/lib/docker/devicemapper/
         └─vgpaas-kubernetes           253:4  0 20G  0 lvm  /mnt/paas/kubernetes/kubelet
            └─sdc                      8:32   0 100G 0 disk
               └─vguser1-user          253:5  0 100G 0 lvm  /tmp2
[root@test-37106 ~]#
```

The following is an API example. There are two data disks. One is used by CCE, and the other is mounted to the **/tmp2** directory.

```
{
  "kind": "Node",
  "apiVersion": "v3",
  "metadata": {
    "name": "test-37106"
  },
  "spec": {
    "flavor": "c3.large.2",
    "az": "",
    "os": "EulerOS 2.9",
    "dataVolumes": [
      {
        "size": 100,
        "volumetype": "SAS"
      }
    ]
  },
}
```

```

    {
      "size": 100,
      "volumetype": "SAS"
    }
  ],
  "billingMode": 0,
  "extendParam": {
    "maxPods": 110
  },
  "nodeNicSpec": {
    "primaryNic": {
      "subnetId": "ca964acf-8468-4735-8229-97940ef6c881"
    }
  },
  "rootVolume": {
    "size": 50,
    "volumetype": "SAS"
  },
  "runtime": {
    "name": "docker"
  },
  "login": {
    "userPassword": {
      "username": "root",
      "password": "*****"
    }
  },
  "storage": {
    "storageSelectors": [
      {
        "name": "cceUse",
        "storageType": "evs",
        "matchLabels": {
          "size": "100",
          "volumeType": "SAS",
          "count": "1"
        }
      },
      {
        "name": "user1",
        "storageType": "evs",
        "matchLabels": {
          "size": "100",
          "volumeType": "SAS",
          "count": "1"
        }
      }
    ],
    "storageGroups": [
      {
        "name": "vgpaas",
        "selectorNames": [
          "cceUse"
        ],
        "cceManaged": true,
        "virtualSpaces": [
          {
            "name": "runtime",
            "size": "80%"
          },
          {
            "name": "kubernetes",
            "size": "20%"
          }
        ]
      }
    ],
    {
      "name": "vguser1",
      "selectorNames": [

```


```

        "user1"
      ],
      "virtualSpaces": [
        {
          "name": "user",
          "size": "100%",
          "lvmConfig": {
            "lvType": "linear",
            "path": "/tmp2"
          }
        }
      ]
    }
  ],
  "count": 1
}

```

### Creating Striped LVs to Improve Disk Performance

Currently, the striped LV function is supported only by calling an API. The following is an example:



1. **storageSelectors** matches all EVS disks in **dataVolumes** because **matchLabels** is not contained in **storageSelectors**.
2. Create a VG named **vgpaas** using LVM.
3. Strip 90% of the **vgpaas** space into runtime LVs.
4. Strip 10% of the **vgpaas** space into Kubernetes LVs.

 NOTE

- Two or more data disks are required for striping.
- When creating a striped LV, ensure that the types and sizes of the PVs added to the VG are the same. Otherwise, the creation will fail.
- When creating a striped LV, use the striping configuration for both the runtime LV and Kubernetes LV. Otherwise, the creation will fail.

Log in to the node and run the following command to view the striping result:

```
[root@test-83773 ~]# lvdisplay -m | grep -C 10 striped
LV Size                36.00 GiB
Current LE             9216
Segments               1
Allocation             inherit
Read ahead sectors    auto
- currently set to    512
Block device          253:0

--- Segments ---
Logical extents 0 to 9215:
  Type                striped
  Stripes              2
  Stripe size         64.00 KiB
  Stripe 0:
    Physical volume   /dev/sdb
    Physical extents  0 to 4607
  Stripe 1:
    Physical volume   /dev/sdc
    Physical extents  0 to 4607

--
LV Size                20.00 GiB
Current LE             5120
Segments               1
Allocation             inherit
Read ahead sectors    auto
- currently set to    8192
Block device          253:4

--- Segments ---
Logical extents 0 to 5119:
  Type                striped
  Stripes              2
  Stripe size         64.00 KiB
  Stripe 0:
    Physical volume   /dev/sdb
```

The following is an API example:

```
{
  "kind": "Node",
  "apiVersion": "v3",
```

```

"metadata": {
  "name": "test-83773"
},
"spec": {
  "flavor": "c3.large.2",
  "az": "",
  "os": "EulerOS 2.9",
  "dataVolumes": [
    {
      "size": 100,
      "volumetype": "SAS"
    },
    {
      "size": 100,
      "volumetype": "SAS"
    }
  ],
  "billingMode": 0,
  "extendParam": {
    "maxPods": 110
  },
  "nodeNicSpec": {
    "primaryNic": {
      "subnetId": "ca964acf-8468-4735-8229-97940ef6c881"
    }
  },
  "rootVolume": {
    "size": 50,
    "volumetype": "SAS"
  },
  "runtime": {
    "name": "docker"
  },
  "login": {
    "userPassword": {
      "username": "root",
      "password": "*****"
    }
  },
  "storage": {
    "storageSelectors": [
      {
        "name": "cceUse",
        "storageType": "evs"
      }
    ],
    "storageGroups": [
      {
        "name": "vgpaas",
        "selectorNames": [
          "cceUse"
        ],
        "cceManaged": true,
        "virtualSpaces": [
          {
            "name": "runtime",
            "size": "90%",
            "runtimeConfig": {
              "lvType": "striped"
            }
          },
          {
            "name": "kubernetes",
            "size": "10%",
            "lvmConfig": {
              "lvType": "striped"
            }
          }
        ]
      }
    ]
  }
}

```

```
    }  
  ]  
},  
"count": 1  
}
```