



Cloud Container Engine

API Reference

Date 2023-03-31

Contents

1 Before You Start.....	1
1.1 Overview.....	1
1.2 API Calling.....	2
1.3 Endpoints.....	2
1.4 Constraints.....	3
1.5 Concepts.....	3
2 API Overview.....	5
3 Calling APIs.....	23
3.1 Making an API Request.....	23
3.2 Authentication.....	28
3.3 Response.....	29
4 APIs.....	31
4.1 API URL.....	31
4.2 Cluster Management.....	32
4.2.1 Creating a Cluster.....	32
4.2.2 Reading a Specified Cluster.....	69
4.2.3 Listing Clusters in a Specified Project.....	85
4.2.4 Updating a Specified Cluster.....	103
4.2.5 Deleting a Cluster.....	122
4.2.6 Hibernating a Cluster.....	142
4.2.7 Waking Up a Cluster.....	143
4.2.8 Obtaining a Cluster Certificate.....	145
4.2.9 Querying a Job.....	149
4.2.10 Binding/Unbinding Public API Server Address.....	154
4.2.11 Obtaining Cluster Access Address.....	158
4.3 Node Management.....	162
4.3.1 Creating a Node.....	162
4.3.2 Reading a Specified Node.....	202
4.3.3 Listing All Nodes in a Cluster.....	219
4.3.4 Updating a Specified Node.....	237
4.3.5 Deleting a Node.....	255
4.3.6 Accepting a Node.....	273

4.3.7 Resetting a Node.....	288
4.3.8 Removing a Node.....	304
4.3.9 Migrating a Node.....	309
4.4 Node Pool Management.....	319
4.4.1 Creating a Node Pool.....	319
4.4.2 Reading a Specified Node Pool.....	364
4.4.3 Listing All Node Pools in a Specified Cluster.....	384
4.4.4 Updating a Specified Node Pool.....	405
4.4.5 Deleting a Node Pool.....	431
4.5 Storage Management.....	450
4.5.1 Creating a PVC (to be discarded).....	451
4.5.2 Deleting a PVC (to be discarded).....	459
4.6 Add-on Management.....	463
4.6.1 Installing an Add-on Instance.....	463
4.6.2 Listing Add-on Templates.....	472
4.6.3 Updating an Add-on Instance.....	478
4.6.4 Deleting an Add-on Instance.....	485
4.6.5 Querying an Add-on Instance.....	487
4.6.6 Listing Add-on Instances.....	491
4.7 Quota Management.....	496
4.7.1 Querying Resource Quotas.....	497
4.8 API Versions.....	498
4.8.1 Listing API Versions.....	499
5 Kubernetes APIs.....	502
6 Permissions Policies and Supported Actions.....	506
7 Appendix.....	513
7.1 Status Code.....	513
7.2 Error Codes.....	517
7.3 Obtaining a Project ID.....	522
7.4 Obtaining the Account ID.....	523
7.5 Specifying Add-ons to Be Installed During Cluster Creation.....	523
7.6 How to Obtain Parameters in the API URI.....	532
7.7 Creating a VPC and Subnet.....	533
7.8 Creating a Key Pair.....	533
7.9 Node Flavor Description.....	534
7.10 Adding a Salt in the password Field When Creating a Node.....	534
7.11 Maximum Number of Pods That Can Be Created on a Node.....	536
7.12 Node OS.....	538
7.13 Data Disk Space Allocation.....	539
7.14 Attaching Disks to a Node.....	542

8 Change History..... 552

1 Before You Start

1.1 Overview

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 to perform 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.

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

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

In addition:

- 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.

NOTICE

Use the methods described in this document to call APIs. If other methods, such as kubectl and SDKs, are used to call APIs, the APIs may fail to work.

This document describes functions, syntax, parameters, and examples of CCE APIs. You can find the information you need according to [Table 1-1](#).

Table 1-1 Document organization

Section	Description
API overview	General introduction to CCE APIs
Environment preparation	Preparations before using APIs, for example, creating a VPC and obtaining request authentication information
Using APIs	REST API messages, calling methods, and examples
API	CCE APIs to serve specific purposes, such as cluster management, storage management, and secret management
Common parameters	Common parameters, status codes, and error codes of CCE APIs

1.2 API Calling

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

1.3 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](#).

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 and storage 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.

Table 1-2 URL parameters

Parameter	Description
{clusterid}	Cluster ID. After a cluster is created, call the API for obtaining a cluster in a specified project to obtain the cluster ID.
Endpoint	Entry (URL) for a web service. Endpoints vary depending on services and regions.

Parameter	Description
uri	Access path of an API for performing a specified operation. Obtain the path from the URI of an API. For example, the resource-path of the API used to obtain a user token is v3/auth/tokens .

1.4 Constraints

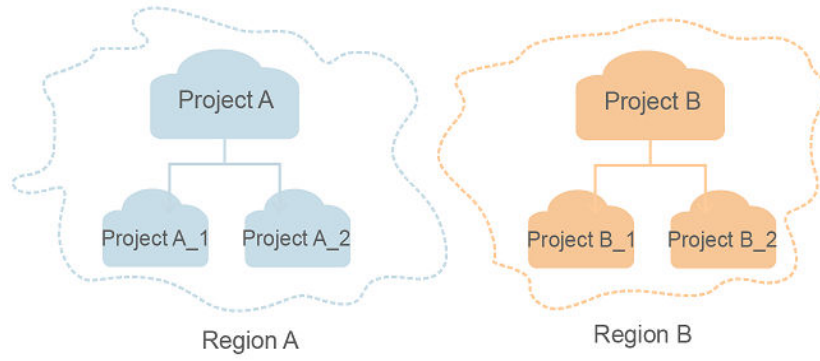
- CCE imposes a quota on the number and capacity of resources that a user can access. By default, you can create a maximum of five clusters in each region and a cluster can have a maximum of 50 nodes.
- For more constraints, see API description.

1.5 Concepts

- Account
An account is created upon successful registration. The account 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 account is a payment entity, which should not be used directly to perform routine management. For security purposes, create Identity and Access Management (IAM) users and grant them permissions for routine management.
- User
An IAM user is created by an account in IAM to use cloud services. Each IAM user has its 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



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	Cluster-related APIs	Manage clusters, including creating and deleting clusters. You can use APIs in this category to create clusters and obtain information about created clusters.
	Node-related APIs	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.
	Node Pool-related APIs	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.
	Add-on-related APIs	Manage add-ons, including querying AddonTemplates and creating, updating, deleting, and obtaining AddonInstances.
	Quota-related API	Query quotas of CCE resources.
Kubernetes-native APIs	-	Kubernetes-native APIs.

 NOTE

The Kubernetes-native APIs called in the current version do not support HTTP persistent connections.

Cluster-related APIs

Table 2-1 Cluster-related APIs

API	Description
Creating a Cluster	Create an empty cluster, which has only master nodes but do not have worker nodes.
Reading a Specified Cluster	Obtain details about a specified cluster.
Listing Clusters in a Specified Project	Obtain details about all clusters in a specified project.
Updating a Specified Cluster	Update information about a specified cluster.
Deleting a Cluster	Delete a specified cluster.
Hibernating a Cluster	Hibernate a specified cluster.
Waking Up a Cluster	Wake up a hibernated cluster.
Obtaining Cluster Certificates	Obtain certificates of a specified cluster.
Obtaining Job Information	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
Creating a Node	Create a node in a specified cluster.
Reading a Specified Node	Obtain details about a node with a specified node ID.
Reading all Nodes in a Cluster	Obtain details about all nodes in a cluster with a specified cluster ID.
Updating a Specified Node	Update information about a specified node.

API	Description
Deleting a Node	Delete a specified node
Accepting a Node	Accept a node into a specified cluster.
Resetting a Node	Reset a node in a specified cluster.
Removing a Node	Remove a node from a specified cluster.
Migrating a Node	Migrate a node from a specified cluster to another cluster.

Node Pool-related APIs

Table 2-3 Node pool-related APIs

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

Add-on-related APIs

Table 2-4 Add-on-related APIs

API	Description
Installing an Add-on Instance	Install an add-on by using the add-on template. One or more instances will be created for the installed add-on.
Listing Add-on Templates	Query add-on information.
Updating an Add-on Instance	Update an add-on instance.
Deleting an Add-on Instance	Delete an add-on instance.
Reading an Add-on Instance	Obtain details about an add-on instance.

API	Description
Listing Add-on Instances	List all add-on instances in the cluster.

Quota-related API

Table 2-5 Quota-related API

API	Description
Querying Resource Quotas	Query resource quotas.

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

API	Function	URI
	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}
	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}

API	Function	URI
	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
	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

API	Function	URI
	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
Daemon Set	Creating a DaemonSet	POST /apis/apps/v1/namespaces/{namespace}/daemonsets
	Deleting a specified DaemonSet	DELETE /apis/apps/v1/namespaces/{namespace}/daemonsets/{name}
	Deleting all DaemonSets	DELETE /apis/apps/v1/namespaces/{namespace}/daemonsets
	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
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

API	Function	URI
	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}
CronJob	Creating a cron job	POST /apis/batch/v1beta1/namespaces/{namespace}/cronjobs
	Deleting a cron job	DELETE /apis/batch/v1beta1/namespaces/{namespace}/cronjobs/{name}
	Deleting all cron jobs	DELETE /apis/batch/v1beta1/namespaces/{namespace}/cronjobs
	Reading a specified cron job	GET /apis/batch/v1beta1/namespaces/{namespace}/cronjobs/{name}
	Reading the status of a specified cron job	GET /apis/batch/v1beta1/namespaces/{namespace}/cronjobs/{name}/status
	Replacing a specified cron job	PUT /apis/batch/v1beta1/namespaces/{namespace}/cronjobs/{name}
	Replacing the status of a specified cron job	PUT /apis/batch/v1beta1/namespaces/{namespace}/cronjobs/{name}/status
	Listing cron jobs under a specified namespace	GET /apis/batch/v1beta1/namespaces/{namespace}/cronjobs
	Listing all cron jobs	GET /apis/batch/v1beta1/cronjobs
	Updating the status of a specified cron job	PATCH /apis/batch/v1beta1/namespaces/{namespace}/cronjobs/{name}/status
	Updating a specified cron job	PATCH /apis/batch/v1beta1/namespaces/{namespace}/cronjobs/{name}
	ReplicaSet	Listing ReplicaSets
Reading a specified ReplicaSet		GET /apis/apps/v1/namespaces/{namespace}/replicasets/{name}
Listing all ReplicaSets		GET /apis/apps/v1/replicasets

API	Function	URI
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}
	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
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

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

API	Function	URI
	Deleting an ingress	DELETE /apis/networking.k8s.io/v1/namespaces/{namespace}/ingresses/{name} (for clusters of v1.21 and later) DELETE /apis/networking.k8s.io/v1beta1/namespaces/{namespace}/ingresses/{name} (for clusters from v1.15 to v1.21) DELETE /apis/extensions/v1beta1/namespaces/{namespace}/ingresses/{name} (for clusters earlier than v1.15)
	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)

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

API	Function	URI
	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
	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 PersistentVolumeClaims	DELETE /api/v1/namespaces/{namespace}/persistentvolumeclaims
	Reading a specified PersistentVolumeClaim	GET /api/v1/namespaces/{namespace}/persistentvolumeclaims/{name}

API	Function	URI
	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 PersistentVolumeClaims in a specified namespace	GET /api/v1/namespaces/{namespace}/persistentvolumeclaims
	Listing all PersistentVolumeClaims	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}

API	Function	URI
	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}
	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}

API	Function	URI
	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}
	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/ RoleBin ding	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

API	Function	URI
	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
	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

API	Function	URI
	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
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 parameter description


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 Regions and Endpoints . For example, the endpoint of IAM in the UAE-Abu Dhabi region is iam.ae-ad-1.myhuaweicloud.com .
resource-path	Access path of an API for performing a specified operation. Obtain the path from the URI of an API. For example, the resource-path of the API used to obtain a user token is /v3/auth/tokens .

Parameter	Description
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, ?limit=10 indicates that a maximum of 10 data records will be displayed.

For example, to obtain an IAM token in the **ae-ad-1** region, obtain the endpoint of IAM (**iam.ae-ad-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.ae-ad-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	Same as GET except that the server must return only the response header.

Method	Description
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.ae-ad-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 https is 443 .	No This field is mandatory for AK/SK authentication.	code.test.com or code.test.com:443

Parameter	Description	Mandatory	Example Value
Content-Type	Specifies the type (or format) of the message body. The default value application/json 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 Obtaining a Project ID .	No This field is mandatory for requests that use AK/SK authentication in the Dedicated Cloud (DeC) scenario or multi-project scenario.	e9993fc787d94b6c886cb aa340f9c0f4
X-Auth-Token	Specifies the user token. It is a response to the API for obtaining a user token (This is the only API that does not require authentication). After the request is processed, the value of X-Subject-Token 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.ae-ad-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.ae-ad-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": "xxxxxxxxxxxxxxxxxxxx"
      }
    }
  }
}
```

```
}  
}
```

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.ae-ad-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-iam-trace-id → 218d45ab-d674-4995-af3a-2d0255ba41b5
x-subject-token
.....
x-ssr-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": "ae-ad-1",
            .....

```

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 and storage 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 API for obtaining a cluster in a specific project to obtain the cluster ID.
Endpoint	URL that is the entry point for a web service. You can obtain it from Endpoints .
uri	Access path of an API for performing a specified operation. Obtain the path from the URI of an API. For example, the resource-path of the API used to obtain a user token is v3/auth/tokens .

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.
- CIDR blocks have been properly set for containers and Services. They are not editable after cluster creation, unless you create a new cluster.
- 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. 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 a maximum of 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 How to Obtain Parameters in the API URI .

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 Obtaining a User Token .

Table 4-4 Request body parameters

Parameter	Mandatory	Type	Description
kind	Yes	String	API type. The value is fixed at Cluster or cluster and cannot be changed.
apiVersion	Yes	String	API version. The value is fixed at v3 and cannot be changed.
metadata	Yes	ClusterMetadata object	Basic information about a cluster. Metadata is a collection of attributes.
spec	Yes	ClusterSpec object	Detailed description of the cluster. CCE creates or updates objects by defining or updating spec .
status	No	ClusterStatus 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 resource. The value is automatically generated after the object is created. A user-defined value will not take effect.
annotations	No	Map<String,String>	Cluster annotations, in the format of key-value pairs. <pre>"annotations": { "key1": "value1", "key2": "value2" }</pre> <p>NOTE</p> <ul style="list-style-type: none"> • annotations: Does not label or select objects. The metadata in annotations may be small or large, structured or unstructured, and may include characters that are not allowed in labels. • This field is not stored in the database and is used only to specify the add-ons to be installed in the cluster. • Install ICAgent during cluster creation by adding the key-value pair "cluster.install.addons.external / install":{"addonTemplateName":"icagent"}".
labels	No	Map<String,String>	Cluster labels, in the format of key-value pairs. <p>NOTE</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>

Parameter	Mandatory	Type	Description
creationTimes tamp	No	String	Time when the cluster was created.
updateTimest amp	No	String	Time when the cluster was updated.

Table 4-6 ClusterSpec

Parameter	Mandatory	Type	Description
category	No	String	Cluster type: <ul style="list-style-type: none"> • CCE: CCE cluster. CCE cluster supports hybrid deployment of VMs and bare-metal servers (BMSs), and heterogeneous nodes such as GPU- and NPU-enabled nodes. You can run your containers in a secure and stable container runtime environment based on a high-performance network model.
type	No	String	Master node architecture: <ul style="list-style-type: none"> • VirtualMachine: x86

Parameter	Mandatory	Type	Description
flavor	Yes	String	<p>Default value: When you create a CCE cluster, the value is cce.s1.small for non-DeC scenarios and cce.dec.s1.small for DeC scenarios.</p> <p>Cluster flavor, which cannot be changed after the cluster is created.</p> <ul style="list-style-type: none"> ● cce.s1.small: small-scale, single-master CCE cluster (≤ 50 nodes) ● cce.s1.medium: medium-scale, single-master CCE cluster (≤ 200 nodes) ● cce.s2.small: small-scale, multi-master CCE cluster (≤ 50 nodes) ● cce.s2.medium: medium-scale, multi-master CCE cluster (≤ 200 nodes) ● cce.s2.large: large-scale, multi-master CCE cluster (≤ 1,000 nodes) ● cce.s2.xlarge: ultra-large-scale, multi-master CCE cluster (≤ 2,000 nodes)

Parameter	Mandatory	Type	Description
			<p>NOTE</p> <ul style="list-style-type: none"> • s1: single-master CCE cluster • s2: multi-master CCE cluster • dec: dedicated CCE cluster. For example, cce.dec.s1.small is a small-scale, single-master, dedicated CCE cluster (≤ 50 nodes). • Values in the parentheses above indicate the maximum number of nodes that can be managed by the cluster. • A single-master cluster has only one master node. If the master node is down, the cluster will become unavailable and stop serving new workloads. However, existing workloads in the cluster are not affected. • A multi-master cluster is highly available. When a master node is faulty, the cluster is still available.
version	No	String	<p>Cluster version, which mirrors the baseline version of the Kubernetes community. The latest version is recommended. You can create clusters of two latest versions on the CCE console. To learn which cluster versions are available, log in to the CCE console, create a cluster, and check the Version parameter. You can call APIs to create clusters of other versions. However, these clusters will be gradually brought offline. For details about the support policy, see the CCE announcement.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If not specified, a cluster of the latest version is created. • If the baseline cluster version is specified but the R version is not specified, the system selects the latest R version of the cluster by default. You are advised not to specify the R version.

Parameter	Mandatory	Type	Description
platformVersion	No	String	<p>CCE cluster platform version, indicating the internal version under the cluster version (version). 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 platformVersion is cce.X.Y.</p> <ul style="list-style-type: none"> • X: internal feature version, indicating changes in features, patches, or OS support in the cluster version. The value starts from 1 and increases monotonically. • Y: 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 0 and increases monotonically.
description	No	String	<p>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 API for updating cluster information or go to the cluster details page on the CCE console. Only UTF-8 encoding is supported.</p>

Parameter	Mandatory	Type	Description
customSan	No	Array of strings	<p>Custom SAN field in the server certificate of the cluster API server, which must comply with the SSL and X509 format specifications.</p> <ol style="list-style-type: none"> 1. Duplicate names are not allowed. 2. Must comply with the IP address and domain name formats. <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>
ipv6enable	No	Boolean	Whether the cluster supports IPv6 addresses. This field is supported in clusters of v1.15 and later versions.
hostNetwork	Yes	HostNetwork 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	ContainerNetwork object	Container networking parameters, including the container network model and container CIDR block.
authentication	No	Authentication object	Configurations of the cluster authentication mode.
billingMode	No	Integer	<p>Billing mode of a cluster.</p> <ul style="list-style-type: none"> • 0: pay-per-use <p>Defaults to pay-per-use.</p>
masters	No	Array of MasterSpec objects	Advanced configurations of master nodes
kubernetesSvcIpRange	No	String	Service CIDR block or the IP address range which kubernetes clusterIp must fall within. This field is available only for clusters of v1.11.7 and later.

Parameter	Mandatory	Type	Description
clusterTags	No	Array of ResourceTag objects	Cluster resource tags.
kubeProxyMode	No	String	Service forwarding mode. Two modes are available: <ul style="list-style-type: none"> • iptables: 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. • ipvs: 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.
az	No	String	AZ. This field is returned only for a query.
extendParam	No	ClusterExtendParam 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.

Parameter	Mandatory	Type	Description
configurations Override	No	Array of PackageConfiguration objects	Whether to overwrite the default component configurations of the cluster. If a component or a parameter that is not supported by the component is specified, this configuration is ignored. For details about the supported configurable components and their parameters, see

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"> • Method 1: Log in to the VPC console and view the VPC ID in the VPC details. • Method 2: Query the VPC ID through the VPC API. <p>NOTE Currently, the VPC network model does not support interconnection with VPCs that contain a secondary CIDR block.</p>
subnet	Yes	String	<p>Network ID of the subnet used to create a master node.</p> <p>Methods:</p> <ul style="list-style-type: none"> • Method 1: Log in to VPC console and click the target subnet on the Subnets tab page. You can view the network ID on the displayed page. • Method 2: Use the VPC API for querying subnets. . .

Parameter	Mandatory	Type	Description
SecurityGroup	No	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.

Table 4-8 ContainerNetwork

Parameter	Mandatory	Type	Description
mode	Yes	String	Container network model. Select one of the following possible values: <ul style="list-style-type: none"> • overlay_l2: an overlay_l2 network (container tunnel network) built for containers by using OpenVSwitch (OVS). • vpc-router: an underlay_l2 network built for containers by using IPVLAN and custom VPC routes.
cidr	No	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. Not editable after the cluster is created. (deprecated. A specified cidrs will make cidr invalid.)

Parameter	Mandatory	Type	Description
cidrs	No	Array of ContainerCIDR objects	List of container CIDR blocks. In clusters of v1.21 and later, the cidrs field is used. When the cluster network type is vpc-router , you can add multiple container CIDR blocks. In versions earlier than v1.21, if the cidrs field is used, the first CIDR element in the array is used as the container CIDR block. This parameter cannot be modified after the cluster is created.

Table 4-9 ContainerCIDR

Parameter	Mandatory	Type	Description
cidr	Yes	String	Container CIDR block. Recommended: 10.0.0.0/12-19, 172.16.0.0/16-19, and 192.168.0.0/16-19

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"> Method 1: Log in to the VPC console and click the target subnet on the Subnets page. You can view the IPv4 subnet ID on the displayed page. Method 2: Use the VPC API for querying subnets.
eniSubnetCIDR	Yes	String	ENI subnet CIDR (being discarded)

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

Table 4-12 Authentication

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

Table 4-13 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 authenticating_proxy . 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 authenticating_proxy .
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 authenticating_proxy .

Table 4-14 MasterSpec

Parameter	Mandatory	Type	Description
availabilityZone	No	String	AZ

Table 4-15 ResourceTag

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

Table 4-16 ClusterExtendParam

Parameter	Mandatory	Type	Description
clusterAZ	No	String	<p>AZ of master nodes in the cluster.</p> <ul style="list-style-type: none"> • multi_az: (Optional) The cluster will span across AZs. This field is configurable only for high-availability clusters. • <i>AZ of the dedicated cloud computing pool:</i> The cluster will be deployed in the AZ of Dedicated Cloud (DeC). This parameter is mandatory for dedicated CCE clusters.

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><rootVol.dssPoolID>.<rootVol.volType>;<dataVol.dssPoolID>.<dataVol.volType></code></p> <p>Field description:</p> <ul style="list-style-type: none"> • rootVol is the system disk. dataVol is the data disk. • dssPoolID indicates the ID of the DSS storage pool. • volType indicates the storage volume type of the DSS storage pool, such as SAS and SSD. <p>Example: c950ee97-587c-4f24-8a74-3367e3da570f.sas; 6edbc2f4-1507-44f8-ac0d-eed1d2608d38.ssd</p> <p>NOTE 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>NOTE</p> <ul style="list-style-type: none"> • An enterprise project can be configured only after the enterprise project function is enabled. • The enterprise project to which the cluster belongs must be the same as that to which other cloud service resources associated with the cluster belong.

Parameter	Mandatory	Type	Description
kubeProxyMode	No	String	<p>Service forwarding mode. Two modes are available:</p> <ul style="list-style-type: none"> • iptables: 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. • ipvs: 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. <p>NOTE This parameter has been deprecated. If this parameter and kubeProxyMode in ClusterSpec 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	<p>Number of mask bits of the fixed IP address pool of the container network model. This field is supported only for the VPC network model (vpc-router).</p> <p>This parameter determines the number of container IP addresses that can be allocated to a node. The maximum number of pods that can be created on a node is decided by this parameter and maxPods set during node creation. For details, see Maximum Number of Pods That Can Be Created on a Node.</p> <p>For integer characters, the value ranges from 24 to 28.</p>
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 secure or normal . If this parameter is not specified, normal is used by default.
kubernetes.io/ cpuManagerPolicy	No	String	<p>Cluster CPU management policy. The value can be none or static. Defaults to none.</p> <ul style="list-style-type: none"> • none: CPU cores will not be exclusively allocated to workload pods. Select this value if you want a large pool of shareable CPU cores. • static: CPU cores can be exclusively allocated to workload pods. Select this value if your workload is sensitive to latency in CPU cache and scheduling.

Parameter	Mandatory	Type	Description
upgradefrom	No	String	Records of how the cluster is upgraded to the current version

Table 4-17 PackageConfiguration

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

Table 4-18 ConfigurationItem

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

Table 4-19 ClusterStatus

Parameter	Mandatory	Type	Description
phase	No	String	Cluster status. Possible values: <ul style="list-style-type: none"> • Available: The cluster is running properly. • Unavailable: The cluster is exhibiting unexpected behavior. Manually delete the cluster. • ScalingUp: Nodes are being added to the cluster. • ScalingDown: The cluster is being downsized to fewer nodes. • Creating: The cluster is being created. • Deleting: The cluster is being deleted. • Upgrading: The cluster is being upgraded. • Resizing: Cluster specifications are being changed. • RollingBack: The cluster is being rolled back. • RollbackFailed: The rollback is abnormal. • Empty: The cluster has no resources.
jobID	No	String	Job ID.
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 ClusterEndpoints objects	Access address of kube-apiserver in the cluster.

Parameter	Mandatory	Type	Description
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 4-20 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"> • Internal: address for internal network access • External: address for external network access

Response Parameters

Status code: 201

Table 4-21 Response body parameters

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

Table 4-22 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 resource. The value is automatically generated after the object is created. A user-defined value will not take effect.
annotations	Map<String,String>	Cluster annotations, in the format of key-value pairs. <pre>"annotations": { "key1" : "value1", "key2" : "value2" }</pre> <p>NOTE</p> <ul style="list-style-type: none"> • annotations: Does not label or select objects. The metadata in annotations may be small or large, structured or unstructured, and may include characters that are not allowed in labels. • This field is not stored in the database and is used only to specify the add-ons to be installed in the cluster. • Install ICAGENT during cluster creation by adding the key-value pair "cluster.install.addons.external/install":{"addonTemplateName":"icagent"}".
labels	Map<String,String>	Cluster labels, in the format of key-value pairs. <p>NOTE</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 4-23 ClusterSpec

Parameter	Type	Description
category	String	Cluster type: <ul style="list-style-type: none"> ● CCE: CCE cluster. CCE cluster supports hybrid deployment of VMs and bare-metal servers (BMSs), and heterogeneous nodes such as GPU- and NPU-enabled nodes. You can run your containers in a secure and stable container runtime environment based on a high-performance network model.
type	String	Master node architecture: <ul style="list-style-type: none"> ● VirtualMachine: x86
flavor	String	Default value: When you create a CCE cluster, the value is cce.s1.small for non-DeC scenarios and cce.dec.s1.small for DeC scenarios. Cluster flavor, which cannot be changed after the cluster is created. <ul style="list-style-type: none"> ● cce.s1.small: small-scale, single-master CCE cluster (≤ 50 nodes) ● cce.s1.medium: medium-scale, single-master CCE cluster (≤ 200 nodes) ● cce.s2.small: small-scale, multi-master CCE cluster (≤ 50 nodes) ● cce.s2.medium: medium-scale, multi-master CCE cluster (≤ 200 nodes) ● cce.s2.large: large-scale, multi-master CCE cluster (≤ 1,000 nodes) ● cce.s2.xlarge: ultra-large-scale, multi-master CCE cluster (≤ 2,000 nodes) <p>NOTE</p> <ul style="list-style-type: none"> ● s1: single-master CCE cluster ● s2: multi-master CCE cluster ● dec: dedicated CCE cluster. For example, cce.dec.s1.small is a small-scale, single-master, dedicated CCE cluster (≤ 50 nodes). ● Values in the parentheses above indicate the maximum number of nodes that can be managed by the cluster. ● A single-master cluster has only one master node. If the master node is down, the cluster will become unavailable and stop serving new workloads. However, existing workloads in the cluster are not affected. ● A multi-master cluster is highly available. When a master node is faulty, the cluster is still available.

Parameter	Type	Description
version	String	<p>Cluster version, which mirrors the baseline version of the Kubernetes community. The latest version is recommended.</p> <p>You can create clusters of two latest versions on the CCE console. To learn which cluster versions are available, log in to the CCE console, create a cluster, and check the Version parameter. You can call APIs to create clusters of other versions. However, these clusters will be gradually brought offline. For details about the support policy, see the CCE announcement.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If not specified, a cluster of the latest version is created. • If the baseline cluster version is specified but the R version is not specified, the system selects the latest R version of the cluster by default. You are advised not to specify the R version.
platformVersion	String	<p>CCE cluster platform version, indicating the internal version under the cluster version (version). 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 platformVersion is cce.X.Y.</p> <ul style="list-style-type: none"> • X: internal feature version, indicating changes in features, patches, or OS support in the cluster version. The value starts from 1 and increases monotonically. • Y: 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 0 and increases monotonically.
description	String	<p>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 API for updating cluster information or go to the cluster details page on the CCE console. Only UTF-8 encoding is supported.</p>

Parameter	Type	Description
customSan	Array of strings	<p>Custom SAN field in the server certificate of the cluster API server, which must comply with the SSL and X509 format specifications.</p> <ol style="list-style-type: none"> 1. Duplicate names are not allowed. 2. Must comply with the IP address and domain name formats. <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 the cluster supports IPv6 addresses. This field is supported in clusters of v1.15 and later versions.
hostNetwork	HostNetwork 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	ContainerNetwork object	Container networking parameters, including the container network model and container CIDR block.
authentication	Authentication object	Configurations of the cluster authentication mode.
billingMode	Integer	<p>Billing mode of a cluster.</p> <ul style="list-style-type: none"> • 0: pay-per-use <p>Defaults to pay-per-use.</p>
masters	Array of MasterSpec objects	Advanced configurations of master nodes
kubernetesServiceIpRange	String	Service CIDR block or the IP address range which kubernetes clusterIp must fall within. This field is available only for clusters of v1.11.7 and later.
clusterTags	Array of ResourceTag objects	Cluster resource tags.

Parameter	Type	Description
kubeProxyMode	String	Service forwarding mode. Two modes are available: <ul style="list-style-type: none"> • iptables: 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. • ipvs: 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.
az	String	AZ. This field is returned only for a query.
extendParam	ClusterExtendedParam 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 PackageConfiguration objects	Whether to overwrite the default component configurations of the cluster. If a component or a parameter that is not supported by the component is specified, this configuration is ignored. For details about the supported configurable components and their parameters, see

Table 4-24 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"> • Method 1: Log in to the VPC console and view the VPC ID in the VPC details. • Method 2: Query the VPC ID through the VPC API. <p>NOTE Currently, the VPC network model does not support interconnection with VPCs that contain a secondary CIDR block.</p>

Parameter	Type	Description
subnet	String	Network ID of the subnet used to create a master node. Methods: <ul style="list-style-type: none"> • Method 1: Log in to VPC console and click the target subnet on the Subnets tab page. You can view the network ID on the displayed page. • Method 2: Use the VPC API for querying subnets. . .
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.

Table 4-25 ContainerNetwork

Parameter	Type	Description
mode	String	Container network model. Select one of the following possible values: <ul style="list-style-type: none"> • overlay_l2: an overlay_l2 network (container tunnel network) built for containers by using OpenVSwitch (OVS). • vpc-router: an underlay_l2 network built for containers by using IPVLAN and custom VPC routes.
cidr	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. Not editable after the cluster is created. (deprecated. A specified cidrs will make cidr invalid.)

Parameter	Type	Description
cidrs	Array of ContainerCIDR objects	List of container CIDR blocks. In clusters of v1.21 and later, the cidrs field is used. When the cluster network type is vpc-router , you can add multiple container CIDR blocks. In versions earlier than v1.21, if the cidrs field is used, the first CIDR element in the array is used as the container CIDR block. This parameter cannot be modified after the cluster is created.

Table 4-26 ContainerCIDR

Parameter	Type	Description
cidr	String	Container CIDR block. Recommended: 10.0.0.0/12-19, 172.16.0.0/16-19, and 192.168.0.0/16-19

Table 4-27 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"> Method 1: Log in to the VPC console and click the target subnet on the Subnets page. You can view the IPv4 subnet ID on the displayed page. Method 2: Use the VPC API for querying subnets.
eniSubnetCIDR	String	ENI subnet CIDR (being discarded)
subnets	Array of NetworkSubnet objects	List of IPv4 subnet IDs

Table 4-28 NetworkSubnet

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

Table 4-29 Authentication

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

Table 4-30 AuthenticatingProxy

Parameter	Type	Description
ca	String	X509 CA certificate (Base64-encoded) configured in <code>authenticating_proxy</code> mode. This field is mandatory when the cluster authentication mode is authenticating_proxy . Maximum size: 1 MB
cert	String	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 value must be Base64-encoded.) This field is mandatory when the cluster authentication mode is authenticating_proxy .

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 authenticating_proxy .

Table 4-31 MasterSpec

Parameter	Type	Description
availabilityZone	String	AZ

Table 4-32 ResourceTag

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

Table 4-33 ClusterExtendParam

Parameter	Type	Description
clusterAZ	String	<p>AZ of master nodes in the cluster.</p> <ul style="list-style-type: none"> • multi_az: (Optional) The cluster will span across AZs. This field is configurable only for high-availability clusters. • <i>AZ of the dedicated cloud computing pool</i>: The cluster will be deployed in the AZ of Dedicated Cloud (DeC). This parameter is mandatory for dedicated CCE clusters.
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><rootVol.dssPoolID>.<rootVol.volType>;<dataVol.dssPoolID>.<dataVol.volType></code></p> <p>Field description:</p> <ul style="list-style-type: none"> • rootVol is the system disk. dataVol is the data disk. • dssPoolID indicates the ID of the DSS storage pool. • volType indicates the storage volume type of the DSS storage pool, such as SAS and SSD. <p>Example: c950ee97-587c-4f24-8a74-3367e3da570f.sas; 6edbc2f4-1507-44f8-ac0d-eed1d2608d38.ssd</p> <p>NOTE 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>NOTE</p> <ul style="list-style-type: none"> • An enterprise project can be configured only after the enterprise project function is enabled. • The enterprise project to which the cluster belongs must be the same as that to which other cloud service resources associated with the cluster belong.

Parameter	Type	Description
kubeProxyMode	String	<p>Service forwarding mode. Two modes are available:</p> <ul style="list-style-type: none"> • iptables: 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. • ipvs: 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. <p>NOTE This parameter has been deprecated. If this parameter and kubeProxyMode in ClusterSpec 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 network model. This field is supported only for the VPC network model (vpc-router).</p> <p>This parameter determines the number of container IP addresses that can be allocated to a node. The maximum number of pods that can be created on a node is decided by this parameter and maxPods set during node creation. For details, see Maximum Number of Pods That Can Be Created on a Node.</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 secure or normal . If this parameter is not specified, normal is used by default.

Parameter	Type	Description
kubernetes.io/ cpuManagerP olicy	String	Cluster CPU management policy. The value can be none or static . Defaults to none . <ul style="list-style-type: none"> • none: CPU cores will not be exclusively allocated to workload pods. Select this value if you want a large pool of shareable CPU cores. • static: CPU cores can be exclusively allocated to workload pods. Select this value if your workload is sensitive to latency in CPU cache and scheduling.
upgradefrom	String	Records of how the cluster is upgraded to the current version

Table 4-34 PackageConfiguration

Parameter	Type	Description
name	String	Component name.
configurations	Array of Configuratio nItem objects	Component configuration item.

Table 4-35 ConfigurationItem

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

Table 4-36 ClusterStatus

Parameter	Type	Description
phase	String	Cluster status. Possible values: <ul style="list-style-type: none"> • Available: The cluster is running properly. • Unavailable: The cluster is exhibiting unexpected behavior. Manually delete the cluster. • ScalingUp: Nodes are being added to the cluster. • ScalingDown: The cluster is being downsized to fewer nodes. • Creating: The cluster is being created. • Deleting: The cluster is being deleted. • Upgrading: The cluster is being upgraded. • Resizing: Cluster specifications are being changed. • RollingBack: The cluster is being rolled back. • RollbackFailed: The rollback is abnormal. • Empty: The cluster has no resources.
jobID	String	Job ID.
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 ClusterEndpoints 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 4-37 ClusterEndpoints

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

Example Requests

- Creating a CCE 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"
    },
    "containerNetwork": {
      "mode": "vpc-router",
      "cidr": "10.0.0.0/16"
    },
    "kubernetesSvcIppRange": "10.247.0.0/16",
    "description": "",
    "billingMode": 0,
    "extendParam": {
      "kubeProxyMode": "iptables",
      "alpha.cce/fixPoolMask": "25",
      "enterpriseProjectId": "0"
    },
    "authentication": {
      "mode": "rbac"
    },
    "ipv6enable": false
  }
}

```

- Creating a CCE cluster with ICAgent installed

```

/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" : "v1.19",
"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
}

```

- **Creating a Cluster and Specifying a Default Node Security Group**

```

/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",
      "SecurityGroup" : "a4ef108c-2ec6-492f-a6c4-7b64e25ae490"
    },
    "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
  }
}

```

Example Responses

Status code: 201

The cluster creation job is successfully delivered.

```
{
  "kind": "Cluster",
  "apiVersion": "v3",
  "metadata": {
    "name": "cluster",
    "uid": "bce956e1-87f3-11ec-b5e5-0255ac101514",
    "creationTimestamp": "2022-02-07 08:55:45.785425492 +0000 UTC",
    "updateTimestamp": "2022-02-07 08:55:45.78542824 +0000 UTC",
    "annotations": {
      "jobid": "bd42f724-87f3-11ec-b5e5-0255ac101514",
      "resourceJobId": "bce99f78-87f3-11ec-b5e5-0255ac101514"
    }
  },
  "spec": {
    "category": "CCE",
    "type": "VirtualMachine",
    "flavor": "cce.s2.small",
    "version": "v1.19.10-r0",
    "platformVersion": "cce.5.0",
    "hostNetwork": {
      "vpc": "030bfb19-5fa7-42ad-8a0d-c0721d268867",
      "subnet": "ca964acf-8468-4735-8229-97940ef6c881"
    },
    "containerNetwork": {
      "mode": "vpc-router",
      "cidr": "10.0.0.0/16",
      "cidrs": [ {
        "cidr": "10.0.0.0/16"
      } ]
    },
    "eniNetwork": { },
    "authentication": {
      "mode": "rbac",
      "authenticatingProxy": { }
    },
    "billingMode": 0,
    "kubernetesSvcIpsRange": "10.247.0.0/16",
    "kubeProxyMode": "iptables",
    "extendParam": {
      "alpha.cce/fixPoolMask": "25",
      "enterpriseProjectId": "0",
      "kubeProxyMode": "iptables",
      "orderID": ""
    }
  },
  "status": {
    "phase": "Creating",
    "jobID": "bd42f724-87f3-11ec-b5e5-0255ac101514"
  }
}
```

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 4-38 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .

Table 4-39 Query Parameters

Parameter	Mandatory	Type	Description
detail	No	String	Whether the details about a cluster are queried. If this parameter is set to true , 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 4-40 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 Obtaining a User Token .

Response Parameters

Status code: 200

Table 4-41 Response body parameters

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

Table 4-42 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 resource. The value is automatically generated after the object is created. A user-defined value will not take effect.
annotations	Map<String,String>	Cluster annotations, in the format of key-value pairs. <pre>"annotations": { "key1": "value1", "key2": "value2" }</pre> <p>NOTE</p> <ul style="list-style-type: none"> • annotations: Does not label or select objects. The metadata in annotations may be small or large, structured or unstructured, and may include characters that are not allowed in labels. • This field is not stored in the database and is used only to specify the add-ons to be installed in the cluster. • Install ICAgent during cluster creation by adding the key-value pair "cluster.install.addons.external/install":{"addonTemplateName":"icagent"}".
labels	Map<String,String>	Cluster labels, in the format of key-value pairs. <p>NOTE</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 4-43 ClusterSpec

Parameter	Type	Description
category	String	Cluster type: <ul style="list-style-type: none"> • CCE: CCE cluster. CCE cluster supports hybrid deployment of VMs and bare-metal servers (BMSs), and heterogeneous nodes such as GPU- and NPU-enabled nodes. You can run your containers in a secure and stable container runtime environment based on a high-performance network model.
type	String	Master node architecture: <ul style="list-style-type: none"> • VirtualMachine: x86

Parameter	Type	Description
flavor	String	<p>Default value: When you create a CCE cluster, the value is cce.s1.small for non-DeC scenarios and cce.dec.s1.small for DeC scenarios.</p> <p>Cluster flavor, which cannot be changed after the cluster is created.</p> <ul style="list-style-type: none"> ● cce.s1.small: small-scale, single-master CCE cluster (≤ 50 nodes) ● cce.s1.medium: medium-scale, single-master CCE cluster (≤ 200 nodes) ● cce.s2.small: small-scale, multi-master CCE cluster (≤ 50 nodes) ● cce.s2.medium: medium-scale, multi-master CCE cluster (≤ 200 nodes) ● cce.s2.large: large-scale, multi-master CCE cluster (≤ 1,000 nodes) ● cce.s2.xlarge: ultra-large-scale, multi-master CCE cluster (≤ 2,000 nodes) <p>NOTE</p> <ul style="list-style-type: none"> ● s1: single-master CCE cluster ● s2: multi-master CCE cluster ● dec: dedicated CCE cluster. For example, cce.dec.s1.small is a small-scale, single-master, dedicated CCE cluster (≤ 50 nodes). ● Values in the parentheses above indicate the maximum number of nodes that can be managed by the cluster. ● A single-master cluster has only one master node. If the master node is down, the cluster will become unavailable and stop serving new workloads. However, existing workloads in the cluster are not affected. ● A multi-master cluster is highly available. When a master node is faulty, the cluster is still available.

Parameter	Type	Description
version	String	<p>Cluster version, which mirrors the baseline version of the Kubernetes community. The latest version is recommended.</p> <p>You can create clusters of two latest versions on the CCE console. To learn which cluster versions are available, log in to the CCE console, create a cluster, and check the Version parameter. You can call APIs to create clusters of other versions. However, these clusters will be gradually brought offline. For details about the support policy, see the CCE announcement.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If not specified, a cluster of the latest version is created. • If the baseline cluster version is specified but the R version is not specified, the system selects the latest R version of the cluster by default. You are advised not to specify the R version.
platformVersion	String	<p>CCE cluster platform version, indicating the internal version under the cluster version (version). 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 platformVersion is cce.X.Y.</p> <ul style="list-style-type: none"> • X: internal feature version, indicating changes in features, patches, or OS support in the cluster version. The value starts from 1 and increases monotonically. • Y: 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 0 and increases monotonically.
description	String	<p>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 API for updating cluster information or go to the cluster details page on the CCE console. Only UTF-8 encoding is supported.</p>

Parameter	Type	Description
customSan	Array of strings	<p>Custom SAN field in the server certificate of the cluster API server, which must comply with the SSL and X509 format specifications.</p> <ol style="list-style-type: none"> 1. Duplicate names are not allowed. 2. Must comply with the IP address and domain name formats. <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 the cluster supports IPv6 addresses. This field is supported in clusters of v1.15 and later versions.
hostNetwork	HostNetwork 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	ContainerNetwork object	Container networking parameters, including the container network model and container CIDR block.
authentication	Authentication object	Configurations of the cluster authentication mode.
billingMode	Integer	<p>Billing mode of a cluster.</p> <ul style="list-style-type: none"> • 0: pay-per-use <p>Defaults to pay-per-use.</p>
masters	Array of MasterSpec objects	Advanced configurations of master nodes
kubernetesServiceIpRange	String	Service CIDR block or the IP address range which kubernetes clusterIp must fall within. This field is available only for clusters of v1.11.7 and later.
clusterTags	Array of ResourceTag objects	Cluster resource tags.

Parameter	Type	Description
kubeProxyMode	String	Service forwarding mode. Two modes are available: <ul style="list-style-type: none"> • iptables: 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. • ipvs: 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.
az	String	AZ. This field is returned only for a query.
extendParam	ClusterExtendedParam 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 PackageConfiguration objects	Whether to overwrite the default component configurations of the cluster. If a component or a parameter that is not supported by the component is specified, this configuration is ignored. For details about the supported configurable components and their parameters, see

Table 4-44 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"> • Method 1: Log in to the VPC console and view the VPC ID in the VPC details. • Method 2: Query the VPC ID through the VPC API. <p>NOTE Currently, the VPC network model does not support interconnection with VPCs that contain a secondary CIDR block.</p>

Parameter	Type	Description
subnet	String	Network ID of the subnet used to create a master node. Methods: <ul style="list-style-type: none"> • Method 1: Log in to VPC console and click the target subnet on the Subnets tab page. You can view the network ID on the displayed page. • Method 2: Use the VPC API for querying subnets. . .
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.

Table 4-45 ContainerNetwork

Parameter	Type	Description
mode	String	Container network model. Select one of the following possible values: <ul style="list-style-type: none"> • overlay_l2: an overlay_l2 network (container tunnel network) built for containers by using OpenVSwitch (OVS). • vpc-router: an underlay_l2 network built for containers by using IPVLAN and custom VPC routes.
cidr	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. Not editable after the cluster is created. (deprecated. A specified cidrs will make cidr invalid.)

Parameter	Type	Description
cidrs	Array of ContainerCIDR objects	List of container CIDR blocks. In clusters of v1.21 and later, the cidrs field is used. When the cluster network type is vpc-router , you can add multiple container CIDR blocks. In versions earlier than v1.21, if the cidrs field is used, the first CIDR element in the array is used as the container CIDR block. This parameter cannot be modified after the cluster is created.

Table 4-46 ContainerCIDR

Parameter	Type	Description
cidr	String	Container CIDR block. Recommended: 10.0.0.0/12-19, 172.16.0.0/16-19, and 192.168.0.0/16-19

Table 4-47 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"> Method 1: Log in to the VPC console and click the target subnet on the Subnets page. You can view the IPv4 subnet ID on the displayed page. Method 2: Use the VPC API for querying subnets.
eniSubnetCIDR	String	ENI subnet CIDR (being discarded)
subnets	Array of NetworkSubnet objects	List of IPv4 subnet IDs

Table 4-48 NetworkSubnet

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

Table 4-49 Authentication

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

Table 4-50 AuthenticatingProxy

Parameter	Type	Description
ca	String	X509 CA certificate (Base64-encoded) configured in <code>authenticating_proxy</code> mode. This field is mandatory when the cluster authentication mode is authenticating_proxy . Maximum size: 1 MB
cert	String	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 value must be Base64-encoded.) This field is mandatory when the cluster authentication mode is authenticating_proxy .

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 authenticating_proxy .

Table 4-51 MasterSpec

Parameter	Type	Description
availabilityZone	String	AZ

Table 4-52 ResourceTag

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

Table 4-53 ClusterExtendParam

Parameter	Type	Description
clusterAZ	String	<p>AZ of master nodes in the cluster.</p> <ul style="list-style-type: none"> • multi_az: (Optional) The cluster will span across AZs. This field is configurable only for high-availability clusters. • <i>AZ of the dedicated cloud computing pool</i>: The cluster will be deployed in the AZ of Dedicated Cloud (DeC). This parameter is mandatory for dedicated CCE clusters.
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><rootVol.dssPoolID>.<rootVol.volType>;<dataVol.dssPoolID>.<dataVol.volType></code></p> <p>Field description:</p> <ul style="list-style-type: none"> • rootVol is the system disk. dataVol is the data disk. • dssPoolID indicates the ID of the DSS storage pool. • volType indicates the storage volume type of the DSS storage pool, such as SAS and SSD. <p>Example: c950ee97-587c-4f24-8a74-3367e3da570f.sas; 6edbc2f4-1507-44f8-ac0d-eed1d2608d38.ssd</p> <p>NOTE 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>NOTE</p> <ul style="list-style-type: none"> • An enterprise project can be configured only after the enterprise project function is enabled. • The enterprise project to which the cluster belongs must be the same as that to which other cloud service resources associated with the cluster belong.

Parameter	Type	Description
kubeProxyMode	String	<p>Service forwarding mode. Two modes are available:</p> <ul style="list-style-type: none"> • iptables: 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. • ipvs: 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. <p>NOTE This parameter has been deprecated. If this parameter and kubeProxyMode in ClusterSpec 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 network model. This field is supported only for the VPC network model (vpc-router).</p> <p>This parameter determines the number of container IP addresses that can be allocated to a node. The maximum number of pods that can be created on a node is decided by this parameter and maxPods set during node creation. For details, see Maximum Number of Pods That Can Be Created on a Node.</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 secure or normal . If this parameter is not specified, normal is used by default.

Parameter	Type	Description
kubernetes.io/ cpuManagerPolicy	String	Cluster CPU management policy. The value can be none or static . Defaults to none . <ul style="list-style-type: none"> • none: CPU cores will not be exclusively allocated to workload pods. Select this value if you want a large pool of shareable CPU cores. • static: CPU cores can be exclusively allocated to workload pods. Select this value if your workload is sensitive to latency in CPU cache and scheduling.
upgradefrom	String	Records of how the cluster is upgraded to the current version

Table 4-54 PackageConfiguration

Parameter	Type	Description
name	String	Component name.
configurations	Array of ConfigurationItem objects	Component configuration item.

Table 4-55 ConfigurationItem

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

Table 4-56 ClusterStatus

Parameter	Type	Description
phase	String	Cluster status. Possible values: <ul style="list-style-type: none"> • Available: The cluster is running properly. • Unavailable: The cluster is exhibiting unexpected behavior. Manually delete the cluster. • ScalingUp: Nodes are being added to the cluster. • ScalingDown: The cluster is being downsized to fewer nodes. • Creating: The cluster is being created. • Deleting: The cluster is being deleted. • Upgrading: The cluster is being upgraded. • Resizing: Cluster specifications are being changed. • RollingBack: The cluster is being rolled back. • RollbackFailed: The rollback is abnormal. • Empty: The cluster has no resources.
jobID	String	Job ID.
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 ClusterEndpoints 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 4-57 ClusterEndpoints

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

Example Requests

None

Example Responses

Status code: 200

Information about the specified cluster is successfully obtained.

```
{
  "kind": "Cluster",
  "apiVersion": "v3",
  "metadata": {
    "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": "v1.7.3-r13",
    "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"
    },
    "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 4-58 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .

Table 4-59 Query Parameters

Parameter	Mandatory	Type	Description
detail	No	String	<p>Whether the details about a cluster are queried.</p> <p>If this parameter is set to true, 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).</p>
status	No	String	<p>Cluster status. Possible values:</p> <ul style="list-style-type: none"> ● Available: The cluster is running properly. ● Unavailable: The cluster is exhibiting unexpected behavior. Manually delete the cluster. ● ScalingUp: Nodes are being added to the cluster. ● ScalingDown: The cluster is being downsized to fewer nodes. ● Creating: The cluster is being created. ● Deleting: The cluster is being deleted. ● Upgrading: The cluster is being upgraded. ● Resizing: Cluster specifications are being changed. ● RollingBack: The cluster is being rolled back. ● RollbackFailed: The rollback is abnormal. ● Empty: The cluster has no resources.

Parameter	Mandatory	Type	Description
type	No	String	Cluster type. Options: <ul style="list-style-type: none"> • VirtualMachine: CCE cluster • ARM64: Kunpeng cluster
version	No	String	Cluster version filtering.

Request Parameters

Table 4-60 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 Obtaining a User Token .

Response Parameters

Status code: 200

Table 4-61 Response body parameters

Parameter	Type	Description
kind	String	Api type
apiVersion	String	API version
items	Array of Cluster objects	A list of details for all clusters in the current project. You can filter clusters by items.metadata.name .

Table 4-62 Cluster

Parameter	Type	Description
kind	String	API type. The value is fixed at Cluster or cluster and cannot be changed.

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

Table 4-63 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 resource. The value is automatically generated after the object is created. A user-defined value will not take effect.
annotations	Map<String,String>	Cluster annotations, in the format of key-value pairs. <pre>"annotations": { "key1": "value1", "key2": "value2" }</pre> <p>NOTE</p> <ul style="list-style-type: none"> • annotations: Does not label or select objects. The metadata in annotations may be small or large, structured or unstructured, and may include characters that are not allowed in labels. • This field is not stored in the database and is used only to specify the add-ons to be installed in the cluster. • Install ICAgent during cluster creation by adding the key-value pair "cluster.install.addons.external/install":{"addonTemplateName":"icagent"}".
labels	Map<String,String>	Cluster labels, in the format of key-value pairs. <p>NOTE</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>

Parameter	Type	Description
creationTimestamp	String	Time when the cluster was created.
updateTimestamp	String	Time when the cluster was updated.

Table 4-64 ClusterSpec

Parameter	Type	Description
category	String	Cluster type: <ul style="list-style-type: none"> • CCE: CCE cluster. CCE cluster supports hybrid deployment of VMs and bare-metal servers (BMSs), and heterogeneous nodes such as GPU- and NPU-enabled nodes. You can run your containers in a secure and stable container runtime environment based on a high-performance network model.
type	String	Master node architecture: <ul style="list-style-type: none"> • VirtualMachine: x86

Parameter	Type	Description
flavor	String	<p>Default value: When you create a CCE cluster, the value is cce.s1.small for non-DeC scenarios and cce.dec.s1.small for DeC scenarios.</p> <p>Cluster flavor, which cannot be changed after the cluster is created.</p> <ul style="list-style-type: none"> ● cce.s1.small: small-scale, single-master CCE cluster (≤ 50 nodes) ● cce.s1.medium: medium-scale, single-master CCE cluster (≤ 200 nodes) ● cce.s2.small: small-scale, multi-master CCE cluster (≤ 50 nodes) ● cce.s2.medium: medium-scale, multi-master CCE cluster (≤ 200 nodes) ● cce.s2.large: large-scale, multi-master CCE cluster (≤ 1,000 nodes) ● cce.s2.xlarge: ultra-large-scale, multi-master CCE cluster (≤ 2,000 nodes) <p>NOTE</p> <ul style="list-style-type: none"> ● s1: single-master CCE cluster ● s2: multi-master CCE cluster ● dec: dedicated CCE cluster. For example, cce.dec.s1.small is a small-scale, single-master, dedicated CCE cluster (≤ 50 nodes). ● Values in the parentheses above indicate the maximum number of nodes that can be managed by the cluster. ● A single-master cluster has only one master node. If the master node is down, the cluster will become unavailable and stop serving new workloads. However, existing workloads in the cluster are not affected. ● A multi-master cluster is highly available. When a master node is faulty, the cluster is still available.

Parameter	Type	Description
version	String	<p>Cluster version, which mirrors the baseline version of the Kubernetes community. The latest version is recommended.</p> <p>You can create clusters of two latest versions on the CCE console. To learn which cluster versions are available, log in to the CCE console, create a cluster, and check the Version parameter. You can call APIs to create clusters of other versions. However, these clusters will be gradually brought offline. For details about the support policy, see the CCE announcement.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If not specified, a cluster of the latest version is created. • If the baseline cluster version is specified but the R version is not specified, the system selects the latest R version of the cluster by default. You are advised not to specify the R version.
platformVersion	String	<p>CCE cluster platform version, indicating the internal version under the cluster version (version). 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 platformVersion is cce.X.Y.</p> <ul style="list-style-type: none"> • X: internal feature version, indicating changes in features, patches, or OS support in the cluster version. The value starts from 1 and increases monotonically. • Y: 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 0 and increases monotonically.
description	String	<p>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 API for updating cluster information or go to the cluster details page on the CCE console. Only UTF-8 encoding is supported.</p>

Parameter	Type	Description
customSan	Array of strings	<p>Custom SAN field in the server certificate of the cluster API server, which must comply with the SSL and X509 format specifications.</p> <ol style="list-style-type: none"> 1. Duplicate names are not allowed. 2. Must comply with the IP address and domain name formats. <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 the cluster supports IPv6 addresses. This field is supported in clusters of v1.15 and later versions.
hostNetwork	HostNetwork 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	ContainerNetwork object	Container networking parameters, including the container network model and container CIDR block.
authentication	Authentication object	Configurations of the cluster authentication mode.
billingMode	Integer	<p>Billing mode of a cluster.</p> <ul style="list-style-type: none"> • 0: pay-per-use <p>Defaults to pay-per-use.</p>
masters	Array of MasterSpec objects	Advanced configurations of master nodes
kubernetesServiceIpRange	String	Service CIDR block or the IP address range which kubernetes clusterIp must fall within. This field is available only for clusters of v1.11.7 and later.
clusterTags	Array of ResourceTag objects	Cluster resource tags.

Parameter	Type	Description
kubeProxyMode	String	Service forwarding mode. Two modes are available: <ul style="list-style-type: none"> • iptables: 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. • ipvs: 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.
az	String	AZ. This field is returned only for a query.
extendParam	ClusterExtendedParam 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 PackageConfiguration objects	Whether to overwrite the default component configurations of the cluster. If a component or a parameter that is not supported by the component is specified, this configuration is ignored. For details about the supported configurable components and their parameters, see

Table 4-65 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"> • Method 1: Log in to the VPC console and view the VPC ID in the VPC details. • Method 2: Query the VPC ID through the VPC API. <p>NOTE Currently, the VPC network model does not support interconnection with VPCs that contain a secondary CIDR block.</p>

Parameter	Type	Description
subnet	String	Network ID of the subnet used to create a master node. Methods: <ul style="list-style-type: none"> • Method 1: Log in to VPC console and click the target subnet on the Subnets tab page. You can view the network ID on the displayed page. • Method 2: Use the VPC API for querying subnets. . .
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.

Table 4-66 ContainerNetwork

Parameter	Type	Description
mode	String	Container network model. Select one of the following possible values: <ul style="list-style-type: none"> • overlay_l2: an overlay_l2 network (container tunnel network) built for containers by using OpenVSwitch (OVS). • vpc-router: an underlay_l2 network built for containers by using IPVLAN and custom VPC routes.
cidr	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. Not editable after the cluster is created. (deprecated. A specified cidrs will make cidr invalid.)

Parameter	Type	Description
cidrs	Array of ContainerCIDR objects	List of container CIDR blocks. In clusters of v1.21 and later, the cidrs field is used. When the cluster network type is vpc-router , you can add multiple container CIDR blocks. In versions earlier than v1.21, if the cidrs field is used, the first CIDR element in the array is used as the container CIDR block. This parameter cannot be modified after the cluster is created.

Table 4-67 ContainerCIDR

Parameter	Type	Description
cidr	String	Container CIDR block. Recommended: 10.0.0.0/12-19, 172.16.0.0/16-19, and 192.168.0.0/16-19

Table 4-68 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"> Method 1: Log in to the VPC console and click the target subnet on the Subnets page. You can view the IPv4 subnet ID on the displayed page. Method 2: Use the VPC API for querying subnets.
eniSubnetCIDR	String	ENI subnet CIDR (being discarded)
subnets	Array of NetworkSubnet objects	List of IPv4 subnet IDs

Table 4-69 NetworkSubnet

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

Table 4-70 Authentication

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

Table 4-71 AuthenticatingProxy

Parameter	Type	Description
ca	String	X509 CA certificate (Base64-encoded) configured in <code>authenticating_proxy</code> mode. This field is mandatory when the cluster authentication mode is authenticating_proxy . Maximum size: 1 MB
cert	String	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 value must be Base64-encoded.) This field is mandatory when the cluster authentication mode is authenticating_proxy .

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 authenticating_proxy .

Table 4-72 MasterSpec

Parameter	Type	Description
availabilityZone	String	AZ

Table 4-73 ResourceTag

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

Table 4-74 ClusterExtendParam

Parameter	Type	Description
clusterAZ	String	<p>AZ of master nodes in the cluster.</p> <ul style="list-style-type: none"> • multi_az: (Optional) The cluster will span across AZs. This field is configurable only for high-availability clusters. • <i>AZ of the dedicated cloud computing pool</i>: The cluster will be deployed in the AZ of Dedicated Cloud (DeC). This parameter is mandatory for dedicated CCE clusters.
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><rootVol.dssPoolID>.<rootVol.volType>;<dataVol.dssPoolID>.<dataVol.volType></code></p> <p>Field description:</p> <ul style="list-style-type: none"> • rootVol is the system disk. dataVol is the data disk. • dssPoolID indicates the ID of the DSS storage pool. • volType indicates the storage volume type of the DSS storage pool, such as SAS and SSD. <p>Example: c950ee97-587c-4f24-8a74-3367e3da570f.sas; 6edbc2f4-1507-44f8-ac0d-eed1d2608d38.ssd</p> <p>NOTE 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>NOTE</p> <ul style="list-style-type: none"> • An enterprise project can be configured only after the enterprise project function is enabled. • The enterprise project to which the cluster belongs must be the same as that to which other cloud service resources associated with the cluster belong.

Parameter	Type	Description
kubeProxyMode	String	<p>Service forwarding mode. Two modes are available:</p> <ul style="list-style-type: none"> • iptables: 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. • ipvs: 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. <p>NOTE This parameter has been deprecated. If this parameter and kubeProxyMode in ClusterSpec 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 network model. This field is supported only for the VPC network model (vpc-router).</p> <p>This parameter determines the number of container IP addresses that can be allocated to a node. The maximum number of pods that can be created on a node is decided by this parameter and maxPods set during node creation. For details, see Maximum Number of Pods That Can Be Created on a Node.</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 secure or normal . If this parameter is not specified, normal is used by default.

Parameter	Type	Description
kubernetes.io/cpuManagerPolicy	String	Cluster CPU management policy. The value can be none or static . Defaults to none . <ul style="list-style-type: none"> • none: CPU cores will not be exclusively allocated to workload pods. Select this value if you want a large pool of shareable CPU cores. • static: CPU cores can be exclusively allocated to workload pods. Select this value if your workload is sensitive to latency in CPU cache and scheduling.
upgradefrom	String	Records of how the cluster is upgraded to the current version

Table 4-75 PackageConfiguration

Parameter	Type	Description
name	String	Component name.
configurations	Array of Configuratio nItem objects	Component configuration item.

Table 4-76 ConfigurationItem

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

Table 4-77 ClusterStatus

Parameter	Type	Description
phase	String	Cluster status. Possible values: <ul style="list-style-type: none"> • Available: The cluster is running properly. • Unavailable: The cluster is exhibiting unexpected behavior. Manually delete the cluster. • ScalingUp: Nodes are being added to the cluster. • ScalingDown: The cluster is being downsized to fewer nodes. • Creating: The cluster is being created. • Deleting: The cluster is being deleted. • Upgrading: The cluster is being upgraded. • Resizing: Cluster specifications are being changed. • RollingBack: The cluster is being rolled back. • RollbackFailed: The rollback is abnormal. • Empty: The cluster has no resources.
jobID	String	Job ID.
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 ClusterEndpoints 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 4-78 ClusterEndpoints

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

Example Requests

None

Example Responses

Status code: 200

The cluster list is successfully obtained.

```
{
  "kind": "Cluster",
  "apiVersion": "v3",
  "items": [ {
    "kind": "Cluster",
    "apiVersion": "v3",
    "metadata": {
      "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": "v1.7.3-r13",
      "description": "awesome cluster",
      "customSan": [ "192.168.1.0", "example.com" ],
      "hostNetwork": {
        "vpc": "f0c12911-4fdb-4284-9230-7ffb0860826a",
        "subnet": "ac274229-fd2e-4695-9f01-a0c1372b8006"
      },
      "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 4-79 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .

Request Parameters

Table 4-80 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 Obtaining a User Token .

Table 4-81 Request body parameters

Parameter	Mandatory	Type	Description
spec	Yes	ClusterInformationSpec object	Detailed cluster parameters

Table 4-82 ClusterInformationSpec

Parameter	Mandatory	Type	Description
description	No	String	Cluster description. <ol style="list-style-type: none"> 1. A maximum of 200 characters are allowed. The value cannot contain the following special characters: ~\$%^&*<>[]{}()'"#\ 2. Only clusters in the Available, ScalingUp, and ScalingDown states can be modified.

Parameter	Mandatory	Type	Description
customSan	No	Array of strings	<p>Custom SAN field in the server certificate of the cluster API server, which must comply with the SSL and X509 format specifications.</p> <ol style="list-style-type: none"> 1. Duplicate names are not allowed. 2. Must comply with the IP address and domain name formats. <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	ContainerNetworkUpdate object	Container networking parameters, including information about the container CIDR block.
eniNetwork	No	EniNetworkUpdate object	Cloud Native Network 2.0 network configuration, including the container subnet information of the CCE Turbo cluster.
hostNetwork	No	hostNetwork object	Node network parameters, including the default security group settings.

Table 4-83 ContainerNetworkUpdate

Parameter	Mandatory	Type	Description
cidrs	No	Array of ContainerCIDR objects	<p>List of container CIDR blocks. For clusters of v1.21 and later, if the cluster uses the VPC network model, container CIDR blocks can be added incrementally.</p> <p>The configuration cannot be changed after the cluster is updated.</p>

Table 4-84 ContainerCIDR

Parameter	Mandatory	Type	Description
cidr	Yes	String	Container CIDR block. Recommended: 10.0.0.0/12-19, 172.16.0.0/16-19, and 192.168.0.0/16-19

Table 4-85 EniNetworkUpdate

Parameter	Mandatory	Type	Description
subnets	No	Array of NetworkSubnet objects	List of IPv4 subnet IDs. CCE Turbo clusters of version 1.19.10 and later support multiple container subnets. In addition, the container subnet list can be incrementally updated. Subnets can be added, not deleted. Exercise caution when selecting a subnet. The request body must contain all existing subnets.

Table 4-86 NetworkSubnet

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

Table 4-87 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.

Response Parameters

Status code: 200

Table 4-88 Response body parameters

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

Table 4-89 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 resource. The value is automatically generated after the object is created. A user-defined value will not take effect.
annotations	Map<String,String>	Cluster annotations, in the format of key-value pairs. <pre>"annotations": { "key1": "value1", "key2": "value2" }</pre> <p>NOTE</p> <ul style="list-style-type: none"> • annotations: Does not label or select objects. The metadata in annotations may be small or large, structured or unstructured, and may include characters that are not allowed in labels. • This field is not stored in the database and is used only to specify the add-ons to be installed in the cluster. • Install ICAgent during cluster creation by adding the key-value pair "cluster.install.addons.external/install":{"addonTemplateName":"icagent"}".
labels	Map<String,String>	Cluster labels, in the format of key-value pairs. <p>NOTE</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 4-90 ClusterSpec

Parameter	Type	Description
category	String	Cluster type: <ul style="list-style-type: none"> • CCE: CCE cluster. CCE cluster supports hybrid deployment of VMs and bare-metal servers (BMSs), and heterogeneous nodes such as GPU- and NPU-enabled nodes. You can run your containers in a secure and stable container runtime environment based on a high-performance network model.
type	String	Master node architecture: <ul style="list-style-type: none"> • VirtualMachine: x86

Parameter	Type	Description
flavor	String	<p>Default value: When you create a CCE cluster, the value is cce.s1.small for non-DeC scenarios and cce.dec.s1.small for DeC scenarios.</p> <p>Cluster flavor, which cannot be changed after the cluster is created.</p> <ul style="list-style-type: none"> ● cce.s1.small: small-scale, single-master CCE cluster (≤ 50 nodes) ● cce.s1.medium: medium-scale, single-master CCE cluster (≤ 200 nodes) ● cce.s2.small: small-scale, multi-master CCE cluster (≤ 50 nodes) ● cce.s2.medium: medium-scale, multi-master CCE cluster (≤ 200 nodes) ● cce.s2.large: large-scale, multi-master CCE cluster (≤ 1,000 nodes) ● cce.s2.xlarge: ultra-large-scale, multi-master CCE cluster (≤ 2,000 nodes) <p>NOTE</p> <ul style="list-style-type: none"> ● s1: single-master CCE cluster ● s2: multi-master CCE cluster ● dec: dedicated CCE cluster. For example, cce.dec.s1.small is a small-scale, single-master, dedicated CCE cluster (≤ 50 nodes). ● Values in the parentheses above indicate the maximum number of nodes that can be managed by the cluster. ● A single-master cluster has only one master node. If the master node is down, the cluster will become unavailable and stop serving new workloads. However, existing workloads in the cluster are not affected. ● A multi-master cluster is highly available. When a master node is faulty, the cluster is still available.

Parameter	Type	Description
version	String	<p>Cluster version, which mirrors the baseline version of the Kubernetes community. The latest version is recommended.</p> <p>You can create clusters of two latest versions on the CCE console. To learn which cluster versions are available, log in to the CCE console, create a cluster, and check the Version parameter. You can call APIs to create clusters of other versions. However, these clusters will be gradually brought offline. For details about the support policy, see the CCE announcement.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If not specified, a cluster of the latest version is created. • If the baseline cluster version is specified but the R version is not specified, the system selects the latest R version of the cluster by default. You are advised not to specify the R version.
platformVersion	String	<p>CCE cluster platform version, indicating the internal version under the cluster version (version). 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 platformVersion is cce.X.Y.</p> <ul style="list-style-type: none"> • X: internal feature version, indicating changes in features, patches, or OS support in the cluster version. The value starts from 1 and increases monotonically. • Y: 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 0 and increases monotonically.
description	String	<p>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 API for updating cluster information or go to the cluster details page on the CCE console. Only UTF-8 encoding is supported.</p>

Parameter	Type	Description
customSan	Array of strings	<p>Custom SAN field in the server certificate of the cluster API server, which must comply with the SSL and X509 format specifications.</p> <ol style="list-style-type: none"> 1. Duplicate names are not allowed. 2. Must comply with the IP address and domain name formats. <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>
ipv6enable	Boolean	Whether the cluster supports IPv6 addresses. This field is supported in clusters of v1.15 and later versions.
hostNetwork	HostNetwork 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	ContainerNetwork object	Container networking parameters, including the container network model and container CIDR block.
authentication	Authentication object	Configurations of the cluster authentication mode.
billingMode	Integer	<p>Billing mode of a cluster.</p> <ul style="list-style-type: none"> • 0: pay-per-use <p>Defaults to pay-per-use.</p>
masters	Array of MasterSpec objects	Advanced configurations of master nodes
kubernetesServiceIpRange	String	Service CIDR block or the IP address range which kubernetes clusterIp must fall within. This field is available only for clusters of v1.11.7 and later.
clusterTags	Array of ResourceTag objects	Cluster resource tags.

Parameter	Type	Description
kubeProxyMode	String	Service forwarding mode. Two modes are available: <ul style="list-style-type: none"> • iptables: 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. • ipvs: 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.
az	String	AZ. This field is returned only for a query.
extendParam	ClusterExtendedParam 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 PackageConfiguration objects	Whether to overwrite the default component configurations of the cluster. If a component or a parameter that is not supported by the component is specified, this configuration is ignored. For details about the supported configurable components and their parameters, see

Table 4-91 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"> • Method 1: Log in to the VPC console and view the VPC ID in the VPC details. • Method 2: Query the VPC ID through the VPC API. <p>NOTE Currently, the VPC network model does not support interconnection with VPCs that contain a secondary CIDR block.</p>

Parameter	Type	Description
subnet	String	Network ID of the subnet used to create a master node. Methods: <ul style="list-style-type: none"> • Method 1: Log in to VPC console and click the target subnet on the Subnets tab page. You can view the network ID on the displayed page. • Method 2: Use the VPC API for querying subnets. . .
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.

Table 4-92 ContainerNetwork

Parameter	Type	Description
mode	String	Container network model. Select one of the following possible values: <ul style="list-style-type: none"> • overlay_l2: an overlay_l2 network (container tunnel network) built for containers by using OpenVSwitch (OVS). • vpc-router: an underlay_l2 network built for containers by using IPVLAN and custom VPC routes.
cidr	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. Not editable after the cluster is created. (deprecated. A specified cidrs will make cidr invalid.)

Parameter	Type	Description
cidrs	Array of ContainerCIDR objects	List of container CIDR blocks. In clusters of v1.21 and later, the cidrs field is used. When the cluster network type is vpc-router , you can add multiple container CIDR blocks. In versions earlier than v1.21, if the cidrs field is used, the first CIDR element in the array is used as the container CIDR block. This parameter cannot be modified after the cluster is created.

Table 4-93 ContainerCIDR

Parameter	Type	Description
cidr	String	Container CIDR block. Recommended: 10.0.0.0/12-19, 172.16.0.0/16-19, and 192.168.0.0/16-19

Table 4-94 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"> Method 1: Log in to the VPC console and click the target subnet on the Subnets page. You can view the IPv4 subnet ID on the displayed page. Method 2: Use the VPC API for querying subnets.
eniSubnetCIDR	String	ENI subnet CIDR (being discarded)
subnets	Array of NetworkSubnet objects	List of IPv4 subnet IDs

Table 4-95 NetworkSubnet

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

Table 4-96 Authentication

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

Table 4-97 AuthenticatingProxy

Parameter	Type	Description
ca	String	X509 CA certificate (Base64-encoded) configured in <code>authenticating_proxy</code> mode. This field is mandatory when the cluster authentication mode is authenticating_proxy . Maximum size: 1 MB
cert	String	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 value must be Base64-encoded.) This field is mandatory when the cluster authentication mode is authenticating_proxy .

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 authenticating_proxy .

Table 4-98 MasterSpec

Parameter	Type	Description
availabilityZone	String	AZ

Table 4-99 ResourceTag

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

Table 4-100 ClusterExtendParam

Parameter	Type	Description
clusterAZ	String	<p>AZ of master nodes in the cluster.</p> <ul style="list-style-type: none"> • multi_az: (Optional) The cluster will span across AZs. This field is configurable only for high-availability clusters. • <i>AZ of the dedicated cloud computing pool</i>: The cluster will be deployed in the AZ of Dedicated Cloud (DeC). This parameter is mandatory for dedicated CCE clusters.
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><rootVol.dssPoolID>.<rootVol.volType>;<dataVol.dssPoolID>.<dataVol.volType></code></p> <p>Field description:</p> <ul style="list-style-type: none"> • rootVol is the system disk. dataVol is the data disk. • dssPoolID indicates the ID of the DSS storage pool. • volType indicates the storage volume type of the DSS storage pool, such as SAS and SSD. <p>Example: c950ee97-587c-4f24-8a74-3367e3da570f.sas; 6edbc2f4-1507-44f8-ac0d-eed1d2608d38.ssd</p> <p>NOTE 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>NOTE</p> <ul style="list-style-type: none"> • An enterprise project can be configured only after the enterprise project function is enabled. • The enterprise project to which the cluster belongs must be the same as that to which other cloud service resources associated with the cluster belong.

Parameter	Type	Description
kubeProxyMode	String	<p>Service forwarding mode. Two modes are available:</p> <ul style="list-style-type: none"> • iptables: 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. • ipvs: 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. <p>NOTE This parameter has been deprecated. If this parameter and kubeProxyMode in ClusterSpec 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 network model. This field is supported only for the VPC network model (vpc-router).</p> <p>This parameter determines the number of container IP addresses that can be allocated to a node. The maximum number of pods that can be created on a node is decided by this parameter and maxPods set during node creation. For details, see Maximum Number of Pods That Can Be Created on a Node.</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 secure or normal . If this parameter is not specified, normal is used by default.

Parameter	Type	Description
kubernetes.io/ cpuManagerPolicy	String	Cluster CPU management policy. The value can be none or static . Defaults to none . <ul style="list-style-type: none"> • none: CPU cores will not be exclusively allocated to workload pods. Select this value if you want a large pool of shareable CPU cores. • static: CPU cores can be exclusively allocated to workload pods. Select this value if your workload is sensitive to latency in CPU cache and scheduling.
upgradefrom	String	Records of how the cluster is upgraded to the current version

Table 4-101 PackageConfiguration

Parameter	Type	Description
name	String	Component name.
configurations	Array of ConfigurationItem objects	Component configuration item.

Table 4-102 ConfigurationItem

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

Table 4-103 ClusterStatus

Parameter	Type	Description
phase	String	Cluster status. Possible values: <ul style="list-style-type: none"> • Available: The cluster is running properly. • Unavailable: The cluster is exhibiting unexpected behavior. Manually delete the cluster. • ScalingUp: Nodes are being added to the cluster. • ScalingDown: The cluster is being downsized to fewer nodes. • Creating: The cluster is being created. • Deleting: The cluster is being deleted. • Upgrading: The cluster is being upgraded. • Resizing: Cluster specifications are being changed. • RollingBack: The cluster is being rolled back. • RollbackFailed: The rollback is abnormal. • Empty: The cluster has no resources.
jobID	String	Job ID.
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 ClusterEndpoints 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 4-104 ClusterEndpoints

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

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" ]
  }
}
```

- Adding a container CIDR block for a cluster of v1.21 or later that uses the VPC network model

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

- Modifying the Security Group of the Default Node in the 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": {
    "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 4-105 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .

Table 4-106 Query Parameters

Parameter	Mandatory	Type	Description
delete_efs	No	String	Whether to delete SFS Turbo volumes. Value options: <ul style="list-style-type: none"> • true or block (The system starts to delete the object. If the deletion fails, subsequent processes are blocked.) • try (The system starts to delete the object. If the deletion fails, no deletion retry is performed, and subsequent processes are not blocked.) • false or skip (The object is not deleted. These are the default value options.)

Parameter	Mandatory	Type	Description
delete_evs	No	String	<p>Whether to delete EVS disks. Value options:</p> <ul style="list-style-type: none"> • true or block (The system starts to delete the object. If the deletion fails, subsequent processes are blocked.) • try (The system starts to delete the object. If the deletion fails, no deletion retry is performed, and subsequent processes are not blocked.) • false or skip (The object is not deleted. These are the default value options.)
delete_net	No	String	<p>Whether to delete cluster Service and ingress resources, such as ELB load balancers. Value options:</p> <ul style="list-style-type: none"> • true or block (The system starts to delete the object. If the deletion fails, subsequent processes are blocked.) • try (The system starts to delete the object. If the deletion fails, no deletion retry is performed, and subsequent processes are not blocked.) • false or skip (Skip the deletion.)

Parameter	Mandatory	Type	Description
delete_obs	No	String	<p>Whether to delete OBS volumes. Value options:</p> <ul style="list-style-type: none"> • true or block (The system starts to delete the object. If the deletion fails, subsequent processes are blocked.) • try (The system starts to delete the object. If the deletion fails, no deletion retry is performed, and subsequent processes are not blocked.) • false or skip (The object is not deleted. These are the default value options.)
delete_sfs	No	String	<p>Whether to delete SFS volumes. Value options:</p> <ul style="list-style-type: none"> • true or block (The system starts to delete the object. If the deletion fails, subsequent processes are blocked.) • try (The system starts to delete the object. If the deletion fails, no deletion retry is performed, and subsequent processes are not blocked.) • false or skip (The object is not deleted. These are the default value options.)

Parameter	Mandatory	Type	Description
delete_sfs30	No	String	Whether to delete an SFS 3.0 volume. Example value: <ul style="list-style-type: none"> • true or block (The system starts to delete the object. If the deletion fails, subsequent processes are blocked.) • try (The system starts to delete the object. If the deletion fails, no deletion retry is performed, and subsequent processes are not blocked.) • false or skip (The object is not deleted. These are the default value options.)

Request Parameters

Table 4-107 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 Obtaining a User Token .

Response Parameters

Status code: 200

Table 4-108 Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at Cluster or cluster and cannot be changed.

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

Table 4-109 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 resource. The value is automatically generated after the object is created. A user-defined value will not take effect.
annotations	Map<String,String>	Cluster annotations, in the format of key-value pairs. <pre>"annotations": { "key1": "value1", "key2": "value2" }</pre> <p>NOTE</p> <ul style="list-style-type: none"> • annotations: Does not label or select objects. The metadata in annotations may be small or large, structured or unstructured, and may include characters that are not allowed in labels. • This field is not stored in the database and is used only to specify the add-ons to be installed in the cluster. • Install ICAgent during cluster creation by adding the key-value pair "cluster.install.addons.external/install":{"addonTemplateName":"icagent"}".
labels	Map<String,String>	Cluster labels, in the format of key-value pairs. <p>NOTE</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>

Parameter	Type	Description
creationTimestamp	String	Time when the cluster was created.
updateTimestamp	String	Time when the cluster was updated.

Table 4-110 ClusterSpec

Parameter	Type	Description
category	String	Cluster type: <ul style="list-style-type: none"> • CCE: CCE cluster. CCE cluster supports hybrid deployment of VMs and bare-metal servers (BMSs), and heterogeneous nodes such as GPU- and NPU-enabled nodes. You can run your containers in a secure and stable container runtime environment based on a high-performance network model.
type	String	Master node architecture: <ul style="list-style-type: none"> • VirtualMachine: x86

Parameter	Type	Description
flavor	String	<p>Default value: When you create a CCE cluster, the value is cce.s1.small for non-DeC scenarios and cce.dec.s1.small for DeC scenarios.</p> <p>Cluster flavor, which cannot be changed after the cluster is created.</p> <ul style="list-style-type: none"> ● cce.s1.small: small-scale, single-master CCE cluster (≤ 50 nodes) ● cce.s1.medium: medium-scale, single-master CCE cluster (≤ 200 nodes) ● cce.s2.small: small-scale, multi-master CCE cluster (≤ 50 nodes) ● cce.s2.medium: medium-scale, multi-master CCE cluster (≤ 200 nodes) ● cce.s2.large: large-scale, multi-master CCE cluster (≤ 1,000 nodes) ● cce.s2.xlarge: ultra-large-scale, multi-master CCE cluster (≤ 2,000 nodes) <p>NOTE</p> <ul style="list-style-type: none"> ● s1: single-master CCE cluster ● s2: multi-master CCE cluster ● dec: dedicated CCE cluster. For example, cce.dec.s1.small is a small-scale, single-master, dedicated CCE cluster (≤ 50 nodes). ● Values in the parentheses above indicate the maximum number of nodes that can be managed by the cluster. ● A single-master cluster has only one master node. If the master node is down, the cluster will become unavailable and stop serving new workloads. However, existing workloads in the cluster are not affected. ● A multi-master cluster is highly available. When a master node is faulty, the cluster is still available.

Parameter	Type	Description
version	String	<p>Cluster version, which mirrors the baseline version of the Kubernetes community. The latest version is recommended.</p> <p>You can create clusters of two latest versions on the CCE console. To learn which cluster versions are available, log in to the CCE console, create a cluster, and check the Version parameter. You can call APIs to create clusters of other versions. However, these clusters will be gradually brought offline. For details about the support policy, see the CCE announcement.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If not specified, a cluster of the latest version is created. • If the baseline cluster version is specified but the R version is not specified, the system selects the latest R version of the cluster by default. You are advised not to specify the R version.
platformVersion	String	<p>CCE cluster platform version, indicating the internal version under the cluster version (version). 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 platformVersion is cce.X.Y.</p> <ul style="list-style-type: none"> • X: internal feature version, indicating changes in features, patches, or OS support in the cluster version. The value starts from 1 and increases monotonically. • Y: 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 0 and increases monotonically.
description	String	<p>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 API for updating cluster information or go to the cluster details page on the CCE console. Only UTF-8 encoding is supported.</p>

Parameter	Type	Description
customSan	Array of strings	<p>Custom SAN field in the server certificate of the cluster API server, which must comply with the SSL and X509 format specifications.</p> <ol style="list-style-type: none"> 1. Duplicate names are not allowed. 2. Must comply with the IP address and domain name formats. <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 the cluster supports IPv6 addresses. This field is supported in clusters of v1.15 and later versions.
hostNetwork	HostNetwork 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	ContainerNetwork object	Container networking parameters, including the container network model and container CIDR block.
authentication	Authentication object	Configurations of the cluster authentication mode.
billingMode	Integer	<p>Billing mode of a cluster.</p> <ul style="list-style-type: none"> • 0: pay-per-use <p>Defaults to pay-per-use.</p>
masters	Array of MasterSpec objects	Advanced configurations of master nodes
kubernetesServiceIpRange	String	Service CIDR block or the IP address range which kubernetes clusterIp must fall within. This field is available only for clusters of v1.11.7 and later.
clusterTags	Array of ResourceTag objects	Cluster resource tags.

Parameter	Type	Description
kubeProxyMode	String	Service forwarding mode. Two modes are available: <ul style="list-style-type: none"> • iptables: 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. • ipvs: 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.
az	String	AZ. This field is returned only for a query.
extendParam	ClusterExtendedParam 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 PackageConfiguration objects	Whether to overwrite the default component configurations of the cluster. If a component or a parameter that is not supported by the component is specified, this configuration is ignored. For details about the supported configurable components and their parameters, see

Table 4-111 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"> • Method 1: Log in to the VPC console and view the VPC ID in the VPC details. • Method 2: Query the VPC ID through the VPC API. NOTE Currently, the VPC network model does not support interconnection with VPCs that contain a secondary CIDR block.

Parameter	Type	Description
subnet	String	Network ID of the subnet used to create a master node. Methods: <ul style="list-style-type: none"> • Method 1: Log in to VPC console and click the target subnet on the Subnets tab page. You can view the network ID on the displayed page. • Method 2: Use the VPC API for querying subnets. . .
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.

Table 4-112 ContainerNetwork

Parameter	Type	Description
mode	String	Container network model. Select one of the following possible values: <ul style="list-style-type: none"> • overlay_l2: an overlay_l2 network (container tunnel network) built for containers by using OpenVSwitch (OVS). • vpc-router: an underlay_l2 network built for containers by using IPVLAN and custom VPC routes.
cidr	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. Not editable after the cluster is created. (deprecated. A specified cidrs will make cidr invalid.)

Parameter	Type	Description
cidrs	Array of ContainerCIDR objects	List of container CIDR blocks. In clusters of v1.21 and later, the cidrs field is used. When the cluster network type is vpc-router , you can add multiple container CIDR blocks. In versions earlier than v1.21, if the cidrs field is used, the first CIDR element in the array is used as the container CIDR block. This parameter cannot be modified after the cluster is created.

Table 4-113 ContainerCIDR

Parameter	Type	Description
cidr	String	Container CIDR block. Recommended: 10.0.0.0/12-19, 172.16.0.0/16-19, and 192.168.0.0/16-19

Table 4-114 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"> Method 1: Log in to the VPC console and click the target subnet on the Subnets page. You can view the IPv4 subnet ID on the displayed page. Method 2: Use the VPC API for querying subnets.
eniSubnetCIDR	String	ENI subnet CIDR (being discarded)
subnets	Array of NetworkSubnet objects	List of IPv4 subnet IDs

Table 4-115 NetworkSubnet

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

Table 4-116 Authentication

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

Table 4-117 AuthenticatingProxy

Parameter	Type	Description
ca	String	X509 CA certificate (Base64-encoded) configured in <code>authenticating_proxy</code> mode. This field is mandatory when the cluster authentication mode is authenticating_proxy . Maximum size: 1 MB
cert	String	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 value must be Base64-encoded.) This field is mandatory when the cluster authentication mode is authenticating_proxy .

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 authenticating_proxy .

Table 4-118 MasterSpec

Parameter	Type	Description
availabilityZone	String	AZ

Table 4-119 ResourceTag

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

Table 4-120 ClusterExtendParam

Parameter	Type	Description
clusterAZ	String	<p>AZ of master nodes in the cluster.</p> <ul style="list-style-type: none"> • multi_az: (Optional) The cluster will span across AZs. This field is configurable only for high-availability clusters. • <i>AZ of the dedicated cloud computing pool</i>: The cluster will be deployed in the AZ of Dedicated Cloud (DeC). This parameter is mandatory for dedicated CCE clusters.
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><rootVol.dssPoolID>.<rootVol.volType>;<dataVol.dssPoolID>.<dataVol.volType></code></p> <p>Field description:</p> <ul style="list-style-type: none"> • rootVol is the system disk. dataVol is the data disk. • dssPoolID indicates the ID of the DSS storage pool. • volType indicates the storage volume type of the DSS storage pool, such as SAS and SSD. <p>Example: c950ee97-587c-4f24-8a74-3367e3da570f.sas; 6edbc2f4-1507-44f8-ac0d-eed1d2608d38.ssd</p> <p>NOTE 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>NOTE</p> <ul style="list-style-type: none"> • An enterprise project can be configured only after the enterprise project function is enabled. • The enterprise project to which the cluster belongs must be the same as that to which other cloud service resources associated with the cluster belong.

Parameter	Type	Description
kubeProxyMode	String	<p>Service forwarding mode. Two modes are available:</p> <ul style="list-style-type: none"> • iptables: 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. • ipvs: 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. <p>NOTE This parameter has been deprecated. If this parameter and kubeProxyMode in ClusterSpec 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 network model. This field is supported only for the VPC network model (vpc-router).</p> <p>This parameter determines the number of container IP addresses that can be allocated to a node. The maximum number of pods that can be created on a node is decided by this parameter and maxPods set during node creation. For details, see Maximum Number of Pods That Can Be Created on a Node.</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 secure or normal . If this parameter is not specified, normal is used by default.

Parameter	Type	Description
kubernetes.io/ cpuManagerPolicy	String	Cluster CPU management policy. The value can be none or static . Defaults to none . <ul style="list-style-type: none"> • none: CPU cores will not be exclusively allocated to workload pods. Select this value if you want a large pool of shareable CPU cores. • static: CPU cores can be exclusively allocated to workload pods. Select this value if your workload is sensitive to latency in CPU cache and scheduling.
upgradefrom	String	Records of how the cluster is upgraded to the current version

Table 4-121 PackageConfiguration

Parameter	Type	Description
name	String	Component name.
configurations	Array of ConfigurationItem objects	Component configuration item.

Table 4-122 ConfigurationItem

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

Table 4-123 ClusterStatus

Parameter	Type	Description
phase	String	Cluster status. Possible values: <ul style="list-style-type: none"> • Available: The cluster is running properly. • Unavailable: The cluster is exhibiting unexpected behavior. Manually delete the cluster. • ScalingUp: Nodes are being added to the cluster. • ScalingDown: The cluster is being downsized to fewer nodes. • Creating: The cluster is being created. • Deleting: The cluster is being deleted. • Upgrading: The cluster is being upgraded. • Resizing: Cluster specifications are being changed. • RollingBack: The cluster is being rolled back. • RollbackFailed: The rollback is abnormal. • Empty: The cluster has no resources.
jobID	String	Job ID.
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 ClusterEndpoints 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 4-124 ClusterEndpoints

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

Example Requests

None

Example Responses

Status code: 200

The job for deleting a cluster is successfully delivered.

```
{
  "kind": "Cluster",
  "apiVersion": "v3",
  "metadata": {
    "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 4-125 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .

Request Parameters

Table 4-126 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 Obtaining a User Token .

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 Hibernation , 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 4-127 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .

Request Parameters

Table 4-128 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 Obtaining a User Token .

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 Available , 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 4-129 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .

Request Parameters

Table 4-130 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 Obtaining a User Token .

Table 4-131 Request body parameters

Parameter	Mandatory	Type	Description
duration	Yes	Integer	Validity period of the cluster certificate, in days. Min: 1. Max: 10950 (30 x 365, assuming there are always 365 days for a year, ignoring leap years). Value -1 means the maximum value, 30 years.

Response Parameters

Status code: 200

Table 4-132 Response header parameters

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

Table 4-133 Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at Config and cannot be changed.
apiVersion	String	API version. The value is fixed at v1 .
preferences	Object	This field is not used currently and is left unspecified by default.
clusters	Array of Clusters objects	Cluster list

Parameter	Type	Description
users	Array of Users objects	Certificate information and client key information of a specified user
contexts	Array of Contexts objects	Context list
current-context	String	Current context. If publicip (VM EIP) exists, the value is external . If publicip does not exist, the value is internal .

Table 4-134 Clusters

Parameter	Type	Description
name	String	Cluster name. <ul style="list-style-type: none"> If publicip 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 internalCluster. If publicip exists (that is, the EIP exists), there are at least two clusters in the cluster list, and the value of this parameter is externalCluster.
cluster	ClusterCert object	Cluster information

Table 4-135 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 externalCluster , the value is true .

Table 4-136 Users

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

Table 4-137 User

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

Table 4-138 Contexts

Parameter	Type	Description
name	String	Context name. <ul style="list-style-type: none"> If publicip 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 internal. If publicip 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 external.
context	Context object	Context information

Table 4-139 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*****FTkQgQ0VSVEIGSUNBVEUtLS0tLQo="    }  
  } ],  
  "users" : [ {  
    "name" : "user",  
    "user" : {  
      "client-certificate-data" : "LS0tLS1CRUdJTiBDR*****QVRFLS0tLS0K",  
      "client-key-data" : "LS0tLS1CRUdJTi*****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 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 4-140 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
job_id	Yes	String	Job ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .

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 Obtaining a User Token .

Response Parameters

Status code: 200

Table 4-142 Response body parameters

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

Table 4-143 JobSpec

Parameter	Type	Description
type	String	Job type. An example value is CreateCluster .
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 Job objects	Subjob list. <ul style="list-style-type: none"> The list contains details about all subjobs. Generally, a cluster/node creation job consists of multiple subjobs. The job is complete only after all subjobs are complete.

Table 4-144 Job

Parameter	Type	Description
kind	String	API type. The value is fixed at Job and cannot be changed.
apiVersion	String	API version. The value is fixed at v3 and cannot be changed.

Parameter	Type	Description
metadata	JobMetadata object	Job metadata
spec	JobSpec object	Detailed job parameters
status	JobStatus object	Job status

Table 4-145 JobMetadata

Parameter	Type	Description
uid	String	Job ID
creationTimesamp	String	Time when the job was created
updateTimesamp	String	Time when the job was updated

Table 4-146 JobStatus

Parameter	Type	Description
phase	String	Job status. Possible values: <ul style="list-style-type: none"> • JobPhaseInitializing JobPhase = "Initializing" • JobPhaseRunning JobPhase = "Running" • JobPhaseFailed JobPhase = "Failed" • JobPhaseSuccess JobPhase = "Success"
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](#).

4.2.10 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 4-147 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .

Request Parameters

Table 4-148 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 Obtaining a User Token .

Table 4-149 Request body parameters

Parameter	Mandatory	Type	Description
spec	Yes	MasterEIPRequestSpec object	Parameters in the request for binding or unbinding the public API Server address of a cluster

Table 4-150 MasterEIPRequestSpec

Parameter	Mandatory	Type	Description
action	No	String	Binding or unbinding. Mandatory. <ul style="list-style-type: none"> • Binding: The value is fixed at {"action": "bind"}. • Unbinding: The value is fixed at {"action": "unbind"}.
spec	No	spec 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 4-151 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 4-152 Response body parameters

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

Table 4-153 Metadata

Parameter	Type	Description
uid	String	Unique ID.
name	String	Add-on name.
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"> For add-on installation, the value is fixed at {"addon.install/type":"install"}. For add-on upgrade, the value is fixed at {"addon.upgrade/type":"upgrade"}.
updateTimestamp	String	Update time.
creationTimestamp	String	Creation time.

Table 4-154 MasterEIPResponseSpec

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

Table 4-155 spec

Parameter	Type	Description
id	String	ENI ID
eip	EipSpec object	EIP details
IsDynamic	Boolean	Dynamic provisioning or not

Table 4-156 EipSpec

Parameter	Type	Description
bandwidth	bandwidth object	Bandwidth information

Table 4-157 bandwidth

Parameter	Type	Description
size	Integer	Bandwidth size
sharetype	String	Bandwidth type

Table 4-158 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

```
{
  "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](#).

4.2.11 Obtaining Cluster Access Address

Function

This API is used to obtain the cluster access addresses by cluster ID, including the private IP (VIP returned by an HA cluster) and public IP.

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}/openapi`

Table 4-159 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .

Request Parameters

Table 4-160 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 Obtaining a User Token .

Response Parameters

Status code: 200

Table 4-161 Response body parameters

Parameter	Type	Description
metadata	Metadata object	Basic information about the object. Metadata is a collection of attributes.
spec	OpenAPISpec object	Parameters for configuring the address for accessing the cluster
status	status object	Status information

Table 4-162 Metadata

Parameter	Type	Description
uid	String	Unique ID.
name	String	Add-on name.
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"> For add-on installation, the value is fixed at {"addon.install/type":"install"}. For add-on upgrade, the value is fixed at {"addon.upgrade/type":"upgrade"}.
updateTimestamp	String	Update time.
creationTimestamp	String	Creation time.

Table 4-163 OpenAPISpec

Parameter	Type	Description
spec	spec object	Address for accessing the cluster

Table 4-164 spec

Parameter	Type	Description
eip	EipSpec object	EIP details
isDynamic	Boolean	Dynamic provisioning or not

Table 4-165 EipSpec

Parameter	Type	Description
bandwidth	bandwidth object	Bandwidth information

Table 4-166 bandwidth

Parameter	Type	Description
size	Integer	Bandwidth size
sharetype	String	Bandwidth type

Table 4-167 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

```
{
  "metadata": { },
  "spec": {
    "spec": {
      "eip": {
        "bandwidth": { }
      },
      "IsDynamic": false
    }
  },
  "status": {
    "privateEndpoint": "https://192.168.3.238:5443",
    "publicEndpoint": ""
  }
}
```

Example Responses

Status code: 200

The cluster access address is obtained successfully.

```
{
  "metadata": { },
  "spec": {
    "spec": {
      "eip": {
        "bandwidth": { }
      },
      "IsDynamic": false
    }
  },
  "status": {
    "privateEndpoint": "https://192.168.3.238:5443",
    "publicEndpoint": ""
  }
}
```

Status Codes

Status Code	Description
200	The cluster access address is obtained successfully.

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 4-168 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .

Table 4-169 Query Parameters

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

Request Parameters

Table 4-170 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 Obtaining a User Token .

Table 4-171 Request body parameters

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

Table 4-172 NodeMetadata

Parameter	Mandatory	Type	Description
name	No	String	<p>Node name.</p> <p>NOTE Enter 1 to 56 characters starting with a letter and not ending with a hyphen (-). Only lowercase letters, digits, and hyphens (-) are allowed. If name is not specified or empty, the node name is generated based on the default rule. The default rule is <i>Cluster name-Random characters</i>. Only the first 36 characters are kept for a too long cluster name. If the number of nodes (count) is greater than 1, some random characters are added to the end of the entered node name according to the default rule. The default rule is <i>User-defined name-Random characters</i>. Only the first 50 characters in the user-defined name are kept, followed by some random characters.</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:</p> <pre>"labels": { "key" : "value" }</pre>

Parameter	Mandatory	Type	Description
annotations	No	Map<String,String>	CCE node annotations, in key-value pairs (not the native Kubernetes annotations). Example: <pre>"annotations": { "key1" : "value1", "key2" : "value2" }</pre> NOTE annotations are not used to identify or select objects. The metadata in annotations may be small or large, structured or unstructured, and may include characters that are not allowed in labels.
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.

Table 4-173 NodeSpec

Parameter	Mandatory	Type	Description
flavor	Yes	String	Node specifications. For details about the node specifications supported by CCE, see Node Flavor Description .
az	Yes	String	Name of the AZ where the node to be created is located. . .

Parameter	Mandatory	Type	Description
os	No	String	<p>Node OS. For details about the supported OSs, see Node OS Description.</p> <p>NOTE</p> <ul style="list-style-type: none"> 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. If alpha.cce/NodeImageID in extendParam is specified during node creation, you do not need to set this field.
login	Yes	Login object	Node login mode. Either the key pair or password must be used for login.
rootVolume	Yes	Volume object	Information about disks on the node
dataVolumes	Yes	Array of Volume objects	Data disk parameters of the node. Currently, you can add the second data disk for your node on the CCE console. This data disk is used by the container runtime and kubelet. Do not uninstall this disk. Otherwise, the node will become unavailable. For DeC nodes, the parameter description is the same as that for rootVolume .

Parameter	Mandatory	Type	Description
storage	No	Storage object	<p>Disk initialization management parameter.</p> <p>This parameter is complex to configure. For details, see Attaching Disks to a Node.</p> <p>If this parameter retains its default, disks are managed based on the <code>DockerLVMConfigOverride</code> (discarded) parameter in <code>extendParam</code>. This parameter is supported by clusters of version 1.15.11 and later.</p> <p>NOTE</p> <p>If a node specification involves local disks and EVS disks at the same time, do not retain the default value of this parameter to prevent unexpected disk partitions.</p>
publicIP	No	NodePublicIP object	<p>EIP of a node.</p> <p>NOTE</p> <p>This parameter is not supported when you add a node to a node pool.</p>
nodeNicSpec	No	NodeNicSpec 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	<p>Node billing mode.</p> <ul style="list-style-type: none"> • 0: pay-per-use

Parameter	Mandatory	Type	Description
taints	No	Array of Taint 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"> • Key: A key must contain 1 to 63 characters starting with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain name can be used as the prefix of a key. • Value: A value must start with a letter or digit and can contain a maximum of 63 characters, including letters, digits, hyphens (-), underscores (_), and periods (.). • Effect: Available options are NoSchedule, PreferNoSchedule, and NoExecute. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In using 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>Example:</p> <pre data-bbox="1029 1608 1428 1843">"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"> • Key: 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 and contain a maximum of 253 characters. Example DNS subdomain: example.com/my-key • Value: The value can be left blank or contain 1 to 63 characters that start with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed in the character string. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In using 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>Example:</p> <pre>"k8sTags": { "key": "value" }</pre>
ecsGroupId	No	String	Cloud server group ID. If this field is specified, the node is created in the specified cloud server group.
dedicatedHostId	No	String	<p>ID of the DeH to which the node is scheduled.</p> <p>NOTE This parameter is not supported when you add a node during node pool creation.</p>

Parameter	Mandatory	Type	Description
userTags	No	Array of UserTag objects	Cloud server tag. The key of a tag must be unique. The maximum number of custom tags supported by CCE depends on the region 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 using 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.
runtime	No	Runtime object	Container runtime. Defaults to docker .

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

Table 4-174 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 4-175 UserPassword

Parameter	Mandatory	Type	Description
username	No	String	Login account. Defaults to root .
password	Yes	String	<p>If a username and a 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"> • Contains 8 to 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. The password field must be salted during node creation. For details, see Adding a Salt in the password Field When Creating a Node.

Table 4-176 Volume

Parameter	Mandatory	Type	Description
size	Yes	Integer	<p>Disk size, in GB.</p> <ul style="list-style-type: none"> • System disk: 40 to 1024 - Value range for data disks: 100 to 32768

Parameter	Mandatory	Type	Description
volumetype	Yes	String	Disk type. For details about possible values, see the description of the root_volume parameter in the API used to create an ECS in the ECS API reference. <ul style="list-style-type: none"> • SAS: high I/O disk • SSD: ultra-high I/O disk • SATA: common I/O disk SATA disks have been removed from EVS. You can find them attached only on existing nodes.
extendParam	No	Map<String,Object>	Extended disk parameters, defined in extendparam 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 dssPoolID , 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 dss . This field is used only for DeC clusters.
hw:passthrough	No	Boolean	<ul style="list-style-type: none"> • Pay attention to this field if your ECS is SDI-compliant. If the value of this field is true, an SCSI disk will be created. • If the node pool type is ElasticBMS, this field must be set to true. • If a node specification involves local disks and EVS disks at the same time, set the disk initialization parameters. For details, see Attaching Disks to a Node.

Parameter	Mandatory	Type	Description
metadata	No	VolumeMetadata 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 4-177 VolumeMetadata

Parameter	Mandatory	Type	Description
<code>__system__encrypted</code>	No	String	Whether the EVS disk is encrypted. The value 0 indicates that the EVS disk is not encrypted, and the value 1 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 metadata . This field is used with __system__encrypted .

Table 4-178 Storage

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

Table 4-179 StorageSelectors

Parameter	Mandatory	Type	Description
name	Yes	String	Selector name, used as the index of selectorNames in storageGroup . Therefore, the name of each selector must be unique.
storageType	Yes	String	Specifies the storage type. Currently, only evs (EVS volumes) and local (local volumes) are supported. The local storage does not support disk selection. All local disks will form a VG. Therefore, only one storageSelector of the local type is allowed.
matchLabels	No	matchLabels object	Matching field of an EVS volume. The size , volumeType , metadataEncrypted , metadataCmkid and count fields are supported.

Table 4-180 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. Currently, SSD, GPSSD and SAS are supported.
metadataEncrypted	No	String	Disk encryption identifier. 0 indicates that the disk is not encrypted, and 1 indicates that the disk is encrypted.
metadataCmkid	No	String	Customer master key ID of an encrypted disk. The value is a 36-byte string.
count	No	String	Number of disks to be selected. If this parameter is left blank, all disks of this type are selected.

Table 4-181 StorageGroups

Parameter	Mandatory	Type	Description
name	Yes	String	Name of a virtual storage group, which must be unique.
cceManaged	No	Boolean	Storage space for Kubernetes and runtime components. Only one group can be set to true . If this parameter is left blank, the default value false is used.
selectorNames	Yes	Array of strings	This parameter corresponds to name in storageSelectors . A group can match multiple selectors, but a selector can match only one group.
virtualSpaces	Yes	Array of VirtualSpace objects	Detailed management of space configuration in a group.

Table 4-182 VirtualSpace

Parameter	Mandatory	Type	Description
name	Yes	String	Name of a virtualSpace. <ul style="list-style-type: none"> • Kubernetes: Kubernetes space configuration. lvmConfig needs to be configured. • runtime: runtime space configuration. runtimeConfig needs to be configured. • user: user space configuration. lvmConfig needs to be configured.
size	Yes	String	Size of a virtualSpace. The value must be an integer in percentage. Example: 90%. <p>NOTE The sum of the percentages of all virtualSpaces in a group cannot exceed 100%.</p>

Parameter	Mandatory	Type	Description
lvmConfig	No	LVMConfig object	LVM configurations, applicable to kubernetes and user spaces. Note that one virtual space supports only one config.
runtimeConfig	No	RuntimeConfig object	runtime configurations, applicable to the runtime space. Note that one virtual space supports only one config.

Table 4-183 LVMConfig

Parameter	Mandatory	Type	Description
lvType	Yes	String	LVM write mode. linear indicates the linear mode. striped indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.
path	No	String	Path to which the disk is attached. This parameter takes effect only in user configuration. The value is an absolute path. Digits, letters, periods (.), hyphens (-), and underscores (_) are allowed.

Table 4-184 RuntimeConfig

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

Table 4-185 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. NOTE If ids has been set, you do not need to set count and eip .
count	No	Integer	Number of EIPs to be dynamically created. NOTE count and eip must be set at the same time.
eip	No	NodeEIPSpec object	EIP configuration.

Table 4-186 NodeEIPSpec

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

Table 4-187 NodeBandwidth

Parameter	Mandatory	Type	Description
chargemode	No	String	<p>Bandwidth billing mode.</p> <ul style="list-style-type: none"> If this field is not specified, the billing is based on bandwidth. - If the field is null, the billing is based on bandwidth. If the field value is traffic, the billing is based on traffic. If the value is out of the preceding options, the cloud server will fail to be created. <p>NOTE</p> <ul style="list-style-type: none"> 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%. 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%.
size	No	Integer	Bandwidth size, specified in bandwidth.size in the API for assigning an EIP. . .
sharetype	No	String	Bandwidth sharing type. Value options: PER (exclusive bandwidth)

Table 4-188 NodeNicSpec

Parameter	Mandatory	Type	Description
primaryNic	No	NicSpec object	Description of the primary NIC.

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

Table 4-189 NicSpec

Parameter	Mandatory	Type	Description
subnetId	No	String	ID of the subnet to which the NIC belongs. If subnetId is not specified when creating the primary NIC, the cluster subnet is used. When creating a secondary NIC, you must specify subnetId .
fixedIps	No	Array of strings	The IP address of the primary NIC is specified by fixedIps . The number of IP addresses cannot be greater than the number of created nodes. fixedIps and ipBlock cannot be specified at the same time.
ipBlock	No	String	CIDR format of the primary NIC IP range. The IP address of the created node falls in this range. fixedIps and ipBlock cannot be specified at the same time.

Table 4-190 Taint

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

Table 4-191 UserTag

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

Table 4-192 Runtime

Parameter	Mandatory	Type	Description
name	No	String	Container runtime. Defaults to docker .

Table 4-193 NodeExtendParam

Parameter	Mandatory	Type	Description
ecs:performancetype	No	String	ECS flavor types. This field is returned in the response.
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. The number of pods that can be created on a node is determined by multiple parameters. For details, see Maximum Number of Pods That Can Be Created on a Node .

Parameter	Mandatory	Type	Description
DockerLVMConfigOverride	No	String	<p>Docker data disk configuration item. (This parameter has been discarded. Use the storage 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 diskType based on the actual drive letter type. The following fields are included:</p> <ul style="list-style-type: none"> • userLV (optional): size of the user space, for example, vgpaas/20%VG. • userPath (optional): mount path of the user space, for example, /home/wqt-test. • diskType: disk type. Currently, only evs, hdd, and ssd are supported. • lvType: type of a logic volume. The value can be linear or striped. • dockerThinpool: Docker space size, for example, vgpaas/60%VG. • kubernetesLV: kubelet space size, for example, vgpaas/20%VG.

Parameter	Mandatory	Type	Description
dockerBaseSize	No	Integer	<p>Available disk space of a single container on a node, in GB.</p> <p>If this parameter is left blank or is set to 0, the default value is used. In Device Mapper mode, the default value is 10. In OverlayFS mode, the available space of a single container is not limited by default, and the dockerBaseSize setting takes effect only on EulerOS nodes in the cluster of the new version.</p> <p>For details about how to allocate the space for the container runtime, see Data Disk Space Allocation.</p> <p>When Device Mapper is used, you are advised to set dockerBaseSize 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>
publicKey	No	String	Public key of a node.
alpha.cce/preInstall	No	String	<p>Pre-installation script.</p> <p>NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)</p>
alpha.cce/postInstall	No	String	<p>Post-installation script.</p> <p>NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)</p>
alpha.cce/NodeImageID	No	String	This parameter is required when a custom image is used to create a BMS node.

Parameter	Mandatory	Type	Description
enterprise_project_id	No	String	ID of the enterprise project to which the node belongs
chargingMode	No	Integer	Billing mode of a node. This parameter has been deprecated. Use the billingMode parameter in NodeSpec .
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.

Response Parameters

Status code: 201

Table 4-194 Response body parameters

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

Table 4-195 NodeMetadata

Parameter	Type	Description
name	String	Node name. NOTE Enter 1 to 56 characters starting with a letter and not ending with a hyphen (-). Only lowercase letters, digits, and hyphens (-) are allowed. If name is not specified or empty, the node name is generated based on the default rule. The default rule is <i>Cluster name-Random characters</i> . Only the first 36 characters are kept for a too long cluster name. If the number of nodes (count) is greater than 1, some random characters are added to the end of the entered node name according to the default rule. The default rule is <i>User-defined name-Random characters</i> . Only the first 50 characters in the user-defined name are kept, followed by some random characters.
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>
annotations	Map<String,String>	CCE node annotations, in key-value pairs (not the native Kubernetes annotations). Example: <pre>"annotations": { "key1" : "value1", "key2" : "value2" }</pre> NOTE annotations are not used to identify or select objects. The metadata in annotations may be small or large, structured or unstructured, and may include characters that are not allowed in labels.
creationTimestamp	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.

Table 4-196 NodeSpec

Parameter	Type	Description
flavor	String	Node specifications. For details about the node specifications supported by CCE, see Node Flavor Description .
az	String	Name of the AZ where the node to be created is located. . .
os	String	Node OS. For details about the supported OSs, see Node OS Description . NOTE <ul style="list-style-type: none"> 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. If <code>alpha.cce/NodeImageID</code> in <code>extendParam</code> is specified during node creation, you do not need to set this field.
login	Login object	Node login mode. Either the key pair or password must be used for login.
rootVolume	Volume object	Information about disks on the node
dataVolumes	Array of Volume objects	Data disk parameters of the node. Currently, you can add the second data disk for your node on the CCE console. This data disk is used by the container runtime and kubelet. Do not uninstall this disk. Otherwise, the node will become unavailable. For DeC nodes, the parameter description is the same as that for rootVolume .
storage	Storage object	Disk initialization management parameter. This parameter is complex to configure. For details, see Attaching Disks to a Node . If this parameter retains its default, disks are managed based on the <code>DockerLVMConfigOverride (discarded)</code> parameter in <code>extendParam</code> . This parameter is supported by clusters of version 1.15.11 and later. NOTE If a node specification involves local disks and EVS disks at the same time, do not retain the default value of this parameter to prevent unexpected disk partitions.
publicIP	NodePublicIP object	EIP of a node. NOTE This parameter is not supported when you add a node to a node pool.

Parameter	Type	Description
nodeNicSpec	NodeNicSpec 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"> • 0: pay-per-use
taints	Array of Taint 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"> • Key: A key must contain 1 to 63 characters starting with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain name can be used as the prefix of a key. • Value: A value must start with a letter or digit and can contain a maximum of 63 characters, including letters, digits, hyphens (-), underscores (_), and periods (.). • Effect: Available options are NoSchedule, PreferNoSchedule, and NoExecute. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In using 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>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"> • Key: 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 and contain a maximum of 253 characters. Example DNS subdomain: example.com/my-key • Value: The value can be left blank or contain 1 to 63 characters that start with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed in the character string. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In using 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>Example:</p> <pre>"k8sTags": { "key": "value" }</pre>
ecsGroupId	String	Cloud server group ID. If this field is specified, the node is created in the specified cloud server group.
dedicatedHostId	String	<p>ID of the DeH to which the node is scheduled.</p> <p>NOTE This parameter is not supported when you add a node during node pool creation.</p>
userTags	Array of UserTag objects	Cloud server tag. The key of a tag must be unique. The maximum number of custom tags supported by CCE depends on the region 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 using 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.
runtime	Runtime object	Container runtime. Defaults to docker .

Parameter	Type	Description
initializedConditions	Array of strings	<p>Custom initialization flag.</p> <p>Before CCE nodes are initialized, they are tainted with node.cloudprovider.kubernetes.io/uninitialized to prevent pods from being scheduled to them.</p> <p>CCE supports custom initialization flags. After receiving the initializedConditions parameter, CCE converts the parameter value into a node label and provisions the label with the node, for example, cloudprovider.openvessel.io/inject-initialized-conditions=CCEInitial_CustomedInitial.</p> <p>After the node is labeled, its status.Conditions is polled to check whether the type of conditions has a flag name, such as CCEInitial and CustomedInitial. If all input flags exist and their status is True, the node initialization is complete and the initialization taint is removed.</p> <ul style="list-style-type: none"> • Use only letters and digits. Max. characters: 20. • Max. flags: 2.
extendParam	NodeExtendParam object	Extended parameters for creating a node.

Table 4-197 Login

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

Table 4-198 UserPassword

Parameter	Type	Description
username	String	Login account. Defaults to root .

Parameter	Type	Description
password	String	<p>If a username and a 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"> • Contains 8 to 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. The password field must be salted during node creation. For details, see Adding a Salt in the password Field When Creating a Node.

Table 4-199 Volume

Parameter	Type	Description
size	Integer	<p>Disk size, in GB.</p> <ul style="list-style-type: none"> • System disk: 40 to 1024 -Value range for data disks: 100 to 32768
volumetype	String	<p>Disk type. For details about possible values, see the description of the root_volume parameter in the API used to create an ECS in the ECS API reference.</p> <ul style="list-style-type: none"> • SAS: high I/O disk • SSD: ultra-high I/O disk • SATA: common I/O disk SATA disks have been removed from EVS. You can find them attached only on existing nodes.
extendParam	Map<String, Object>	<p>Extended disk parameters, defined in extendparam 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 dssPoolID, 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 dss. This field is used only for DeC clusters.</p>

Parameter	Type	Description
hw:passthrough	Boolean	<ul style="list-style-type: none"> Pay attention to this field if your ECS is SDI-compliant. If the value of this field is true, an SCSI disk will be created. If the node pool type is ElasticBMS, this field must be set to true. If a node specification involves local disks and EVS disks at the same time, set the disk initialization parameters. For details, see Attaching Disks to a Node.
metadata	VolumeMetadata 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 4-200 VolumeMetadata

Parameter	Type	Description
__system__encrypted	String	Whether the EVS disk is encrypted. The value 0 indicates that the EVS disk is not encrypted, and the value 1 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 metadata . This field is used with __system__encrypted .

Table 4-201 Storage

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

Table 4-202 StorageSelectors

Parameter	Type	Description
name	String	Selector name, used as the index of selectorNames in storageGroup . Therefore, the name of each selector must be unique.
storageType	String	Specifies the storage type. Currently, only evs (EVS volumes) and local (local volumes) are supported. The local storage does not support disk selection. All local disks will form a VG. Therefore, only one storageSelector of the local type is allowed.
matchLabels	matchLabels object	Matching field of an EVS volume. The size , volumeType , metadataEncrypted , metadataCmkid and count fields are supported.

Table 4-203 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. Currently, SSD, GPSSD and SAS are supported.
metadataEncrypted	String	Disk encryption identifier. 0 indicates that the disk is not encrypted, and 1 indicates that the disk is encrypted.
metadataCmkid	String	Customer master key ID of an encrypted disk. The value is a 36-byte string.
count	String	Number of disks to be selected. If this parameter is left blank, all disks of this type are selected.

Table 4-204 StorageGroups

Parameter	Type	Description
name	String	Name of a virtual storage group, which must be unique.

Parameter	Type	Description
cceManaged	Boolean	Storage space for Kubernetes and runtime components. Only one group can be set to true . If this parameter is left blank, the default value false is used.
selectorNames	Array of strings	This parameter corresponds to name in storageSelectors . A group can match multiple selectors, but a selector can match only one group.
virtualSpaces	Array of VirtualSpace objects	Detailed management of space configuration in a group.

Table 4-205 VirtualSpace

Parameter	Type	Description
name	String	Name of a virtualSpace. <ul style="list-style-type: none"> • Kubernetes: Kubernetes space configuration. lvmConfig needs to be configured. • runtime: runtime space configuration. runtimeConfig needs to be configured. • user: user space configuration. lvmConfig needs to be configured.
size	String	Size of a virtualSpace. The value must be an integer in percentage. Example: 90%. NOTE The sum of the percentages of all virtualSpaces in a group cannot exceed 100%.
lvmConfig	LVMConfig object	LVM configurations, applicable to kubernetes and user spaces. Note that one virtual space supports only one config.
runtimeConfig	RuntimeConfig object	runtime configurations, applicable to the runtime space. Note that one virtual space supports only one config.

Table 4-206 LVMConfig

Parameter	Type	Description
lvType	String	LVM write mode. linear indicates the linear mode. striped indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.
path	String	Path to which the disk is attached. This parameter takes effect only in user configuration. The value is an absolute path. Digits, letters, periods (.), hyphens (-), and underscores (_) are allowed.

Table 4-207 RuntimeConfig

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

Table 4-208 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. NOTE If ids has been set, you do not need to set count and eip .
count	Integer	Number of EIPs to be dynamically created. NOTE count and eip must be set at the same time.
eip	NodeEIPSpec object	EIP configuration.

Table 4-209 NodeEIPSpec

Parameter	Type	Description
iptype	String	EIP type, specified in publicip.type in the API for assigning an EIP. . .

Parameter	Type	Description
bandwidth	NodeBandwidth object	Bandwidth parameters of the EIP

Table 4-210 NodeBandwidth

Parameter	Type	Description
chargemode	String	Bandwidth billing mode. <ul style="list-style-type: none"> If this field is not specified, the billing is based on bandwidth. - If the field is null, the billing is based on bandwidth. If the field value is traffic, the billing is based on traffic. If the value is out of the preceding options, the cloud server will fail to be created. <p>NOTE</p> <ul style="list-style-type: none"> 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%. 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%.
size	Integer	Bandwidth size, specified in bandwidth.size in the API for assigning an EIP. . .
sharetype	String	Bandwidth sharing type. Value options: PER (exclusive bandwidth)

Table 4-211 NodeNicSpec

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

Table 4-212 NicSpec

Parameter	Type	Description
subnetId	String	ID of the subnet to which the NIC belongs. If subnetId is not specified when creating the primary NIC, the cluster subnet is used. When creating a secondary NIC, you must specify subnetId .
fixedIps	Array of strings	The IP address of the primary NIC is specified by fixedIps . The number of IP addresses cannot be greater than the number of created nodes. fixedIps and ipBlock cannot be specified at the same time.
ipBlock	String	CIDR format of the primary NIC IP range. The IP address of the created node falls in this range. fixedIps and ipBlock cannot be specified at the same time.

Table 4-213 Taint

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

Table 4-214 UserTag

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

Table 4-215 Runtime

Parameter	Type	Description
name	String	Container runtime. Defaults to docker .

Table 4-216 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 Maximum Number of Pods That Can Be Created on a Node.</p>
DockerLVMDiskConfigOverride	String	<p>Docker data disk configuration item. (This parameter has been discarded. Use the storage field instead.) Example default configuration: <code>"DockerLVMDiskConfigOverride": "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 diskType based on the actual drive letter type. The following fields are included:</p> <ul style="list-style-type: none"> • userLV (optional): size of the user space, for example, vgpaas/20%VG. • userPath (optional): mount path of the user space, for example, /home/wqt-test. • diskType: disk type. Currently, only evs, hdd, and ssd are supported. • lvType: type of a logic volume. The value can be linear or striped. • dockerThinpool: Docker space size, for example, vgpaas/60%VG. • kubernetesLV: kubelet space size, for example, vgpaas/20%VG.

Parameter	Type	Description
dockerBaseSize	Integer	<p>Available disk space of a single container on a node, in GB.</p> <p>If this parameter is left blank or is set to 0, the default value is used. In Device Mapper mode, the default value is 10. In OverlayFS mode, the available space of a single container is not limited by default, and the dockerBaseSize setting takes effect only on EulerOS nodes in the cluster of the new version.</p> <p>For details about how to allocate the space for the container runtime, see Data Disk Space Allocation.</p> <p>When Device Mapper is used, you are advised to set dockerBaseSize 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>
publicKey	String	Public key of a node.
alpha.cce/preInstall	String	<p>Pre-installation script.</p> <p>NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)</p>
alpha.cce/postInstall	String	<p>Post-installation script.</p> <p>NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)</p>
alpha.cce/NodeImageID	String	This parameter is required when a custom image is used to create a BMS node.
enterprise_project_id	String	ID of the enterprise project to which the node belongs
chargingMode	Integer	Billing mode of a node. This parameter has been deprecated. Use the billingMode parameter in NodeSpec .
agency_name	String	<p>Name of an agency</p> <p>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.</p>

Table 4-217 NodeStatus

Parameter	Type	Description
phase	String	Node status
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 Sync Node Data on the console to manually update the data.
deleteStatus	DeleteStatus object	Resource status during resource deletion.

Table 4-218 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.
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

Creating a pay-per-use node

```
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
}

```

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" : "5ecfdfe-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",
    "enterprise_project_id": "0",
    "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 4-219 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
node_id	Yes	String	Node ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .

Request Parameters

Table 4-220 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 Obtaining a User Token .

Response Parameters

Status code: 200

Table 4-221 Response body parameters

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

Table 4-222 NodeMetadata

Parameter	Type	Description
name	String	Node name. NOTE Enter 1 to 56 characters starting with a letter and not ending with a hyphen (-). Only lowercase letters, digits, and hyphens (-) are allowed. If name is not specified or empty, the node name is generated based on the default rule. The default rule is <i>Cluster name-Random characters</i> . Only the first 36 characters are kept for a too long cluster name. If the number of nodes (count) is greater than 1, some random characters are added to the end of the entered node name according to the default rule. The default rule is <i>User-defined name-Random characters</i> . Only the first 50 characters in the user-defined name are kept, followed by some random characters.
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>
annotations	Map<String,String>	CCE node annotations, in key-value pairs (not the native Kubernetes annotations). Example: <pre>"annotations": { "key1" : "value1", "key2" : "value2" }</pre> NOTE annotations are not used to identify or select objects. The metadata in annotations may be small or large, structured or unstructured, and may include characters that are not allowed in labels.
creationTimestamp	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.

Table 4-223 NodeSpec

Parameter	Type	Description
flavor	String	Node specifications. For details about the node specifications supported by CCE, see Node Flavor Description .
az	String	Name of the AZ where the node to be created is located. . .
os	String	Node OS. For details about the supported OSs, see Node OS Description . NOTE <ul style="list-style-type: none"> 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. If <code>alpha.cce/NodeImageID</code> in <code>extendParam</code> is specified during node creation, you do not need to set this field.
login	Login object	Node login mode. Either the key pair or password must be used for login.
rootVolume	Volume object	Information about disks on the node
dataVolumes	Array of Volume objects	Data disk parameters of the node. Currently, you can add the second data disk for your node on the CCE console. This data disk is used by the container runtime and kubelet. Do not uninstall this disk. Otherwise, the node will become unavailable. For DeC nodes, the parameter description is the same as that for rootVolume .
storage	Storage object	Disk initialization management parameter. This parameter is complex to configure. For details, see Attaching Disks to a Node . If this parameter retains its default, disks are managed based on the <code>DockerLVMConfigOverride (discarded)</code> parameter in <code>extendParam</code> . This parameter is supported by clusters of version 1.15.11 and later. NOTE If a node specification involves local disks and EVS disks at the same time, do not retain the default value of this parameter to prevent unexpected disk partitions.
publicIP	NodePublicIP object	EIP of a node. NOTE This parameter is not supported when you add a node to a node pool.

Parameter	Type	Description
nodeNicSpec	NodeNicSpec 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"> 0: pay-per-use
taints	Array of Taint 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"> Key: A key must contain 1 to 63 characters starting with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain name can be used as the prefix of a key. Value: A value must start with a letter or digit and can contain a maximum of 63 characters, including letters, digits, hyphens (-), underscores (_), and periods (.). Effect: Available options are NoSchedule, PreferNoSchedule, and NoExecute. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In using 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>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"> • Key: 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 and contain a maximum of 253 characters. Example DNS subdomain: example.com/my-key • Value: The value can be left blank or contain 1 to 63 characters that start with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed in the character string. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In using 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>Example:</p> <pre>"k8sTags": { "key": "value" }</pre>
ecsGroupId	String	Cloud server group ID. If this field is specified, the node is created in the specified cloud server group.
dedicatedHostId	String	<p>ID of the DeH to which the node is scheduled.</p> <p>NOTE This parameter is not supported when you add a node during node pool creation.</p>
userTags	Array of UserTag objects	Cloud server tag. The key of a tag must be unique. The maximum number of custom tags supported by CCE depends on the region 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 using 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.
runtime	Runtime object	Container runtime. Defaults to docker .

Parameter	Type	Description
initializedConditions	Array of strings	<p>Custom initialization flag.</p> <p>Before CCE nodes are initialized, they are tainted with node.cloudprovider.kubernetes.io/uninitialized to prevent pods from being scheduled to them.</p> <p>CCE supports custom initialization flags. After receiving the initializedConditions parameter, CCE converts the parameter value into a node label and provisions the label with the node, for example, cloudprovider.openvessel.io/inject-initialized-conditions=CCEInitial_CustomedInitial.</p> <p>After the node is labeled, its status.Conditions is polled to check whether the type of conditions has a flag name, such as CCEInitial and CustomedInitial. If all input flags exist and their status is True, the node initialization is complete and the initialization taint is removed.</p> <ul style="list-style-type: none"> • Use only letters and digits. Max. characters: 20. • Max. flags: 2.
extendParam	NodeExtendParam object	Extended parameters for creating a node.

Table 4-224 Login

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

Table 4-225 UserPassword

Parameter	Type	Description
username	String	Login account. Defaults to root .

Parameter	Type	Description
password	String	<p>If a username and a 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"> • Contains 8 to 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. The password field must be salted during node creation. For details, see Adding a Salt in the password Field When Creating a Node.

Table 4-226 Volume

Parameter	Type	Description
size	Integer	<p>Disk size, in GB.</p> <ul style="list-style-type: none"> • System disk: 40 to 1024 -Value range for data disks: 100 to 32768
volumetype	String	<p>Disk type. For details about possible values, see the description of the root_volume parameter in the API used to create an ECS in the ECS API reference.</p> <ul style="list-style-type: none"> • SAS: high I/O disk • SSD: ultra-high I/O disk • SATA: common I/O disk SATA disks have been removed from EVS. You can find them attached only on existing nodes.
extendParam	Map<String, Object>	<p>Extended disk parameters, defined in extendparam 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 dssPoolID, 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 dss. This field is used only for DeC clusters.</p>

Parameter	Type	Description
hw:passthrough	Boolean	<ul style="list-style-type: none"> Pay attention to this field if your ECS is SDI-compliant. If the value of this field is true, an SCSI disk will be created. If the node pool type is ElasticBMS, this field must be set to true. If a node specification involves local disks and EVS disks at the same time, set the disk initialization parameters. For details, see Attaching Disks to a Node.
metadata	VolumeMetadata 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 4-227 VolumeMetadata

Parameter	Type	Description
__system__encrypted	String	Whether the EVS disk is encrypted. The value 0 indicates that the EVS disk is not encrypted, and the value 1 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 metadata . This field is used with __system__encrypted .

Table 4-228 Storage

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

Table 4-229 StorageSelectors

Parameter	Type	Description
name	String	Selector name, used as the index of selectorNames in storageGroup . Therefore, the name of each selector must be unique.
storageType	String	Specifies the storage type. Currently, only evs (EVS volumes) and local (local volumes) are supported. The local storage does not support disk selection. All local disks will form a VG. Therefore, only one storageSelector of the local type is allowed.
matchLabels	matchLabels object	Matching field of an EVS volume. The size , volumeType , metadataEncrypted , metadataCmkid and count fields are supported.

Table 4-230 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. Currently, SSD, GPSSD and SAS are supported.
metadataEncrypted	String	Disk encryption identifier. 0 indicates that the disk is not encrypted, and 1 indicates that the disk is encrypted.
metadataCmkid	String	Customer master key ID of an encrypted disk. The value is a 36-byte string.
count	String	Number of disks to be selected. If this parameter is left blank, all disks of this type are selected.

Table 4-231 StorageGroups

Parameter	Type	Description
name	String	Name of a virtual storage group, which must be unique.

Parameter	Type	Description
cceManaged	Boolean	Storage space for Kubernetes and runtime components. Only one group can be set to true . If this parameter is left blank, the default value false is used.
selectorNames	Array of strings	This parameter corresponds to name in storageSelectors . A group can match multiple selectors, but a selector can match only one group.
virtualSpaces	Array of VirtualSpace objects	Detailed management of space configuration in a group.

Table 4-232 VirtualSpace

Parameter	Type	Description
name	String	Name of a virtualSpace. <ul style="list-style-type: none"> • Kubernetes: Kubernetes space configuration. lvmConfig needs to be configured. • runtime: runtime space configuration. runtimeConfig needs to be configured. • user: user space configuration. lvmConfig needs to be configured.
size	String	Size of a virtualSpace. The value must be an integer in percentage. Example: 90%. NOTE The sum of the percentages of all virtualSpaces in a group cannot exceed 100%.
lvmConfig	LVMConfig object	LVM configurations, applicable to kubernetes and user spaces. Note that one virtual space supports only one config.
runtimeConfig	RuntimeConfig object	runtime configurations, applicable to the runtime space. Note that one virtual space supports only one config.

Table 4-233 LVMConfig

Parameter	Type	Description
lvType	String	LVM write mode. linear indicates the linear mode. striped indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.
path	String	Path to which the disk is attached. This parameter takes effect only in user configuration. The value is an absolute path. Digits, letters, periods (.), hyphens (-), and underscores (_) are allowed.

Table 4-234 RuntimeConfig

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

Table 4-235 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. NOTE If ids has been set, you do not need to set count and eip .
count	Integer	Number of EIPs to be dynamically created. NOTE count and eip must be set at the same time.
eip	NodeEIPSpec object	EIP configuration.

Table 4-236 NodeEIPSpec

Parameter	Type	Description
iptype	String	EIP type, specified in publicip.type in the API for assigning an EIP. . .

Parameter	Type	Description
bandwidth	NodeBandwidth object	Bandwidth parameters of the EIP

Table 4-237 NodeBandwidth

Parameter	Type	Description
chargemode	String	<p>Bandwidth billing mode.</p> <ul style="list-style-type: none"> If this field is not specified, the billing is based on bandwidth. - If the field is null, the billing is based on bandwidth. If the field value is traffic, the billing is based on traffic. If the value is out of the preceding options, the cloud server will fail to be created. <p>NOTE</p> <ul style="list-style-type: none"> 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%. 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%.
size	Integer	Bandwidth size, specified in bandwidth.size in the API for assigning an EIP. . .
sharetype	String	Bandwidth sharing type. Value options: PER (exclusive bandwidth)

Table 4-238 NodeNicSpec

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

Table 4-239 NicSpec

Parameter	Type	Description
subnetId	String	ID of the subnet to which the NIC belongs. If subnetId is not specified when creating the primary NIC, the cluster subnet is used. When creating a secondary NIC, you must specify subnetId .
fixedIps	Array of strings	The IP address of the primary NIC is specified by fixedIps . The number of IP addresses cannot be greater than the number of created nodes. fixedIps and ipBlock cannot be specified at the same time.
ipBlock	String	CIDR format of the primary NIC IP range. The IP address of the created node falls in this range. fixedIps and ipBlock cannot be specified at the same time.

Table 4-240 Taint

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

Table 4-241 UserTag

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

Table 4-242 Runtime

Parameter	Type	Description
name	String	Container runtime. Defaults to docker .

Table 4-243 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 Maximum Number of Pods That Can Be Created on a Node.</p>
DockerLVMDockerConfigOverride	String	<p>Docker data disk configuration item. (This parameter has been discarded. Use the storage field instead.) Example default configuration: <code>"DockerLVMDockerConfigOverride": "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 diskType based on the actual drive letter type. The following fields are included:</p> <ul style="list-style-type: none"> • userLV (optional): size of the user space, for example, vgpaas/20%VG. • userPath (optional): mount path of the user space, for example, /home/wqt-test. • diskType: disk type. Currently, only evs, hdd, and ssd are supported. • lvType: type of a logic volume. The value can be linear or striped. • dockerThinpool: Docker space size, for example, vgpaas/60%VG. • kubernetesLV: kubelet space size, for example, vgpaas/20%VG.

Parameter	Type	Description
dockerBaseSize	Integer	<p>Available disk space of a single container on a node, in GB.</p> <p>If this parameter is left blank or is set to 0, the default value is used. In Device Mapper mode, the default value is 10. In OverlayFS mode, the available space of a single container is not limited by default, and the dockerBaseSize setting takes effect only on EulerOS nodes in the cluster of the new version.</p> <p>For details about how to allocate the space for the container runtime, see Data Disk Space Allocation.</p> <p>When Device Mapper is used, you are advised to set dockerBaseSize 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>
publicKey	String	Public key of a node.
alpha.cce/preInstall	String	<p>Pre-installation script.</p> <p>NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)</p>
alpha.cce/postInstall	String	<p>Post-installation script.</p> <p>NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)</p>
alpha.cce/NodeImageID	String	This parameter is required when a custom image is used to create a BMS node.
enterprise_project_id	String	ID of the enterprise project to which the node belongs
chargingMode	Integer	Billing mode of a node. This parameter has been deprecated. Use the billingMode parameter in NodeSpec .
agency_name	String	<p>Name of an agency</p> <p>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.</p>

Table 4-244 NodeStatus

Parameter	Type	Description
phase	String	Node status
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 Sync Node Data on the console to manually update the data.
deleteStatus	DeleteStatus object	Resource status during resource deletion.

Table 4-245 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.
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 4-246 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .

Request Parameters

Table 4-247 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 Obtaining a User Token .

Response Parameters

Status code: 200

Table 4-248 Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at List .

Parameter	Type	Description
apiVersion	String	API version. The value is fixed at v3 .
items	Array of Node objects	List of details for all nodes in the current cluster. You can filter nodes by items.metadata.name .

Table 4-249 Node

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

Table 4-250 NodeMetadata

Parameter	Type	Description
name	String	Node name. NOTE Enter 1 to 56 characters starting with a letter and not ending with a hyphen (-). Only lowercase letters, digits, and hyphens (-) are allowed. If name is not specified or empty, the node name is generated based on the default rule. The default rule is <i>Cluster name-Random characters</i> . Only the first 36 characters are kept for a too long cluster name. If the number of nodes (count) is greater than 1, some random characters are added to the end of the entered node name according to the default rule. The default rule is <i>User-defined name-Random characters</i> . Only the first 50 characters in the user-defined name are kept, followed by some random characters.
uid	String	Node ID, which is unique and automatically generated after the resource is created. A user-defined ID will not take effect.

Parameter	Type	Description
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>NOTE annotations are not used to identify or select objects. The metadata in annotations may be small or large, structured or unstructured, and may include characters that are not allowed in labels.</p>
creationTimestamp	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.

Table 4-251 NodeSpec

Parameter	Type	Description
flavor	String	Node specifications. For details about the node specifications supported by CCE, see Node Flavor Description .
az	String	Name of the AZ where the node to be created is located. . .
os	String	<p>Node OS. For details about the supported OSs, see Node OS Description.</p> <p>NOTE</p> <ul style="list-style-type: none"> 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. If <code>alpha.cce/NodeImageID</code> in <code>extendParam</code> is specified during node creation, you do not need to set this field.

Parameter	Type	Description
login	Login object	Node login mode. Either the key pair or password must be used for login.
rootVolume	Volume object	Information about disks on the node
dataVolumes	Array of Volume objects	Data disk parameters of the node. Currently, you can add the second data disk for your node on the CCE console. This data disk is used by the container runtime and kubelet. Do not uninstall this disk. Otherwise, the node will become unavailable. For DeC nodes, the parameter description is the same as that for rootVolume .
storage	Storage object	<p>Disk initialization management parameter. This parameter is complex to configure. For details, see Attaching Disks to a Node.</p> <p>If this parameter retains its default, disks are managed based on the DockerLVMConfigOverride (discarded) parameter in extendParam. This parameter is supported by clusters of version 1.15.11 and later.</p> <p>NOTE If a node specification involves local disks and EVS disks at the same time, do not retain the default value of this parameter to prevent unexpected disk partitions.</p>
publicIP	NodePublicIP object	<p>EIP of a node.</p> <p>NOTE This parameter is not supported when you add a node to a node pool.</p>
nodeNicSpec	NodeNicSpec 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"> • 0: pay-per-use

Parameter	Type	Description
taints	Array of Taint 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"> • Key: A key must contain 1 to 63 characters starting with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain name can be used as the prefix of a key. • Value: A value must start with a letter or digit and can contain a maximum of 63 characters, including letters, digits, hyphens (-), underscores (_), and periods (.). • Effect: Available options are NoSchedule, PreferNoSchedule, and NoExecute. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In using 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>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"> • Key: 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 and contain a maximum of 253 characters. Example DNS subdomain: example.com/my-key • Value: The value can be left blank or contain 1 to 63 characters that start with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed in the character string. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In using 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>Example:</p> <pre>"k8sTags": { "key": "value" }</pre>
ecsGroupId	String	Cloud server group ID. If this field is specified, the node is created in the specified cloud server group.
dedicatedHostId	String	<p>ID of the DeH to which the node is scheduled.</p> <p>NOTE This parameter is not supported when you add a node during node pool creation.</p>
userTags	Array of UserTag objects	Cloud server tag. The key of a tag must be unique. The maximum number of custom tags supported by CCE depends on the region 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 using 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.
runtime	Runtime object	Container runtime. Defaults to docker .

Parameter	Type	Description
initializedConditions	Array of strings	<p>Custom initialization flag.</p> <p>Before CCE nodes are initialized, they are tainted with node.cloudprovider.kubernetes.io/uninitialized to prevent pods from being scheduled to them.</p> <p>CCE supports custom initialization flags. After receiving the initializedConditions parameter, CCE converts the parameter value into a node label and provisions the label with the node, for example, cloudprovider.openvessel.io/inject-initialized-conditions=CCEInitial_CustomedInitial.</p> <p>After the node is labeled, its status.Conditions is polled to check whether the type of conditions has a flag name, such as CCEInitial and CustomedInitial. If all input flags exist and their status is True, the node initialization is complete and the initialization taint is removed.</p> <ul style="list-style-type: none"> • Use only letters and digits. Max. characters: 20. • Max. flags: 2.
extendParam	NodeExtendParam object	Extended parameters for creating a node.

Table 4-252 Login

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

Table 4-253 UserPassword

Parameter	Type	Description
username	String	Login account. Defaults to root .

Parameter	Type	Description
password	String	<p>If a username and a 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"> • Contains 8 to 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. The password field must be salted during node creation. For details, see Adding a Salt in the password Field When Creating a Node.

Table 4-254 Volume

Parameter	Type	Description
size	Integer	<p>Disk size, in GB.</p> <ul style="list-style-type: none"> • System disk: 40 to 1024 -Value range for data disks: 100 to 32768
volumetype	String	<p>Disk type. For details about possible values, see the description of the root_volume parameter in the API used to create an ECS in the ECS API reference.</p> <ul style="list-style-type: none"> • SAS: high I/O disk • SSD: ultra-high I/O disk • SATA: common I/O disk SATA disks have been removed from EVS. You can find them attached only on existing nodes.
extendParam	Map<String, Object>	<p>Extended disk parameters, defined in extendparam 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 dssPoolID, 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 dss. This field is used only for DeC clusters.</p>

Parameter	Type	Description
hw:passthrough	Boolean	<ul style="list-style-type: none"> Pay attention to this field if your ECS is SDI-compliant. If the value of this field is true, an SCSI disk will be created. If the node pool type is ElasticBMS, this field must be set to true. If a node specification involves local disks and EVS disks at the same time, set the disk initialization parameters. For details, see Attaching Disks to a Node.
metadata	VolumeMetadata 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 4-255 VolumeMetadata

Parameter	Type	Description
__system__encrypted	String	Whether the EVS disk is encrypted. The value 0 indicates that the EVS disk is not encrypted, and the value 1 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 metadata . This field is used with __system__encrypted .

Table 4-256 Storage

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

Table 4-257 StorageSelectors

Parameter	Type	Description
name	String	Selector name, used as the index of selectorNames in storageGroup . Therefore, the name of each selector must be unique.
storageType	String	Specifies the storage type. Currently, only evs (EVS volumes) and local (local volumes) are supported. The local storage does not support disk selection. All local disks will form a VG. Therefore, only one storageSelector of the local type is allowed.
matchLabels	matchLabels object	Matching field of an EVS volume. The size , volumeType , metadataEncrypted , metadataCmkid and count fields are supported.

Table 4-258 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. Currently, SSD, GPSSD and SAS are supported.
metadataEncrypted	String	Disk encryption identifier. 0 indicates that the disk is not encrypted, and 1 indicates that the disk is encrypted.
metadataCmkid	String	Customer master key ID of an encrypted disk. The value is a 36-byte string.
count	String	Number of disks to be selected. If this parameter is left blank, all disks of this type are selected.

Table 4-259 StorageGroups

Parameter	Type	Description
name	String	Name of a virtual storage group, which must be unique.

Parameter	Type	Description
cceManaged	Boolean	Storage space for Kubernetes and runtime components. Only one group can be set to true . If this parameter is left blank, the default value false is used.
selectorNames	Array of strings	This parameter corresponds to name in storageSelectors . A group can match multiple selectors, but a selector can match only one group.
virtualSpaces	Array of VirtualSpace objects	Detailed management of space configuration in a group.

Table 4-260 VirtualSpace

Parameter	Type	Description
name	String	Name of a virtualSpace. <ul style="list-style-type: none"> • Kubernetes: Kubernetes space configuration. lvmConfig needs to be configured. • runtime: runtime space configuration. runtimeConfig needs to be configured. • user: user space configuration. lvmConfig needs to be configured.
size	String	Size of a virtualSpace. The value must be an integer in percentage. Example: 90%. NOTE The sum of the percentages of all virtualSpaces in a group cannot exceed 100%.
lvmConfig	LVMConfig object	LVM configurations, applicable to kubernetes and user spaces. Note that one virtual space supports only one config.
runtimeConfig	RuntimeConfig object	runtime configurations, applicable to the runtime space. Note that one virtual space supports only one config.

Table 4-261 LVMConfig

Parameter	Type	Description
lvType	String	LVM write mode. linear indicates the linear mode. striped indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.
path	String	Path to which the disk is attached. This parameter takes effect only in user configuration. The value is an absolute path. Digits, letters, periods (.), hyphens (-), and underscores (_) are allowed.

Table 4-262 RuntimeConfig

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

Table 4-263 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. NOTE If ids has been set, you do not need to set count and eip .
count	Integer	Number of EIPs to be dynamically created. NOTE count and eip must be set at the same time.
eip	NodeEIPSpec object	EIP configuration.

Table 4-264 NodeEIPSpec

Parameter	Type	Description
iptype	String	EIP type, specified in publicip.type in the API for assigning an EIP. . .

Parameter	Type	Description
bandwidth	NodeBandwidth object	Bandwidth parameters of the EIP

Table 4-265 NodeBandwidth

Parameter	Type	Description
chargemode	String	<p>Bandwidth billing mode.</p> <ul style="list-style-type: none"> If this field is not specified, the billing is based on bandwidth. - If the field is null, the billing is based on bandwidth. If the field value is traffic, the billing is based on traffic. If the value is out of the preceding options, the cloud server will fail to be created. <p>NOTE</p> <ul style="list-style-type: none"> 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%. 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%.
size	Integer	Bandwidth size, specified in bandwidth.size in the API for assigning an EIP. . .
sharetype	String	Bandwidth sharing type. Value options: PER (exclusive bandwidth)

Table 4-266 NodeNicSpec

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

Table 4-267 NicSpec

Parameter	Type	Description
subnetId	String	ID of the subnet to which the NIC belongs. If subnetId is not specified when creating the primary NIC, the cluster subnet is used. When creating a secondary NIC, you must specify subnetId .
fixedIps	Array of strings	The IP address of the primary NIC is specified by fixedIps . The number of IP addresses cannot be greater than the number of created nodes. fixedIps and ipBlock cannot be specified at the same time.
ipBlock	String	CIDR format of the primary NIC IP range. The IP address of the created node falls in this range. fixedIps and ipBlock cannot be specified at the same time.

Table 4-268 Taint

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

Table 4-269 UserTag

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

Table 4-270 Runtime

Parameter	Type	Description
name	String	Container runtime. Defaults to docker .

Table 4-271 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 Maximum Number of Pods That Can Be Created on a Node.</p>
DockerLVMDockerConfigOverride	String	<p>Docker data disk configuration item. (This parameter has been discarded. Use the storage field instead.) Example default configuration: "DockerLVMDockerConfigOverride":"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 diskType based on the actual drive letter type. The following fields are included:</p> <ul style="list-style-type: none"> • userLV (optional): size of the user space, for example, vgpaas/20%VG. • userPath (optional): mount path of the user space, for example, /home/wqt-test. • diskType: disk type. Currently, only evs, hdd, and ssd are supported. • lvType: type of a logic volume. The value can be linear or striped. • dockerThinpool: Docker space size, for example, vgpaas/60%VG. • kubernetesLV: kubelet space size, for example, vgpaas/20%VG.

Parameter	Type	Description
dockerBaseSize	Integer	<p>Available disk space of a single container on a node, in GB.</p> <p>If this parameter is left blank or is set to 0, the default value is used. In Device Mapper mode, the default value is 10. In OverlayFS mode, the available space of a single container is not limited by default, and the dockerBaseSize setting takes effect only on EulerOS nodes in the cluster of the new version.</p> <p>For details about how to allocate the space for the container runtime, see Data Disk Space Allocation.</p> <p>When Device Mapper is used, you are advised to set dockerBaseSize 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>
publicKey	String	Public key of a node.
alpha.cce/preInstall	String	<p>Pre-installation script.</p> <p>NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)</p>
alpha.cce/postInstall	String	<p>Post-installation script.</p> <p>NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)</p>
alpha.cce/NodeImageID	String	This parameter is required when a custom image is used to create a BMS node.
enterprise_project_id	String	ID of the enterprise project to which the node belongs
chargingMode	Integer	Billing mode of a node. This parameter has been deprecated. Use the billingMode parameter in NodeSpec .
agency_name	String	<p>Name of an agency</p> <p>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.</p>

Table 4-272 NodeStatus

Parameter	Type	Description
phase	String	Node status
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 Sync Node Data on the console to manually update the data.
deleteStatus	DeleteStatus object	Resource status during resource deletion.

Table 4-273 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.
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 4-274 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
node_id	Yes	String	Node ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .

Request Parameters

Table 4-275 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 Obtaining a User Token .

Table 4-276 Request body parameters

Parameter	Mandatory	Type	Description
metadata	Yes	ClusterNodeInformationMetadata object	Node metadata, which is a collection of attributes.

Table 4-277 ClusterNodeInformationMetadata

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

Response Parameters

Status code: 200

Table 4-278 Response body parameters

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

Table 4-279 NodeMetadata

Parameter	Type	Description
name	String	Node name. NOTE Enter 1 to 56 characters starting with a letter and not ending with a hyphen (-). Only lowercase letters, digits, and hyphens (-) are allowed. If name is not specified or empty, the node name is generated based on the default rule. The default rule is <i>Cluster name-Random characters</i> . Only the first 36 characters are kept for a too long cluster name. If the number of nodes (count) is greater than 1, some random characters are added to the end of the entered node name according to the default rule. The default rule is <i>User-defined name-Random characters</i> . Only the first 50 characters in the user-defined name are kept, followed by some random characters.
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>
annotations	Map<String,String>	CCE node annotations, in key-value pairs (not the native Kubernetes annotations). Example: <pre>"annotations": { "key1" : "value1", "key2" : "value2" }</pre> NOTE annotations are not used to identify or select objects. The metadata in annotations may be small or large, structured or unstructured, and may include characters that are not allowed in labels.
creationTimestamp	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.

Table 4-280 NodeSpec

Parameter	Type	Description
flavor	String	Node specifications. For details about the node specifications supported by CCE, see Node Flavor Description .
az	String	Name of the AZ where the node to be created is located. . .
os	String	Node OS. For details about the supported OSs, see Node OS Description . NOTE <ul style="list-style-type: none"> 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. If <code>alpha.cce/NodeImageID</code> in <code>extendParam</code> is specified during node creation, you do not need to set this field.
login	Login object	Node login mode. Either the key pair or password must be used for login.
rootVolume	Volume object	Information about disks on the node
dataVolumes	Array of Volume objects	Data disk parameters of the node. Currently, you can add the second data disk for your node on the CCE console. This data disk is used by the container runtime and kubelet. Do not uninstall this disk. Otherwise, the node will become unavailable. For DeC nodes, the parameter description is the same as that for rootVolume .
storage	Storage object	Disk initialization management parameter. This parameter is complex to configure. For details, see Attaching Disks to a Node . If this parameter retains its default, disks are managed based on the <code>DockerLVMConfigOverride (discarded)</code> parameter in <code>extendParam</code> . This parameter is supported by clusters of version 1.15.11 and later. NOTE If a node specification involves local disks and EVS disks at the same time, do not retain the default value of this parameter to prevent unexpected disk partitions.
publicIP	NodePublicIP object	EIP of a node. NOTE This parameter is not supported when you add a node to a node pool.

Parameter	Type	Description
nodeNicSpec	NodeNicSpec 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"> • 0: pay-per-use
taints	Array of Taint 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"> • Key: A key must contain 1 to 63 characters starting with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain name can be used as the prefix of a key. • Value: A value must start with a letter or digit and can contain a maximum of 63 characters, including letters, digits, hyphens (-), underscores (_), and periods (.). • Effect: Available options are NoSchedule, PreferNoSchedule, and NoExecute. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In using 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>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"> • Key: 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 and contain a maximum of 253 characters. Example DNS subdomain: example.com/my-key • Value: The value can be left blank or contain 1 to 63 characters that start with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed in the character string. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In using 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>Example: <pre>"k8sTags": { "key": "value" }</pre> </p>
ecsGroupId	String	Cloud server group ID. If this field is specified, the node is created in the specified cloud server group.
dedicatedHostId	String	<p>ID of the DeH to which the node is scheduled.</p> <p>NOTE This parameter is not supported when you add a node during node pool creation.</p>
userTags	Array of UserTag objects	Cloud server tag. The key of a tag must be unique. The maximum number of custom tags supported by CCE depends on the region 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 using 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.
runtime	Runtime object	Container runtime. Defaults to docker .

Parameter	Type	Description
initializedConditions	Array of strings	<p>Custom initialization flag.</p> <p>Before CCE nodes are initialized, they are tainted with node.cloudprovider.kubernetes.io/uninitialized to prevent pods from being scheduled to them.</p> <p>CCE supports custom initialization flags. After receiving the initializedConditions parameter, CCE converts the parameter value into a node label and provisions the label with the node, for example, cloudprovider.openvessel.io/inject-initialized-conditions=CCEInitial_CustomedInitial.</p> <p>After the node is labeled, its status.Conditions is polled to check whether the type of conditions has a flag name, such as CCEInitial and CustomedInitial. If all input flags exist and their status is True, the node initialization is complete and the initialization taint is removed.</p> <ul style="list-style-type: none"> • Use only letters and digits. Max. characters: 20. • Max. flags: 2.
extendParam	NodeExtendParam object	Extended parameters for creating a node.

Table 4-281 Login

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

Table 4-282 UserPassword

Parameter	Type	Description
username	String	Login account. Defaults to root .

Parameter	Type	Description
password	String	<p>If a username and a 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"> • Contains 8 to 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. The password field must be salted during node creation. For details, see Adding a Salt in the password Field When Creating a Node.

Table 4-283 Volume

Parameter	Type	Description
size	Integer	<p>Disk size, in GB.</p> <ul style="list-style-type: none"> • System disk: 40 to 1024 -Value range for data disks: 100 to 32768
volumetype	String	<p>Disk type. For details about possible values, see the description of the root_volume parameter in the API used to create an ECS in the ECS API reference.</p> <ul style="list-style-type: none"> • SAS: high I/O disk • SSD: ultra-high I/O disk • SATA: common I/O disk SATA disks have been removed from EVS. You can find them attached only on existing nodes.
extendParam	Map<String, Object>	<p>Extended disk parameters, defined in extendparam 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 dssPoolID, 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 dss. This field is used only for DeC clusters.</p>

Parameter	Type	Description
hw:passthrough	Boolean	<ul style="list-style-type: none"> Pay attention to this field if your ECS is SDI-compliant. If the value of this field is true, an SCSI disk will be created. If the node pool type is ElasticBMS, this field must be set to true. If a node specification involves local disks and EVS disks at the same time, set the disk initialization parameters. For details, see Attaching Disks to a Node.
metadata	VolumeMetadata 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 4-284 VolumeMetadata

Parameter	Type	Description
__system__encrypted	String	Whether the EVS disk is encrypted. The value 0 indicates that the EVS disk is not encrypted, and the value 1 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 metadata . This field is used with __system__encrypted .

Table 4-285 Storage

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

Table 4-286 StorageSelectors

Parameter	Type	Description
name	String	Selector name, used as the index of selectorNames in storageGroup . Therefore, the name of each selector must be unique.
storageType	String	Specifies the storage type. Currently, only evs (EVS volumes) and local (local volumes) are supported. The local storage does not support disk selection. All local disks will form a VG. Therefore, only one storageSelector of the local type is allowed.
matchLabels	matchLabels object	Matching field of an EVS volume. The size , volumeType , metadataEncrypted , metadataCmkid and count fields are supported.

Table 4-287 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. Currently, SSD, GPSSD and SAS are supported.
metadataEncrypted	String	Disk encryption identifier. 0 indicates that the disk is not encrypted, and 1 indicates that the disk is encrypted.
metadataCmkid	String	Customer master key ID of an encrypted disk. The value is a 36-byte string.
count	String	Number of disks to be selected. If this parameter is left blank, all disks of this type are selected.

Table 4-288 StorageGroups

Parameter	Type	Description
name	String	Name of a virtual storage group, which must be unique.

Parameter	Type	Description
cceManaged	Boolean	Storage space for Kubernetes and runtime components. Only one group can be set to true . If this parameter is left blank, the default value false is used.
selectorNames	Array of strings	This parameter corresponds to name in storageSelectors . A group can match multiple selectors, but a selector can match only one group.
virtualSpaces	Array of VirtualSpace objects	Detailed management of space configuration in a group.

Table 4-289 VirtualSpace

Parameter	Type	Description
name	String	Name of a virtualSpace. <ul style="list-style-type: none"> • Kubernetes: Kubernetes space configuration. lvmConfig needs to be configured. • runtime: runtime space configuration. runtimeConfig needs to be configured. • user: user space configuration. lvmConfig needs to be configured.
size	String	Size of a virtualSpace. The value must be an integer in percentage. Example: 90%. NOTE The sum of the percentages of all virtualSpaces in a group cannot exceed 100%.
lvmConfig	LVMConfig object	LVM configurations, applicable to kubernetes and user spaces. Note that one virtual space supports only one config.
runtimeConfig	RuntimeConfig object	runtime configurations, applicable to the runtime space. Note that one virtual space supports only one config.

Table 4-290 LVMConfig

Parameter	Type	Description
lvType	String	LVM write mode. linear indicates the linear mode. striped indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.
path	String	Path to which the disk is attached. This parameter takes effect only in user configuration. The value is an absolute path. Digits, letters, periods (.), hyphens (-), and underscores (_) are allowed.

Table 4-291 RuntimeConfig

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

Table 4-292 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. NOTE If ids has been set, you do not need to set count and eip .
count	Integer	Number of EIPs to be dynamically created. NOTE count and eip must be set at the same time.
eip	NodeEIPSpec object	EIP configuration.

Table 4-293 NodeEIPSpec

Parameter	Type	Description
iptype	String	EIP type, specified in publicip.type in the API for assigning an EIP. . .

Parameter	Type	Description
bandwidth	NodeBandwidth object	Bandwidth parameters of the EIP

Table 4-294 NodeBandwidth

Parameter	Type	Description
chargemode	String	<p>Bandwidth billing mode.</p> <ul style="list-style-type: none"> If this field is not specified, the billing is based on bandwidth. - If the field is null, the billing is based on bandwidth. If the field value is traffic, the billing is based on traffic. If the value is out of the preceding options, the cloud server will fail to be created. <p>NOTE</p> <ul style="list-style-type: none"> 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%. 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%.
size	Integer	Bandwidth size, specified in bandwidth.size in the API for assigning an EIP. . .
sharetype	String	Bandwidth sharing type. Value options: PER (exclusive bandwidth)

Table 4-295 NodeNicSpec

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

Table 4-296 NicSpec

Parameter	Type	Description
subnetId	String	ID of the subnet to which the NIC belongs. If subnetId is not specified when creating the primary NIC, the cluster subnet is used. When creating a secondary NIC, you must specify subnetId .
fixedIps	Array of strings	The IP address of the primary NIC is specified by fixedIps . The number of IP addresses cannot be greater than the number of created nodes. fixedIps and ipBlock cannot be specified at the same time.
ipBlock	String	CIDR format of the primary NIC IP range. The IP address of the created node falls in this range. fixedIps and ipBlock cannot be specified at the same time.

Table 4-297 Taint

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

Table 4-298 UserTag

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

Table 4-299 Runtime

Parameter	Type	Description
name	String	Container runtime. Defaults to docker .

Table 4-300 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 Maximum Number of Pods That Can Be Created on a Node.</p>
DockerLVMDockerConfigOverride	String	<p>Docker data disk configuration item. (This parameter has been discarded. Use the storage field instead.) Example default configuration: <code>"DockerLVMDockerConfigOverride": "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 diskType based on the actual drive letter type. The following fields are included:</p> <ul style="list-style-type: none"> • userLV (optional): size of the user space, for example, vgpaas/20%VG. • userPath (optional): mount path of the user space, for example, /home/wqt-test. • diskType: disk type. Currently, only evs, hdd, and ssd are supported. • lvType: type of a logic volume. The value can be linear or striped. • dockerThinpool: Docker space size, for example, vgpaas/60%VG. • kubernetesLV: kubelet space size, for example, vgpaas/20%VG.

Parameter	Type	Description
dockerBaseSize	Integer	<p>Available disk space of a single container on a node, in GB.</p> <p>If this parameter is left blank or is set to 0, the default value is used. In Device Mapper mode, the default value is 10. In OverlayFS mode, the available space of a single container is not limited by default, and the dockerBaseSize setting takes effect only on EulerOS nodes in the cluster of the new version.</p> <p>For details about how to allocate the space for the container runtime, see Data Disk Space Allocation.</p> <p>When Device Mapper is used, you are advised to set dockerBaseSize 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>
publicKey	String	Public key of a node.
alpha.cce/preInstall	String	<p>Pre-installation script.</p> <p>NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)</p>
alpha.cce/postInstall	String	<p>Post-installation script.</p> <p>NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)</p>
alpha.cce/NodeImageID	String	This parameter is required when a custom image is used to create a BMS node.
enterprise_project_id	String	ID of the enterprise project to which the node belongs
chargingMode	Integer	Billing mode of a node. This parameter has been deprecated. Use the billingMode parameter in NodeSpec .
agency_name	String	<p>Name of an agency</p> <p>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.</p>

Table 4-301 NodeStatus

Parameter	Type	Description
phase	String	Node status
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 Sync Node Data on the console to manually update the data.
deleteStatus	DeleteStatus object	Resource status during resource deletion.

Table 4-302 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.
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

```
{
  "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",
      "diskSize": 40
    },
    "dataVolumes": [ {
      "volumeType": "SAS",
      "diskSize": 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 4-303 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
node_id	Yes	String	Node ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .

Table 4-304 Query Parameters

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

Request Parameters

Table 4-305 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 Obtaining a User Token .

Response Parameters

Status code: 200

Table 4-306 Response body parameters

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

Table 4-307 NodeMetadata

Parameter	Type	Description
name	String	Node name. NOTE Enter 1 to 56 characters starting with a letter and not ending with a hyphen (-). Only lowercase letters, digits, and hyphens (-) are allowed. If name is not specified or empty, the node name is generated based on the default rule. The default rule is <i>Cluster name-Random characters</i> . Only the first 36 characters are kept for a too long cluster name. If the number of nodes (count) is greater than 1, some random characters are added to the end of the entered node name according to the default rule. The default rule is <i>User-defined name-Random characters</i> . Only the first 50 characters in the user-defined name are kept, followed by some random characters.
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>
annotations	Map<String,String>	CCE node annotations, in key-value pairs (not the native Kubernetes annotations). Example: <pre>"annotations": { "key1" : "value1", "key2" : "value2" }</pre> NOTE annotations are not used to identify or select objects. The metadata in annotations may be small or large, structured or unstructured, and may include characters that are not allowed in labels.
creationTimestamp	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.

Table 4-308 NodeSpec

Parameter	Type	Description
flavor	String	Node specifications. For details about the node specifications supported by CCE, see Node Flavor Description .
az	String	Name of the AZ where the node to be created is located. . .
os	String	Node OS. For details about the supported OSs, see Node OS Description . NOTE <ul style="list-style-type: none"> 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. If <code>alpha.cce/NodeImageID</code> in <code>extendParam</code> is specified during node creation, you do not need to set this field.
login	Login object	Node login mode. Either the key pair or password must be used for login.
rootVolume	Volume object	Information about disks on the node
dataVolumes	Array of Volume objects	Data disk parameters of the node. Currently, you can add the second data disk for your node on the CCE console. This data disk is used by the container runtime and kubelet. Do not uninstall this disk. Otherwise, the node will become unavailable. For DeC nodes, the parameter description is the same as that for rootVolume .
storage	Storage object	Disk initialization management parameter. This parameter is complex to configure. For details, see Attaching Disks to a Node . If this parameter retains its default, disks are managed based on the <code>DockerLVMConfigOverride</code> (discarded) parameter in <code>extendParam</code> . This parameter is supported by clusters of version 1.15.11 and later. NOTE If a node specification involves local disks and EVS disks at the same time, do not retain the default value of this parameter to prevent unexpected disk partitions.
publicIP	NodePublicIP object	EIP of a node. NOTE This parameter is not supported when you add a node to a node pool.

Parameter	Type	Description
nodeNicSpec	NodeNicSpec 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"> 0: pay-per-use
taints	Array of Taint 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"> Key: A key must contain 1 to 63 characters starting with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain name can be used as the prefix of a key. Value: A value must start with a letter or digit and can contain a maximum of 63 characters, including letters, digits, hyphens (-), underscores (_), and periods (.). Effect: Available options are NoSchedule, PreferNoSchedule, and NoExecute. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In using 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>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"> • Key: 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 and contain a maximum of 253 characters. Example DNS subdomain: example.com/my-key • Value: The value can be left blank or contain 1 to 63 characters that start with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed in the character string. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In using 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>Example:</p> <pre>"k8sTags": { "key": "value" }</pre>
ecsGroupId	String	Cloud server group ID. If this field is specified, the node is created in the specified cloud server group.
dedicatedHostId	String	<p>ID of the DeH to which the node is scheduled.</p> <p>NOTE This parameter is not supported when you add a node during node pool creation.</p>
userTags	Array of UserTag objects	Cloud server tag. The key of a tag must be unique. The maximum number of custom tags supported by CCE depends on the region 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 using 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.
runtime	Runtime object	Container runtime. Defaults to docker .

Parameter	Type	Description
initializedConditions	Array of strings	<p>Custom initialization flag.</p> <p>Before CCE nodes are initialized, they are tainted with node.cloudprovider.kubernetes.io/uninitialized to prevent pods from being scheduled to them.</p> <p>CCE supports custom initialization flags. After receiving the initializedConditions parameter, CCE converts the parameter value into a node label and provisions the label with the node, for example, cloudprovider.openvessel.io/inject-initialized-conditions=CCEInitial_CustomedInitial.</p> <p>After the node is labeled, its status.Conditions is polled to check whether the type of conditions has a flag name, such as CCEInitial and CustomedInitial. If all input flags exist and their status is True, the node initialization is complete and the initialization taint is removed.</p> <ul style="list-style-type: none"> • Use only letters and digits. Max. characters: 20. • Max. flags: 2.
extendParam	NodeExtendParam object	Extended parameters for creating a node.

Table 4-309 Login

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

Table 4-310 UserPassword

Parameter	Type	Description
username	String	Login account. Defaults to root .

Parameter	Type	Description
password	String	<p>If a username and a 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"> • Contains 8 to 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. The password field must be salted during node creation. For details, see Adding a Salt in the password Field When Creating a Node.

Table 4-311 Volume

Parameter	Type	Description
size	Integer	<p>Disk size, in GB.</p> <ul style="list-style-type: none"> • System disk: 40 to 1024 -Value range for data disks: 100 to 32768
volumetype	String	<p>Disk type. For details about possible values, see the description of the root_volume parameter in the API used to create an ECS in the ECS API reference.</p> <ul style="list-style-type: none"> • SAS: high I/O disk • SSD: ultra-high I/O disk • SATA: common I/O disk SATA disks have been removed from EVS. You can find them attached only on existing nodes.
extendParam	Map<String, Object>	<p>Extended disk parameters, defined in extendparam 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 dssPoolID, 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 dss. This field is used only for DeC clusters.</p>

Parameter	Type	Description
hw:passthrough	Boolean	<ul style="list-style-type: none"> Pay attention to this field if your ECS is SDI-compliant. If the value of this field is true, an SCSI disk will be created. If the node pool type is ElasticBMS, this field must be set to true. If a node specification involves local disks and EVS disks at the same time, set the disk initialization parameters. For details, see Attaching Disks to a Node.
metadata	VolumeMetadata 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 4-312 VolumeMetadata

Parameter	Type	Description
__system__encrypted	String	Whether the EVS disk is encrypted. The value 0 indicates that the EVS disk is not encrypted, and the value 1 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 metadata . This field is used with __system__encrypted .

Table 4-313 Storage

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

Table 4-314 StorageSelectors

Parameter	Type	Description
name	String	Selector name, used as the index of selectorNames in storageGroup . Therefore, the name of each selector must be unique.
storageType	String	Specifies the storage type. Currently, only evs (EVS volumes) and local (local volumes) are supported. The local storage does not support disk selection. All local disks will form a VG. Therefore, only one storageSelector of the local type is allowed.
matchLabels	matchLabels object	Matching field of an EVS volume. The size , volumeType , metadataEncrypted , metadataCmkid and count fields are supported.

Table 4-315 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. Currently, SSD, GPSSD and SAS are supported.
metadataEncrypted	String	Disk encryption identifier. 0 indicates that the disk is not encrypted, and 1 indicates that the disk is encrypted.
metadataCmkid	String	Customer master key ID of an encrypted disk. The value is a 36-byte string.
count	String	Number of disks to be selected. If this parameter is left blank, all disks of this type are selected.

Table 4-316 StorageGroups

Parameter	Type	Description
name	String	Name of a virtual storage group, which must be unique.

Parameter	Type	Description
cceManaged	Boolean	Storage space for Kubernetes and runtime components. Only one group can be set to true . If this parameter is left blank, the default value false is used.
selectorNames	Array of strings	This parameter corresponds to name in storageSelectors . A group can match multiple selectors, but a selector can match only one group.
virtualSpaces	Array of VirtualSpace objects	Detailed management of space configuration in a group.

Table 4-317 VirtualSpace

Parameter	Type	Description
name	String	Name of a virtualSpace. <ul style="list-style-type: none"> • Kubernetes: Kubernetes space configuration. lvmConfig needs to be configured. • runtime: runtime space configuration. runtimeConfig needs to be configured. • user: user space configuration. lvmConfig needs to be configured.
size	String	Size of a virtualSpace. The value must be an integer in percentage. Example: 90%. NOTE The sum of the percentages of all virtualSpaces in a group cannot exceed 100%.
lvmConfig	LVMConfig object	LVM configurations, applicable to kubernetes and user spaces. Note that one virtual space supports only one config.
runtimeConfig	RuntimeConfig object	runtime configurations, applicable to the runtime space. Note that one virtual space supports only one config.

Table 4-318 LVMConfig

Parameter	Type	Description
lvType	String	LVM write mode. linear indicates the linear mode. striped indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.
path	String	Path to which the disk is attached. This parameter takes effect only in user configuration. The value is an absolute path. Digits, letters, periods (.), hyphens (-), and underscores (_) are allowed.

Table 4-319 RuntimeConfig

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

Table 4-320 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. NOTE If ids has been set, you do not need to set count and eip .
count	Integer	Number of EIPs to be dynamically created. NOTE count and eip must be set at the same time.
eip	NodeEIPSpec object	EIP configuration.

Table 4-321 NodeEIPSpec

Parameter	Type	Description
iptype	String	EIP type, specified in publicip.type in the API for assigning an EIP. . .

Parameter	Type	Description
bandwidth	NodeBandwidth object	Bandwidth parameters of the EIP

Table 4-322 NodeBandwidth

Parameter	Type	Description
chargemode	String	<p>Bandwidth billing mode.</p> <ul style="list-style-type: none"> If this field is not specified, the billing is based on bandwidth. - If the field is null, the billing is based on bandwidth. If the field value is traffic, the billing is based on traffic. If the value is out of the preceding options, the cloud server will fail to be created. <p>NOTE</p> <ul style="list-style-type: none"> 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%. 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%.
size	Integer	Bandwidth size, specified in bandwidth.size in the API for assigning an EIP. . .
sharetype	String	Bandwidth sharing type. Value options: PER (exclusive bandwidth)

Table 4-323 NodeNicSpec

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

Table 4-324 NicSpec

Parameter	Type	Description
subnetId	String	ID of the subnet to which the NIC belongs. If subnetId is not specified when creating the primary NIC, the cluster subnet is used. When creating a secondary NIC, you must specify subnetId .
fixedIps	Array of strings	The IP address of the primary NIC is specified by fixedIps . The number of IP addresses cannot be greater than the number of created nodes. fixedIps and ipBlock cannot be specified at the same time.
ipBlock	String	CIDR format of the primary NIC IP range. The IP address of the created node falls in this range. fixedIps and ipBlock cannot be specified at the same time.

Table 4-325 Taint

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

Table 4-326 UserTag

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

Table 4-327 Runtime

Parameter	Type	Description
name	String	Container runtime. Defaults to docker .

Table 4-328 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 Maximum Number of Pods That Can Be Created on a Node.</p>
DockerLVMDockerConfigOverride	String	<p>Docker data disk configuration item. (This parameter has been discarded. Use the storage field instead.) Example default configuration: <code>"DockerLVMDockerConfigOverride": "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 diskType based on the actual drive letter type. The following fields are included:</p> <ul style="list-style-type: none"> • userLV (optional): size of the user space, for example, vgpaas/20%VG. • userPath (optional): mount path of the user space, for example, /home/wqt-test. • diskType: disk type. Currently, only evs, hdd, and ssd are supported. • lvType: type of a logic volume. The value can be linear or striped. • dockerThinpool: Docker space size, for example, vgpaas/60%VG. • kubernetesLV: kubelet space size, for example, vgpaas/20%VG.

Parameter	Type	Description
dockerBaseSize	Integer	<p>Available disk space of a single container on a node, in GB.</p> <p>If this parameter is left blank or is set to 0, the default value is used. In Device Mapper mode, the default value is 10. In OverlayFS mode, the available space of a single container is not limited by default, and the dockerBaseSize setting takes effect only on EulerOS nodes in the cluster of the new version.</p> <p>For details about how to allocate the space for the container runtime, see Data Disk Space Allocation.</p> <p>When Device Mapper is used, you are advised to set dockerBaseSize 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>
publicKey	String	Public key of a node.
alpha.cce/preInstall	String	<p>Pre-installation script.</p> <p>NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)</p>
alpha.cce/postInstall	String	<p>Post-installation script.</p> <p>NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)</p>
alpha.cce/NodeImageID	String	This parameter is required when a custom image is used to create a BMS node.
enterprise_project_id	String	ID of the enterprise project to which the node belongs
chargingMode	Integer	Billing mode of a node. This parameter has been deprecated. Use the billingMode parameter in NodeSpec .
agency_name	String	<p>Name of an agency</p> <p>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.</p>

Table 4-329 NodeStatus

Parameter	Type	Description
phase	String	Node status
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 Sync Node Data on the console to manually update the data.
deleteStatus	DeleteStatus object	Resource status during resource deletion.

Table 4-330 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.
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" : "Active",
  "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 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.

URI

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

Table 4-331 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .

Request Parameters

Table 4-332 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 Obtaining a User Token .

Table 4-333 Request body parameters

Parameter	Mandatory	Type	Description
apiVersion	Yes	String	API version. The value is fixed at v3 .
kind	Yes	String	API type. The value is fixed at List .
nodeList	Yes	Array of AddNode objects	List of the nodes to be accepted.

Table 4-334 AddNode

Parameter	Mandatory	Type	Description
serverID	Yes	String	Server ID. For details about how to obtain the server ID, see the ECS or BMS documentation.
spec	Yes	ReinstallNodeSpec object	Node reinstallation configuration parameters. Currently, accepted nodes cannot be added into node pools.

Table 4-335 ReinstallNodeSpec

Parameter	Mandatory	Type	Description
os	Yes	String	Operating system. If you specify a custom image, the actual OS version in the IMS image is used. Select an OS version supported by the current cluster, for example, EulerOS 2.5, CentOS 7.6, or EulerOS 2.8.
login	Yes	Login object	Node login mode. Either the key pair or password must be used for login.
name	No	String	Node name. NOTE 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 letter and not ending with a hyphen (-).
serverConfig	No	ReinstallServerConfig object	Server configuration.
volumeConfig	No	ReinstallVolumeConfig object	Volume management configuration.

Parameter	Mandatory	Type	Description
runtimeConfig	No	ReinstallRuntimeConfig object	Container runtime configuration.
k8sOptions	No	ReinstallK8sOptionsConfig object	Kubernetes node configuration.
lifecycle	No	NodeLifecycleConfig object	Customized lifecycle configuration of a node.
initializedConditions	No	Array of strings	<p>Custom initialization flag. Before CCE nodes are initialized, they are tainted with node.cloudprovider.kubernetes.io/uninitialized to prevent pods from being scheduled to them.</p> <p>CCE supports custom initialization flags. After receiving the initializedConditions parameter, CCE converts the parameter value into a node label and provisions the label with the node, for example, cloudprovider.openvessel.io/inject-initialized-conditions=CCEInitial_CustomedInitial.</p> <p>After the node is labeled, its status.Conditions is polled to check whether the type of conditions has a flag name, such as CCEInitial and CustomedInitial. If all input flags exist and their status is True, the node initialization is complete and the initialization taint is removed.</p> <ul style="list-style-type: none"> • Use only letters and digits. Max. characters: 20. • Max. flags: 2.
extendParam	No	ReinstallExtendParam object	Extended reinstallation parameter, which is discarded.

Table 4-336 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 4-337 UserPassword

Parameter	Mandatory	Type	Description
username	No	String	Login account. Defaults to root .
password	Yes	String	<p>If a username and a 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"> • Contains 8 to 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. The password field must be salted during node creation. For details, see Adding a Salt in the password Field When Creating a Node.

Table 4-338 ReinstallServerConfig

Parameter	Mandatory	Type	Description
userTags	No	Array of UserTag objects	Cloud server labels. The key of a label must be unique. The maximum number of user-defined labels supported by CCE depends on the region. In the region that supports the least number of labels, you can still create up to 5 labels for a cloud server.
rootVolume	No	ReinstallVolumeSpec object	System disk configurations used in reinstallation.

Table 4-339 UserTag

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

Table 4-340 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 4-341 ReinstallVolumeConfig

Parameter	Mandatory	Type	Description
lvmConfig	No	String	<p>Docker data disk configurations.</p> <p>The following is an example 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"> • userLV: size of the user space, for example, vgpaas/20%VG. • userPath: mount path of the user space, for example, /home/wqt-test. • diskType: disk type. Currently, only evs, hdd, and ssd are supported. • lvType: type of a logic volume. The value can be linear or striped. • dockerThinpool: Docker space size, for example, vgpaas/60%VG. • kubernetesLV: kubelet space size, for example, vgpaas/20%VG.
storage	No	Storage object	<p>Disk initialization management parameter.</p> <p>This parameter is complex to configure. For details, see Attaching Disks to a Node.</p> <p>If this parameter retains its default, disks are managed based on the <code>DockerLVMConfigOverride</code> (discarded) parameter in <code>extendParam</code>. This parameter is supported by clusters of version 1.15.11 and later.</p> <p>NOTE If a node specification involves local disks and EVS disks at the same time, do not retain the default value of this parameter to prevent unexpected disk partitions.</p>

Table 4-342 Storage

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

Table 4-343 StorageSelectors

Parameter	Mandatory	Type	Description
name	Yes	String	Selector name, used as the index of selectorNames in storageGroup . Therefore, the name of each selector must be unique.
storageType	Yes	String	Specifies the storage type. Currently, only evs (EVS volumes) and local (local volumes) are supported. The local storage does not support disk selection. All local disks will form a VG. Therefore, only one storageSelector of the local type is allowed.
matchLabels	No	matchLabels object	Matching field of an EVS volume. The size , volumeType , metadataEncrypted , metadataCmkid and count fields are supported.

Table 4-344 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

Parameter	Mandatory	Type	Description
volumeType	No	String	EVS disk type. Currently, SSD, GPSSD and SAS are supported.
metadataEncrypted	No	String	Disk encryption identifier. 0 indicates that the disk is not encrypted, and 1 indicates that the disk is encrypted.
metadataCmkid	No	String	Customer master key ID of an encrypted disk. The value is a 36-byte string.
count	No	String	Number of disks to be selected. If this parameter is left blank, all disks of this type are selected.

Table 4-345 StorageGroups

Parameter	Mandatory	Type	Description
name	Yes	String	Name of a virtual storage group, which must be unique.
cceManaged	No	Boolean	Storage space for Kubernetes and runtime components. Only one group can be set to true . If this parameter is left blank, the default value false is used.
selectorNames	Yes	Array of strings	This parameter corresponds to name in storageSelectors . A group can match multiple selectors, but a selector can match only one group.
virtualSpaces	Yes	Array of VirtualSpace objects	Detailed management of space configuration in a group.

Table 4-346 VirtualSpace

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

Table 4-347 LVMConfig

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

Parameter	Mandatory	Type	Description
path	No	String	Path to which the disk is attached. This parameter takes effect only in user configuration. The value is an absolute path. Digits, letters, periods (.), hyphens (-), and underscores (_) are allowed.

Table 4-348 RuntimeConfig

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

Table 4-349 ReinstallRuntimeConfig

Parameter	Mandatory	Type	Description
dockerBaseSize	No	Integer	<p>Available disk space of a single container on a node, in GB.</p> <p>If this parameter is left blank or is set to 0, the default value is used. In Device Mapper mode, the default value is 10. In OverlayFS mode, the available space of a single container is not limited by default, and the dockerBaseSize setting takes effect only on EulerOS nodes in the cluster of the new version.</p> <p>For details about how to allocate the space for the container runtime, see Data Disk Space Allocation.</p> <p>When Device Mapper is used, you are advised to set dockerBaseSize 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>
runtime	No	Runtime object	Container runtime. Defaults to docker .

Table 4-350 Runtime

Parameter	Mandatory	Type	Description
name	No	String	Container runtime. Defaults to docker .

Table 4-351 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"> Key: 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 and contain a maximum of 253 characters. Example DNS subdomain: example.com/my-key Value: The value can be left blank or contain 1 to 63 characters that start with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed in the character string. <p>Example:</p> <pre>"k8sTags": { "key": "value" }</pre>

Parameter	Mandatory	Type	Description
taints	No	Array of Taint objects	<p>Taints can be added for anti-affinity when creating nodes. A maximum of 20 taints can be added. Each taint contains the following parameters:</p> <ul style="list-style-type: none"> • Key: A key must contain 1 to 63 characters starting with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain name can be used as the prefix of a key. • Value: A value must start with a letter or digit and can contain a maximum of 63 characters, including letters, digits, hyphens (-), underscores (_), and periods (.). • Effect: Available options are NoSchedule, PreferNoSchedule, and NoExecute. <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 4-352 Taint

Parameter	Mandatory	Type	Description
key	Yes	String	Key.

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

Table 4-353 NodeLifecycleConfig

Parameter	Mandatory	Type	Description
preInstall	No	String	Pre-installation script. NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)
postInstall	No	String	Post-installation script. NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)

Table 4-354 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 imageID in ReinstallVolumeSpec . The original value will be overwritten.

Response Parameters

Status code: 200

Table 4-355 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

Accepting a Node

```
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](#).

4.3.7 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 4-356 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .

Request Parameters

Table 4-357 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 Obtaining a User Token .

Table 4-358 Request body parameters

Parameter	Mandatory	Type	Description
apiVersion	Yes	String	API version. The value is fixed at v3 .
kind	Yes	String	API type. The value is fixed at List .
nodeList	Yes	Array of ResetNode objects	List of the nodes to be reset.

Table 4-359 ResetNode

Parameter	Mandatory	Type	Description
nodeID	Yes	String	Node ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
spec	Yes	ReinstallNodeSpec 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 4-360 ReinstallNodeSpec

Parameter	Mandatory	Type	Description
os	Yes	String	Operating system. If you specify a custom image, the actual OS version in the IMS image is used. Select an OS version supported by the current cluster, for example, EulerOS 2.5, CentOS 7.6, or EulerOS 2.8.
login	Yes	Login object	Node login mode. Either the key pair or password must be used for login.
name	No	String	Node name. NOTE 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 letter and not ending with a hyphen (-).
serverConfig	No	ReinstallServerConfig object	Server configuration.

Parameter	Mandatory	Type	Description
volumeConfig	No	ReinstallVolumeConfig object	Volume management configuration.
runtimeConfig	No	ReinstallRuntimeConfig object	Container runtime configuration.
k8sOptions	No	ReinstallK8sOptionsConfig object	Kubernetes node configuration.
lifecycle	No	NodeLifecycleConfig object	Customized lifecycle configuration of a node.
initializedConditions	No	Array of strings	<p>Custom initialization flag. Before CCE nodes are initialized, they are tainted with node.cloudprovider.kubernetes.io/uninitialized to prevent pods from being scheduled to them.</p> <p>CCE supports custom initialization flags. After receiving the initializedConditions parameter, CCE converts the parameter value into a node label and provisions the label with the node, for example, cloudprovider.openvessel.io/inject-initialized-conditions=CCEInitial_CustomedInitial.</p> <p>After the node is labeled, its status.Conditions is polled to check whether the type of conditions has a flag name, such as CCEInitial and CustomedInitial. If all input flags exist and their status is True, the node initialization is complete and the initialization taint is removed.</p> <ul style="list-style-type: none"> • Use only letters and digits. Max. characters: 20. • Max. flags: 2.

Parameter	Mandatory	Type	Description
extendParam	No	ReinstallExtendParam object	Extended reinstallation parameter, which is discarded.

Table 4-361 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 4-362 UserPassword

Parameter	Mandatory	Type	Description
username	No	String	Login account. Defaults to root .
password	Yes	String	<p>If a username and a 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"> • Contains 8 to 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. The password field must be salted during node creation. For details, see Adding a Salt in the password Field When Creating a Node.

Table 4-363 ReinstallServerConfig

Parameter	Mandatory	Type	Description
userTags	No	Array of UserTag objects	Cloud server labels. The key of a label must be unique. The maximum number of user-defined labels supported by CCE depends on the region. In the region that supports the least number of labels, you can still create up to 5 labels for a cloud server.
rootVolume	No	ReinstallVolumeSpec object	System disk configurations used in reinstallation.

Table 4-364 UserTag

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

Table 4-365 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 4-366 ReinstallVolumeConfig

Parameter	Mandatory	Type	Description
lvmConfig	No	String	<p>Docker data disk configurations.</p> <p>The following is an example 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"> • userLV: size of the user space, for example, vgpaas/20%VG. • userPath: mount path of the user space, for example, /home/wqt-test. • diskType: disk type. Currently, only evs, hdd, and ssd are supported. • lvType: type of a logic volume. The value can be linear or striped. • dockerThinpool: Docker space size, for example, vgpaas/60%VG. • kubernetesLV: kubelet space size, for example, vgpaas/20%VG.
storage	No	Storage object	<p>Disk initialization management parameter.</p> <p>This parameter is complex to configure. For details, see Attaching Disks to a Node.</p> <p>If this parameter retains its default, disks are managed based on the <code>DockerLVMConfigOverride</code> (discarded) parameter in <code>extendParam</code>. This parameter is supported by clusters of version 1.15.11 and later.</p> <p>NOTE If a node specification involves local disks and EVS disks at the same time, do not retain the default value of this parameter to prevent unexpected disk partitions.</p>

Table 4-367 Storage

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

Table 4-368 StorageSelectors

Parameter	Mandatory	Type	Description
name	Yes	String	Selector name, used as the index of selectorNames in storageGroup . Therefore, the name of each selector must be unique.
storageType	Yes	String	Specifies the storage type. Currently, only evs (EVS volumes) and local (local volumes) are supported. The local storage does not support disk selection. All local disks will form a VG. Therefore, only one storageSelector of the local type is allowed.
matchLabels	No	matchLabels object	Matching field of an EVS volume. The size , volumeType , metadataEncrypted , metadataCmkid and count fields are supported.

Table 4-369 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

Parameter	Mandatory	Type	Description
volumeType	No	String	EVS disk type. Currently, SSD, GPSSD and SAS are supported.
metadataEncrypted	No	String	Disk encryption identifier. 0 indicates that the disk is not encrypted, and 1 indicates that the disk is encrypted.
metadataCmkid	No	String	Customer master key ID of an encrypted disk. The value is a 36-byte string.
count	No	String	Number of disks to be selected. If this parameter is left blank, all disks of this type are selected.

Table 4-370 StorageGroups

Parameter	Mandatory	Type	Description
name	Yes	String	Name of a virtual storage group, which must be unique.
cceManaged	No	Boolean	Storage space for Kubernetes and runtime components. Only one group can be set to true . If this parameter is left blank, the default value false is used.
selectorNames	Yes	Array of strings	This parameter corresponds to name in storageSelectors . A group can match multiple selectors, but a selector can match only one group.
virtualSpaces	Yes	Array of VirtualSpace objects	Detailed management of space configuration in a group.

Table 4-371 VirtualSpace

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

Table 4-372 LVMConfig

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

Parameter	Mandatory	Type	Description
path	No	String	Path to which the disk is attached. This parameter takes effect only in user configuration. The value is an absolute path. Digits, letters, periods (.), hyphens (-), and underscores (_) are allowed.

Table 4-373 RuntimeConfig

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

Table 4-374 ReinstallRuntimeConfig

Parameter	Mandatory	Type	Description
dockerBaseSize	No	Integer	<p>Available disk space of a single container on a node, in GB.</p> <p>If this parameter is left blank or is set to 0, the default value is used. In Device Mapper mode, the default value is 10. In OverlayFS mode, the available space of a single container is not limited by default, and the dockerBaseSize setting takes effect only on EulerOS nodes in the cluster of the new version.</p> <p>For details about how to allocate the space for the container runtime, see Data Disk Space Allocation.</p> <p>When Device Mapper is used, you are advised to set dockerBaseSize 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>
runtime	No	Runtime object	Container runtime. Defaults to docker .

Table 4-375 Runtime

Parameter	Mandatory	Type	Description
name	No	String	Container runtime. Defaults to docker .

Table 4-376 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"> Key: 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 and contain a maximum of 253 characters. Example DNS subdomain: example.com/my-key Value: The value can be left blank or contain 1 to 63 characters that start with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed in the character string. <p>Example:</p> <pre>"k8sTags": { "key": "value" }</pre>

Parameter	Mandatory	Type	Description
taints	No	Array of Taint objects	<p>Taints can be added for anti-affinity when creating nodes. A maximum of 20 taints can be added. Each taint contains the following parameters:</p> <ul style="list-style-type: none"> • Key: A key must contain 1 to 63 characters starting with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain name can be used as the prefix of a key. • Value: A value must start with a letter or digit and can contain a maximum of 63 characters, including letters, digits, hyphens (-), underscores (_), and periods (.). • Effect: Available options are NoSchedule, PreferNoSchedule, and NoExecute. <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 4-377 Taint

Parameter	Mandatory	Type	Description
key	Yes	String	Key.

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

Table 4-378 NodeLifecycleConfig

Parameter	Mandatory	Type	Description
preInstall	No	String	Pre-installation script. NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)
postInstall	No	String	Post-installation script. NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)

Table 4-379 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 imageID in ReinstallVolumeSpec . The original value will be overwritten.

Response Parameters

Status code: 200

Table 4-380 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

- Resetting nodes 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](#).

4.3.8 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 4-381 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .

Request Parameters

Table 4-382 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 Obtaining a User Token .

Table 4-383 Request body parameters

Parameter	Mandatory	Type	Description
apiVersion	No	String	API version. The value is fixed at v3 .
kind	No	String	API type. The value is fixed at RemoveNodesTask .
spec	Yes	RemoveNodeSpec object	Configuration information.
status	No	TaskStatus object	Job status.

Table 4-384 RemoveNodesSpec

Parameter	Mandatory	Type	Description
login	Yes	Login object	Node login mode. Either the key pair or password must be used for login.
nodes	Yes	Array of NodeItem objects	List of nodes to be operated.

Table 4-385 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 4-386 UserPassword

Parameter	Mandatory	Type	Description
username	No	String	Login account. Defaults to root .

Parameter	Mandatory	Type	Description
password	Yes	String	<p>If a username and a 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"> • Contains 8 to 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. The password field must be salted during node creation. For details, see Adding a Salt in the password Field When Creating a Node.

Table 4-387 Nodeltem

Parameter	Mandatory	Type	Description
uid	Yes	String	Node ID.

Table 4-388 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 4-389 Response body parameters

Parameter	Type	Description
apiVersion	String	API version. The value is fixed at v3 .
kind	String	API type. The value is fixed at RemoveNodesTask .
spec	RemoveNodeSpec object	Configuration information.
status	TaskStatus object	Job status.

Table 4-390 RemoveNodesSpec

Parameter	Type	Description
login	Login object	Node login mode. Either the key pair or password must be used for login.
nodes	Array of NodeItem objects	List of nodes to be operated.

Table 4-391 Login

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

Table 4-392 UserPassword

Parameter	Type	Description
username	String	Login account. Defaults to root .

Parameter	Type	Description
password	String	<p>If a username and a 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"> • Contains 8 to 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. The password field must be salted during node creation. For details, see Adding a Salt in the password Field When Creating a Node.

Table 4-393 NodeItem

Parameter	Type	Description
uid	String	Node ID.

Table 4-394 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-yyyyyyyyyyyy"
    } ]
  },
  "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](#).

4.3.9 Migrating a Node

Function

This API is used to migrate a node from a specified cluster to another cluster (both clusters must in the same VPC).

 **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

PUT /api/v3/projects/{project_id}/clusters/{cluster_id}/nodes/operation/migrateto/{target_cluster_id}

Table 4-395 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
target_cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .

Request Parameters

Table 4-396 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 Obtaining a User Token .

Table 4-397 Request body parameters

Parameter	Mandatory	Type	Description
apiVersion	No	String	API version. The value is fixed at v3 .
kind	No	String	API type. The value is fixed at MigrateNodesTask .
spec	Yes	MigrateNodeSpec object	Configuration data.
status	No	TaskStatus object	Job status.

Table 4-398 MigrateNodesSpec

Parameter	Mandatory	Type	Description
os	Yes	String	Operating system type, which must be accurate to the version number. When alpha.cce/NodeImageID is specified, the value of os must be the same as the OS of the custom image.
extendParam	No	MigrateNodeExtendParam object	Extended parameters for node migration.
login	Yes	Login object	Node login mode. Either the key pair or password must be used for login.
runtime	No	Runtime object	Container runtime
nodes	Yes	Array of NodeItem objects	List of nodes to be operated.

Table 4-399 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 configuration item. (This parameter has been discarded. Use the storage field instead.)</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>Example default configuration:</p> <pre>"DockerLVMConfigOverride":"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"> • userLV (optional): size of the user space, for example, vgpaas/20%VG. • userPath (optional): mount path of the user space, for example, /home/wqt-test. • diskType: disk type. Currently, only evs, hdd, and ssd are supported. • lvType: type of a logic volume. The value can be linear or striped. • dockerThinpool: Docker space size, for example, vgpaas/60%VG. • kubernetesLV: kubelet space size, for example, vgpaas/20%VG.
alpha.cce/preInstall	No	String	<p>Pre-installation script.</p> <p>NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)</p>

Parameter	Mandatory	Type	Description
alpha.cce/postInstall	No	String	Post-installation script. NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)
alpha.cce/NodeImageID	No	String	ID of the user image to run the target OS. When alpha.cce/NodeImageID is specified, the value of os must be the same as the OS of the custom image.

Table 4-400 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 4-401 UserPassword

Parameter	Mandatory	Type	Description
username	No	String	Login account. Defaults to root .

Parameter	Mandatory	Type	Description
password	Yes	String	<p>If a username and a 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"> • Contains 8 to 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. The password field must be salted during node creation. For details, see Adding a Salt in the password Field When Creating a Node.

Table 4-402 Runtime

Parameter	Mandatory	Type	Description
name	No	String	<p>Container runtime:</p> <ul style="list-style-type: none"> • Clusters of v1.25 or earlier: docker. • Clusters of v1.25 or later: Container runtime varies with the OS. 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.

Table 4-403 Nodeltem

Parameter	Mandatory	Type	Description
uid	Yes	String	Node ID.

Table 4-404 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 4-405 Response body parameters

Parameter	Type	Description
apiVersion	String	API version. The value is fixed at v3 .
kind	String	API type. The value is fixed at MigrateNodesTask .
spec	MigrateNodesSpec object	Configuration data.
status	TaskStatus object	Job status.

Table 4-406 MigrateNodesSpec

Parameter	Type	Description
os	String	Operating system type, which must be accurate to the version number. When alpha.cce/NodeImageID is specified, the value of os must be the same as the OS of the custom image.
extendParam	MigrateNodesExtendParam object	Extended parameters for node migration.
login	Login object	Node login mode. Either the key pair or password must be used for login.
runtime	Runtime object	Container runtime

Parameter	Type	Description
nodes	Array of NodeItem objects	List of nodes to be operated.

Table 4-407 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 configuration item. (This parameter has been discarded. Use the storage field instead.)</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>Example default configuration: <code>"DockerLVMConfigOverride":"dockerThinpool=vgpaas/90%VG;kubernetesLV=vgpaas/10%VG;diskType=evs;lvType=linear"</code></p> <p>The following fields are included:</p> <ul style="list-style-type: none"> • userLV (optional): size of the user space, for example, vgpaas/20%VG. • userPath (optional): mount path of the user space, for example, /home/wqt-test. • diskType: disk type. Currently, only evs, hdd, and ssd are supported. • lvType: type of a logic volume. The value can be linear or striped. • dockerThinpool: Docker space size, for example, vgpaas/60%VG. • kubernetesLV: kubelet space size, for example, vgpaas/20%VG.
alpha.cce/preInstall	String	<p>Pre-installation script.</p> <p>NOTE The input value must be Base64-encoded. (Command: <code>echo -n "Content to be encoded" base64</code>)</p>

Parameter	Type	Description
alpha.cce/postInstall	String	Post-installation script. NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)
alpha.cce/NodeImageID	String	ID of the user image to run the target OS. When alpha.cce/NodeImageID is specified, the value of os must be the same as the OS of the custom image.

Table 4-408 Login

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

Table 4-409 UserPassword

Parameter	Type	Description
username	String	Login account. Defaults to root .
password	String	If a username and a password are used to create a node, this field is shielded in the response body. A password must meet the following complexity requirements: <ul style="list-style-type: none"> • Contains 8 to 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. The password field must be salted during node creation. For details, see Adding a Salt in the password Field When Creating a Node.

Table 4-410 Runtime

Parameter	Type	Description
name	String	Container runtime: <ul style="list-style-type: none"> • Clusters of v1.25 or earlier: docker. • Clusters of v1.25 or later: Container runtime varies with the OS. 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.

Table 4-411 NodeItem

Parameter	Type	Description
uid	String	Node ID.

Table 4-412 TaskStatus

Parameter	Type	Description
jobID	String	Job ID, which is used by the caller to query the job progress.

Example Requests

- Migrating a Node

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"
  } ]
}
```

- Migrating a node using a specified user image ID

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-xxxxxxxxxxxxx"
  }, {
    "uid" : "yyyyyyyy-yyyy-yyyy-yyyy-yyy-yyyyyyyyy"
  } ]
}
```

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"
    },
    "nodes" : [ {
      "uid" : "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxxx"
    }, {
      "uid" : "yyyyyyyy-yyyy-yyyy-yyyy-yyy-yyyyyyyyy"
    } ]
  },
  "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 cluster of v1.21, 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**. 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}/nodepools

Table 4-413 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .

Request Parameters

Table 4-414 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 Obtaining a User Token .

Table 4-415 Request body parameters

Parameter	Mandatory	Type	Description
kind	Yes	String	API type. The value is fixed at NodePool .

Parameter	Mandatory	Type	Description
apiVersion	Yes	String	API version. The value is fixed at v3 .
metadata	Yes	NodePoolMetadata object	Metadata information of the node pool
spec	Yes	NodePoolSpec object	Node pool specifications
status	No	NodePoolStatus object	Node pool status

Table 4-416 NodePoolMetadata

Parameter	Mandatory	Type	Description
name	Yes	String	Node pool name. NOTE 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"> You cannot create node pools named DefaultPool.
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
updateTimestamp	No	String	Update time.
creationTimestamp	No	String	Creation time.

Table 4-417 NodePoolSpec

Parameter	Mandatory	Type	Description
type	No	String	Node pool type. If this parameter is left blank, the value vm is used by default. <ul style="list-style-type: none"> • vm: ECS • ElasticBMS: C6 general computing-plus BMS. An example specification is c6.22xlarge.2.physical.
nodeTemplate	Yes	NodeSpec object	Detailed parameters of the node pool template.
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	NodePoolNodeAutoscaling object	Auto scaling parameters
nodeManagement	No	NodeManagement object	Node management configuration
podSecurityGroups	No	Array of SecurityID objects	Security groups configuration
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"> • Specifying no security group ID will add the new nodes to the default security group of the worker nodes. • Specifying a valid security group ID will put new nodes in that security group. • When specifying a security group, do not modify the rules of the port on which CCE running depends.

Table 4-418 NodeSpec

Parameter	Mandatory	Type	Description
flavor	Yes	String	Node specifications. For details about the node specifications supported by CCE, see Node Flavor Description .
az	Yes	String	Name of the AZ where the node to be created is located. . .
os	No	String	Node OS. For details about the supported OSs, see Node OS Description . NOTE <ul style="list-style-type: none"> 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. If <code>alpha.cce/NodeImageID</code> in <code>extendParam</code> is specified during node creation, you do not need to set this field.
login	Yes	Login object	Node login mode. Either the key pair or password must be used for login.
rootVolume	Yes	Volume object	Information about disks on the node
dataVolumes	Yes	Array of Volume objects	Data disk parameters of the node. Currently, you can add the second data disk for your node on the CCE console. This data disk is used by the container runtime and kubelet. Do not uninstall this disk. Otherwise, the node will become unavailable. For DeC nodes, the parameter description is the same as that for <code>rootVolume</code> .

Parameter	Mandatory	Type	Description
storage	No	Storage object	<p>Disk initialization management parameter.</p> <p>This parameter is complex to configure. For details, see Attaching Disks to a Node.</p> <p>If this parameter retains its default, disks are managed based on the <code>DockerLVMConfigOverride</code> (discarded) parameter in <code>extendParam</code>. This parameter is supported by clusters of version 1.15.11 and later.</p> <p>NOTE</p> <p>If a node specification involves local disks and EVS disks at the same time, do not retain the default value of this parameter to prevent unexpected disk partitions.</p>
publicIP	No	NodePublicIP object	<p>EIP of a node.</p> <p>NOTE</p> <p>This parameter is not supported when you add a node to a node pool.</p>
nodeNicSpec	No	NodeNicSpec 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	<p>Node billing mode.</p> <ul style="list-style-type: none"> • 0: pay-per-use

Parameter	Mandatory	Type	Description
taints	No	Array of Taint 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"> • Key: A key must contain 1 to 63 characters starting with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain name can be used as the prefix of a key. • Value: A value must start with a letter or digit and can contain a maximum of 63 characters, including letters, digits, hyphens (-), underscores (_), and periods (.). • Effect: Available options are NoSchedule, PreferNoSchedule, and NoExecute. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In using 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>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"> • Key: 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 and contain a maximum of 253 characters. Example DNS subdomain: example.com/my-key • Value: The value can be left blank or contain 1 to 63 characters that start with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed in the character string. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In using 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>Example:</p> <pre>"k8sTags": { "key": "value" }</pre>
ecsGroupId	No	String	Cloud server group ID. If this field is specified, the node is created in the specified cloud server group.
dedicatedHostId	No	String	<p>ID of the DeH to which the node is scheduled.</p> <p>NOTE This parameter is not supported when you add a node during node pool creation.</p>

Parameter	Mandatory	Type	Description
userTags	No	Array of UserTag objects	Cloud server tag. The key of a tag must be unique. The maximum number of custom tags supported by CCE depends on the region 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 using 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.
runtime	No	Runtime object	Container runtime. Defaults to docker .

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

Table 4-419 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 4-420 UserPassword

Parameter	Mandatory	Type	Description
username	No	String	Login account. Defaults to root .
password	Yes	String	<p>If a username and a 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"> • Contains 8 to 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. The password field must be salted during node creation. For details, see Adding a Salt in the password Field When Creating a Node.

Table 4-421 Volume

Parameter	Mandatory	Type	Description
size	Yes	Integer	<p>Disk size, in GB.</p> <ul style="list-style-type: none"> • System disk: 40 to 1024 - Value range for data disks: 100 to 32768

Parameter	Mandatory	Type	Description
volumetype	Yes	String	Disk type. For details about possible values, see the description of the root_volume parameter in the API used to create an ECS in the ECS API reference. <ul style="list-style-type: none"> • SAS: high I/O disk • SSD: ultra-high I/O disk • SATA: common I/O disk SATA disks have been removed from EVS. You can find them attached only on existing nodes.
extendParam	No	Map<String,Object>	Extended disk parameters, defined in extendparam 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 dssPoolID , 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 dss . This field is used only for DeC clusters.
hw:passthrough	No	Boolean	<ul style="list-style-type: none"> • Pay attention to this field if your ECS is SDI-compliant. If the value of this field is true, an SCSI disk will be created. • If the node pool type is ElasticBMS, this field must be set to true. • If a node specification involves local disks and EVS disks at the same time, set the disk initialization parameters. For details, see Attaching Disks to a Node.

Parameter	Mandatory	Type	Description
metadata	No	VolumeMetadata 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 4-422 VolumeMetadata

Parameter	Mandatory	Type	Description
<code>__system__encrypted</code>	No	String	Whether the EVS disk is encrypted. The value 0 indicates that the EVS disk is not encrypted, and the value 1 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 metadata . This field is used with __system__encrypted .

Table 4-423 Storage

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

Table 4-424 StorageSelectors

Parameter	Mandatory	Type	Description
name	Yes	String	Selector name, used as the index of selectorNames in storageGroup . Therefore, the name of each selector must be unique.
storageType	Yes	String	Specifies the storage type. Currently, only evs (EVS volumes) and local (local volumes) are supported. The local storage does not support disk selection. All local disks will form a VG. Therefore, only one storageSelector of the local type is allowed.
matchLabels	No	matchLabels object	Matching field of an EVS volume. The size , volumeType , metadataEncrypted , metadataCmkid and count fields are supported.

Table 4-425 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. Currently, SSD, GPSSD and SAS are supported.
metadataEncrypted	No	String	Disk encryption identifier. 0 indicates that the disk is not encrypted, and 1 indicates that the disk is encrypted.
metadataCmkid	No	String	Customer master key ID of an encrypted disk. The value is a 36-byte string.
count	No	String	Number of disks to be selected. If this parameter is left blank, all disks of this type are selected.

Table 4-426 StorageGroups

Parameter	Mandatory	Type	Description
name	Yes	String	Name of a virtual storage group, which must be unique.
cceManaged	No	Boolean	Storage space for Kubernetes and runtime components. Only one group can be set to true . If this parameter is left blank, the default value false is used.
selectorNames	Yes	Array of strings	This parameter corresponds to name in storageSelectors . A group can match multiple selectors, but a selector can match only one group.
virtualSpaces	Yes	Array of VirtualSpace objects	Detailed management of space configuration in a group.

Table 4-427 VirtualSpace

Parameter	Mandatory	Type	Description
name	Yes	String	Name of a virtualSpace. <ul style="list-style-type: none"> • Kubernetes: Kubernetes space configuration. lvmConfig needs to be configured. • runtime: runtime space configuration. runtimeConfig needs to be configured. • user: user space configuration. lvmConfig needs to be configured.
size	Yes	String	Size of a virtualSpace. The value must be an integer in percentage. Example: 90%. <p>NOTE The sum of the percentages of all virtualSpaces in a group cannot exceed 100%.</p>

Parameter	Mandatory	Type	Description
lvmConfig	No	LVMConfig object	LVM configurations, applicable to kubernetes and user spaces. Note that one virtual space supports only one config.
runtimeConfig	No	RuntimeConfig object	runtime configurations, applicable to the runtime space. Note that one virtual space supports only one config.

Table 4-428 LVMConfig

Parameter	Mandatory	Type	Description
lvType	Yes	String	LVM write mode. linear indicates the linear mode. striped indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.
path	No	String	Path to which the disk is attached. This parameter takes effect only in user configuration. The value is an absolute path. Digits, letters, periods (.), hyphens (-), and underscores (_) are allowed.

Table 4-429 RuntimeConfig

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

Table 4-430 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. NOTE If ids has been set, you do not need to set count and eip .
count	No	Integer	Number of EIPs to be dynamically created. NOTE count and eip must be set at the same time.
eip	No	NodeEIPSpec object	EIP configuration.

Table 4-431 NodeEIPSpec

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

Table 4-432 NodeBandwidth

Parameter	Mandatory	Type	Description
chargemode	No	String	<p>Bandwidth billing mode.</p> <ul style="list-style-type: none"> If this field is not specified, the billing is based on bandwidth. - If the field is null, the billing is based on bandwidth. If the field value is traffic, the billing is based on traffic. If the value is out of the preceding options, the cloud server will fail to be created. <p>NOTE</p> <ul style="list-style-type: none"> 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%. 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%.
size	No	Integer	Bandwidth size, specified in bandwidth.size in the API for assigning an EIP. . .
sharetype	No	String	Bandwidth sharing type. Value options: PER (exclusive bandwidth)

Table 4-433 NodeNicSpec

Parameter	Mandatory	Type	Description
primaryNic	No	NicSpec object	Description of the primary NIC.

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

Table 4-434 NicSpec

Parameter	Mandatory	Type	Description
subnetId	No	String	ID of the subnet to which the NIC belongs. If subnetId is not specified when creating the primary NIC, the cluster subnet is used. When creating a secondary NIC, you must specify subnetId .
fixedIps	No	Array of strings	The IP address of the primary NIC is specified by fixedIps . The number of IP addresses cannot be greater than the number of created nodes. fixedIps and ipBlock cannot be specified at the same time.
ipBlock	No	String	CIDR format of the primary NIC IP range. The IP address of the created node falls in this range. fixedIps and ipBlock cannot be specified at the same time.

Table 4-435 Taint

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

Table 4-436 UserTag

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

Table 4-437 Runtime

Parameter	Mandatory	Type	Description
name	No	String	Container runtime. Defaults to docker .

Table 4-438 NodeExtendParam

Parameter	Mandatory	Type	Description
ecs:performancetype	No	String	ECS flavor types. This field is returned in the response.
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. The number of pods that can be created on a node is determined by multiple parameters. For details, see Maximum Number of Pods That Can Be Created on a Node .

Parameter	Mandatory	Type	Description
DockerLVMConfigOverride	No	String	<p>Docker data disk configuration item. (This parameter has been discarded. Use the storage 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 diskType based on the actual drive letter type. The following fields are included:</p> <ul style="list-style-type: none"> • userLV (optional): size of the user space, for example, vgpaas/20%VG. • userPath (optional): mount path of the user space, for example, /home/wqt-test. • diskType: disk type. Currently, only evs, hdd, and ssd are supported. • lvType: type of a logic volume. The value can be linear or striped. • dockerThinpool: Docker space size, for example, vgpaas/60%VG. • kubernetesLV: kubelet space size, for example, vgpaas/20%VG.

Parameter	Mandatory	Type	Description
dockerBaseSize	No	Integer	<p>Available disk space of a single container on a node, in GB.</p> <p>If this parameter is left blank or is set to 0, the default value is used. In Device Mapper mode, the default value is 10. In OverlayFS mode, the available space of a single container is not limited by default, and the dockerBaseSize setting takes effect only on EulerOS nodes in the cluster of the new version.</p> <p>For details about how to allocate the space for the container runtime, see Data Disk Space Allocation.</p> <p>When Device Mapper is used, you are advised to set dockerBaseSize 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>
publicKey	No	String	Public key of a node.
alpha.cce/preInstall	No	String	<p>Pre-installation script.</p> <p>NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)</p>
alpha.cce/postInstall	No	String	<p>Post-installation script.</p> <p>NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)</p>
alpha.cce/NodeImageID	No	String	This parameter is required when a custom image is used to create a BMS node.

Parameter	Mandatory	Type	Description
enterprise_project_id	No	String	ID of the enterprise project to which the node belongs
chargingMode	No	Integer	Billing mode of a node. This parameter has been deprecated. Use the billingMode parameter in NodeSpec .
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.

Table 4-439 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 minNodeCount 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.

Parameter	Mandatory	Type	Description
priority	No	Integer	Weight of a node pool. A node pool with a higher weight has a higher priority during scaling.

Table 4-440 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 4-441 SecurityID

Parameter	Mandatory	Type	Description
id	No	String	Security group ID

Table 4-442 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"> • Null: available (the number of current nodes in the node pool has reached the expected value, and no node scaling is being performed.) • 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.) • Synchronized: pending scaling (the number of current nodes in the node pool does not reach the expected value, or node scaling is being performed.) • 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.) <p>NOTE 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"> • Node pool scaling status: You can obtain accurate status of the current node pool based on parameters such as currentNode, creatingNode, or deletingNode. • Node pool scale-out: You can use conditions to obtain the detailed status of a node pool. Scalable can replace SoldOut. • Deleting: The object is being deleted. • Error: An error occurs.

Parameter	Mandatory	Type	Description
jobId	No	String	ID of a job executed on the node pool
conditions	No	Array of NodePoolCondition objects	Node pool status details. For details, see the definition of Condition .

Table 4-443 NodePoolCondition

Parameter	Mandatory	Type	Description
type	No	String	Condition type. The options are as follows: <ul style="list-style-type: none"> • Scalable: whether a node pool can be scaled. If the status is False, node pool scaling will not be triggered again. • QuotaInsufficient: Quotas on which node pool scaling depends are insufficient, affecting the node pool scaling status. • ResourceInsufficient: Resources on which node pool scaling depends are insufficient, affecting the node pool scaling status. • UnexpectedError: The node pool fails to be scaled out due to unexpected reasons, affecting the node pool scaling status. • Error: A node pool error occurs. A common trigger is deletion failure.
status	No	String	Current status of Condition . The options are as follows: <ul style="list-style-type: none"> • "True" • "False"
lastProbeTime	No	String	Time when the status was last checked
lastTransitTime	No	String	Time when the status was last changed

Parameter	Mandatory	Type	Description
reason	No	String	Reason why the status was last changed
message	No	String	Detailed condition description

Response Parameters

Status code: 201

Table 4-444 Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at NodePool .
apiVersion	String	API version. The value is fixed at v3 .
metadata	NodePoolMetadata object	Metadata information of the node pool
spec	NodePoolSpec object	Node pool specifications
status	NodePoolStatus object	Node pool status

Table 4-445 NodePoolMetadata

Parameter	Type	Description
name	String	Node pool name. NOTE 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"> You cannot create node pools named DefaultPool.
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
updateTimestamp	String	Update time.

Parameter	Type	Description
creationTimes stamp	String	Creation time.

Table 4-446 NodePoolSpec

Parameter	Type	Description
type	String	Node pool type. If this parameter is left blank, the value vm is used by default. <ul style="list-style-type: none"> • vm: ECS • ElasticBMS: C6 general computing-plus BMS. An example specification is c6.22xlarge.2.physical.
nodeTemplate	NodeSpec 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	NodePoolNodeAutoscaling object	Auto scaling parameters
nodeManagement	NodeManagement object	Node management configuration
podSecurityGroups	Array of SecurityID objects	Security groups configuration
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"> • Specifying no security group ID will add the new nodes to the default security group of the worker nodes. • Specifying a valid security group ID will put new nodes in that security group. • When specifying a security group, do not modify the rules of the port on which CCE running depends.

Table 4-447 NodeSpec

Parameter	Type	Description
flavor	String	Node specifications. For details about the node specifications supported by CCE, see Node Flavor Description .
az	String	Name of the AZ where the node to be created is located. . .
os	String	Node OS. For details about the supported OSs, see Node OS Description . NOTE <ul style="list-style-type: none"> 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. If <code>alpha.cce/NodeImageID</code> in <code>extendParam</code> is specified during node creation, you do not need to set this field.
login	Login object	Node login mode. Either the key pair or password must be used for login.
rootVolume	Volume object	Information about disks on the node
dataVolumes	Array of Volume objects	Data disk parameters of the node. Currently, you can add the second data disk for your node on the CCE console. This data disk is used by the container runtime and kubelet. Do not uninstall this disk. Otherwise, the node will become unavailable. For DeC nodes, the parameter description is the same as that for rootVolume .
storage	Storage object	Disk initialization management parameter. This parameter is complex to configure. For details, see Attaching Disks to a Node . If this parameter retains its default, disks are managed based on the <code>DockerLVMConfigOverride (discarded)</code> parameter in <code>extendParam</code> . This parameter is supported by clusters of version 1.15.11 and later. NOTE If a node specification involves local disks and EVS disks at the same time, do not retain the default value of this parameter to prevent unexpected disk partitions.
publicIP	NodePublicIP object	EIP of a node. NOTE This parameter is not supported when you add a node to a node pool.

Parameter	Type	Description
nodeNicSpec	NodeNicSpec 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"> 0: pay-per-use
taints	Array of Taint 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"> Key: A key must contain 1 to 63 characters starting with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain name can be used as the prefix of a key. Value: A value must start with a letter or digit and can contain a maximum of 63 characters, including letters, digits, hyphens (-), underscores (_), and periods (.). Effect: Available options are NoSchedule, PreferNoSchedule, and NoExecute. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In using 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>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"> • Key: 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 and contain a maximum of 253 characters. Example DNS subdomain: example.com/my-key • Value: The value can be left blank or contain 1 to 63 characters that start with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed in the character string. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In using 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>Example:</p> <pre>"k8sTags": { "key": "value" }</pre>
ecsGroupId	String	Cloud server group ID. If this field is specified, the node is created in the specified cloud server group.
dedicatedHostId	String	<p>ID of the DeH to which the node is scheduled.</p> <p>NOTE This parameter is not supported when you add a node during node pool creation.</p>
userTags	Array of UserTag objects	Cloud server tag. The key of a tag must be unique. The maximum number of custom tags supported by CCE depends on the region 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 using 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.
runtime	Runtime object	Container runtime. Defaults to docker .

Parameter	Type	Description
initializedConditions	Array of strings	<p>Custom initialization flag.</p> <p>Before CCE nodes are initialized, they are tainted with node.cloudprovider.kubernetes.io/uninitialized to prevent pods from being scheduled to them.</p> <p>CCE supports custom initialization flags. After receiving the initializedConditions parameter, CCE converts the parameter value into a node label and provisions the label with the node, for example, cloudprovider.openvessel.io/inject-initialized-conditions=CCEInitial_CustomedInitial.</p> <p>After the node is labeled, its status.Conditions is polled to check whether the type of conditions has a flag name, such as CCEInitial and CustomedInitial. If all input flags exist and their status is True, the node initialization is complete and the initialization taint is removed.</p> <ul style="list-style-type: none"> • Use only letters and digits. Max. characters: 20. • Max. flags: 2.
extendParam	NodeExtendParam object	Extended parameters for creating a node.

Table 4-448 Login

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

Table 4-449 UserPassword

Parameter	Type	Description
username	String	Login account. Defaults to root .

Parameter	Type	Description
password	String	<p>If a username and a 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"> • Contains 8 to 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. The password field must be salted during node creation. For details, see Adding a Salt in the password Field When Creating a Node.

Table 4-450 Volume

Parameter	Type	Description
size	Integer	<p>Disk size, in GB.</p> <ul style="list-style-type: none"> • System disk: 40 to 1024 -Value range for data disks: 100 to 32768
volumetype	String	<p>Disk type. For details about possible values, see the description of the root_volume parameter in the API used to create an ECS in the ECS API reference.</p> <ul style="list-style-type: none"> • SAS: high I/O disk • SSD: ultra-high I/O disk • SATA: common I/O disk SATA disks have been removed from EVS. You can find them attached only on existing nodes.
extendParam	Map<String, Object>	<p>Extended disk parameters, defined in extendparam 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 dssPoolID, 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 dss. This field is used only for DeC clusters.</p>

Parameter	Type	Description
hw:passthrough	Boolean	<ul style="list-style-type: none"> Pay attention to this field if your ECS is SDI-compliant. If the value of this field is true, an SCSI disk will be created. If the node pool type is ElasticBMS, this field must be set to true. If a node specification involves local disks and EVS disks at the same time, set the disk initialization parameters. For details, see Attaching Disks to a Node.
metadata	VolumeMetadata 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 4-451 VolumeMetadata

Parameter	Type	Description
__system__encrypted	String	Whether the EVS disk is encrypted. The value 0 indicates that the EVS disk is not encrypted, and the value 1 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 metadata . This field is used with __system__encrypted .

Table 4-452 Storage

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

Table 4-453 StorageSelectors

Parameter	Type	Description
name	String	Selector name, used as the index of selectorNames in storageGroup . Therefore, the name of each selector must be unique.
storageType	String	Specifies the storage type. Currently, only evs (EVS volumes) and local (local volumes) are supported. The local storage does not support disk selection. All local disks will form a VG. Therefore, only one storageSelector of the local type is allowed.
matchLabels	matchLabels object	Matching field of an EVS volume. The size , volumeType , metadataEncrypted , metadataCmkid and count fields are supported.

Table 4-454 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. Currently, SSD, GPSSD and SAS are supported.
metadataEncrypted	String	Disk encryption identifier. 0 indicates that the disk is not encrypted, and 1 indicates that the disk is encrypted.
metadataCmkid	String	Customer master key ID of an encrypted disk. The value is a 36-byte string.
count	String	Number of disks to be selected. If this parameter is left blank, all disks of this type are selected.

Table 4-455 StorageGroups

Parameter	Type	Description
name	String	Name of a virtual storage group, which must be unique.

Parameter	Type	Description
cceManaged	Boolean	Storage space for Kubernetes and runtime components. Only one group can be set to true . If this parameter is left blank, the default value false is used.
selectorNames	Array of strings	This parameter corresponds to name in storageSelectors . A group can match multiple selectors, but a selector can match only one group.
virtualSpaces	Array of VirtualSpace objects	Detailed management of space configuration in a group.

Table 4-456 VirtualSpace

Parameter	Type	Description
name	String	Name of a virtualSpace. <ul style="list-style-type: none"> • Kubernetes: Kubernetes space configuration. lvmConfig needs to be configured. • runtime: runtime space configuration. runtimeConfig needs to be configured. • user: user space configuration. lvmConfig needs to be configured.
size	String	Size of a virtualSpace. The value must be an integer in percentage. Example: 90%. NOTE The sum of the percentages of all virtualSpaces in a group cannot exceed 100%.
lvmConfig	LVMConfig object	LVM configurations, applicable to kubernetes and user spaces. Note that one virtual space supports only one config.
runtimeConfig	RuntimeConfig object	runtime configurations, applicable to the runtime space. Note that one virtual space supports only one config.

Table 4-457 LVMConfig

Parameter	Type	Description
lvType	String	LVM write mode. linear indicates the linear mode. striped indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.
path	String	Path to which the disk is attached. This parameter takes effect only in user configuration. The value is an absolute path. Digits, letters, periods (.), hyphens (-), and underscores (_) are allowed.

Table 4-458 RuntimeConfig

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

Table 4-459 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. NOTE If ids has been set, you do not need to set count and eip .
count	Integer	Number of EIPs to be dynamically created. NOTE count and eip must be set at the same time.
eip	NodeEIPSpec object	EIP configuration.

Table 4-460 NodeEIPSpec

Parameter	Type	Description
iptype	String	EIP type, specified in publicip.type in the API for assigning an EIP. . .

Parameter	Type	Description
bandwidth	NodeBandwidth object	Bandwidth parameters of the EIP

Table 4-461 NodeBandwidth

Parameter	Type	Description
chargemode	String	<p>Bandwidth billing mode.</p> <ul style="list-style-type: none"> If this field is not specified, the billing is based on bandwidth. - If the field is null, the billing is based on bandwidth. If the field value is traffic, the billing is based on traffic. If the value is out of the preceding options, the cloud server will fail to be created. <p>NOTE</p> <ul style="list-style-type: none"> 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%. 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%.
size	Integer	Bandwidth size, specified in bandwidth.size in the API for assigning an EIP. . .
sharetype	String	Bandwidth sharing type. Value options: PER (exclusive bandwidth)

Table 4-462 NodeNicSpec

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

Table 4-463 NicSpec

Parameter	Type	Description
subnetId	String	ID of the subnet to which the NIC belongs. If subnetId is not specified when creating the primary NIC, the cluster subnet is used. When creating a secondary NIC, you must specify subnetId .
fixedIps	Array of strings	The IP address of the primary NIC is specified by fixedIps . The number of IP addresses cannot be greater than the number of created nodes. fixedIps and ipBlock cannot be specified at the same time.
ipBlock	String	CIDR format of the primary NIC IP range. The IP address of the created node falls in this range. fixedIps and ipBlock cannot be specified at the same time.

Table 4-464 Taint

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

Table 4-465 UserTag

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

Table 4-466 Runtime

Parameter	Type	Description
name	String	Container runtime. Defaults to docker .

Table 4-467 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 Maximum Number of Pods That Can Be Created on a Node.</p>
DockerLVMDiskConfigOverride	String	<p>Docker data disk configuration item. (This parameter has been discarded. Use the storage field instead.) Example default configuration: <code>"DockerLVMDiskConfigOverride": "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 diskType based on the actual drive letter type. The following fields are included:</p> <ul style="list-style-type: none"> • userLV (optional): size of the user space, for example, vgpaas/20%VG. • userPath (optional): mount path of the user space, for example, /home/wqt-test. • diskType: disk type. Currently, only evs, hdd, and ssd are supported. • lvType: type of a logic volume. The value can be linear or striped. • dockerThinpool: Docker space size, for example, vgpaas/60%VG. • kubernetesLV: kubelet space size, for example, vgpaas/20%VG.

Parameter	Type	Description
dockerBaseSize	Integer	<p>Available disk space of a single container on a node, in GB.</p> <p>If this parameter is left blank or is set to 0, the default value is used. In Device Mapper mode, the default value is 10. In OverlayFS mode, the available space of a single container is not limited by default, and the dockerBaseSize setting takes effect only on EulerOS nodes in the cluster of the new version.</p> <p>For details about how to allocate the space for the container runtime, see Data Disk Space Allocation.</p> <p>When Device Mapper is used, you are advised to set dockerBaseSize 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>
publicKey	String	Public key of a node.
alpha.cce/preInstall	String	<p>Pre-installation script.</p> <p>NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)</p>
alpha.cce/postInstall	String	<p>Post-installation script.</p> <p>NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)</p>
alpha.cce/NodeImageID	String	This parameter is required when a custom image is used to create a BMS node.
enterprise_project_id	String	ID of the enterprise project to which the node belongs
chargingMode	Integer	Billing mode of a node. This parameter has been deprecated. Use the billingMode parameter in NodeSpec .
agency_name	String	<p>Name of an agency</p> <p>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.</p>

Table 4-468 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 minNodeCount 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 4-469 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 4-470 SecurityID

Parameter	Type	Description
id	String	Security group ID

Table 4-471 NodePoolStatus

Parameter	Type	Description
currentNode	Integer	Total number of nodes in the current node pool (excluding the nodes that are being deleted)

Parameter	Type	Description
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.
phase	String	<p>Node pool status.</p> <ul style="list-style-type: none"> • Null: available (the number of current nodes in the node pool has reached the expected value, and no node scaling is being performed.) • 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.) • Synchronized: pending scaling (the number of current nodes in the node pool does not reach the expected value, or node scaling is being performed.) • 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.) <p>NOTE 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"> • Node pool scaling status: You can obtain accurate status of the current node pool based on parameters such as currentNode, creatingNode, or deletingNode. • Node pool scale-out: You can use conditions to obtain the detailed status of a node pool. Scalable can replace SoldOut. • Deleting: The object is being deleted. • Error: An error occurs.
jobId	String	ID of a job executed on the node pool
conditions	Array of NodePoolCondition objects	Node pool status details. For details, see the definition of Condition .

Table 4-472 NodePoolCondition

Parameter	Type	Description
type	String	Condition type. The options are as follows: <ul style="list-style-type: none"> • Scalable: whether a node pool can be scaled. If the status is False, node pool scaling will not be triggered again. • QuotaInsufficient: Quotas on which node pool scaling depends are insufficient, affecting the node pool scaling status. • ResourceInsufficient: Resources on which node pool scaling depends are insufficient, affecting the node pool scaling status. • UnexpectedError: The node pool fails to be scaled out due to unexpected reasons, affecting the node pool scaling status. • Error: A node pool error occurs. A common trigger is deletion failure.
status	String	Current status of Condition . The options are as follows: <ul style="list-style-type: none"> • "True" • "False"
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

```
{
  "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 4-473 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
nodepool_id	Yes	String	Node pool ID.

Request Parameters

Table 4-474 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 Obtaining a User Token .

Response Parameters

Status code: 200

Table 4-475 Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at NodePool .
apiVersion	String	API version. The value is fixed at v3 .
metadata	NodePoolMetadata object	Metadata information of the node pool
spec	NodePoolSpec object	Node pool specifications

Parameter	Type	Description
status	NodePoolStatus object	Node pool status

Table 4-476 NodePoolMetadata

Parameter	Type	Description
name	String	Node pool name. NOTE 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"> You cannot create node pools named DefaultPool.
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
updateTimestamp	String	Update time.
creationTimestamp	String	Creation time.

Table 4-477 NodePoolSpec

Parameter	Type	Description
type	String	Node pool type. If this parameter is left blank, the value vm is used by default. <ul style="list-style-type: none"> vm: ECS ElasticBMS: C6 general computing-plus BMS. An example specification is c6.22xlarge.2.physical.
nodeTemplate	NodeSpec 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.

Parameter	Type	Description
autoscaling	NodePoolNo deAutoscaling object	Auto scaling parameters
nodeManagement	NodeManagement object	Node management configuration
podSecurityGroups	Array of SecurityID objects	Security groups configuration
customSecurityGroups	Array of strings	<p>Custom security group settings for a node pool. New nodes scaled out in a node pool can be bound to a specified security group.</p> <ul style="list-style-type: none"> • Specifying no security group ID will add the new nodes to the default security group of the worker nodes. • Specifying a valid security group ID will put new nodes in that security group. • When specifying a security group, do not modify the rules of the port on which CCE running depends.

Table 4-478 NodeSpec

Parameter	Type	Description
flavor	String	Node specifications. For details about the node specifications supported by CCE, see Node Flavor Description .
az	String	Name of the AZ where the node to be created is located. . .
os	String	<p>Node OS. For details about the supported OSs, see Node OS Description.</p> <p>NOTE</p> <ul style="list-style-type: none"> • 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. • If <code>alpha.cce/NodeImageID</code> in <code>extendParam</code> is specified during node creation, you do not need to set this field.
login	Login object	Node login mode. Either the key pair or password must be used for login.
rootVolume	Volume object	Information about disks on the node

Parameter	Type	Description
dataVolumes	Array of Volume objects	Data disk parameters of the node. Currently, you can add the second data disk for your node on the CCE console. This data disk is used by the container runtime and kubelet. Do not uninstall this disk. Otherwise, the node will become unavailable. For DeC nodes, the parameter description is the same as that for rootVolume .
storage	Storage object	Disk initialization management parameter. This parameter is complex to configure. For details, see Attaching Disks to a Node . If this parameter retains its default, disks are managed based on the DockerLVMConfigOverride (discarded) parameter in extendParam. This parameter is supported by clusters of version 1.15.11 and later. NOTE If a node specification involves local disks and EVS disks at the same time, do not retain the default value of this parameter to prevent unexpected disk partitions.
publicIP	NodePublicIP object	EIP of a node. NOTE This parameter is not supported when you add a node to a node pool.
nodeNicSpec	NodeNicSpec 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"> 0: pay-per-use

Parameter	Type	Description
taints	Array of Taint 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"> • Key: A key must contain 1 to 63 characters starting with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain name can be used as the prefix of a key. • Value: A value must start with a letter or digit and can contain a maximum of 63 characters, including letters, digits, hyphens (-), underscores (_), and periods (.). • Effect: Available options are NoSchedule, PreferNoSchedule, and NoExecute. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In using 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>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"> • Key: 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 and contain a maximum of 253 characters. Example DNS subdomain: example.com/my-key • Value: The value can be left blank or contain 1 to 63 characters that start with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed in the character string. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In using 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>Example:</p> <pre>"k8sTags": { "key": "value" }</pre>
ecsGroupId	String	Cloud server group ID. If this field is specified, the node is created in the specified cloud server group.
dedicatedHostId	String	<p>ID of the DeH to which the node is scheduled.</p> <p>NOTE This parameter is not supported when you add a node during node pool creation.</p>
userTags	Array of UserTag objects	Cloud server tag. The key of a tag must be unique. The maximum number of custom tags supported by CCE depends on the region 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 using 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.
runtime	Runtime object	Container runtime. Defaults to docker .

Parameter	Type	Description
initializedConditions	Array of strings	<p>Custom initialization flag.</p> <p>Before CCE nodes are initialized, they are tainted with node.cloudprovider.kubernetes.io/uninitialized to prevent pods from being scheduled to them.</p> <p>CCE supports custom initialization flags. After receiving the initializedConditions parameter, CCE converts the parameter value into a node label and provisions the label with the node, for example, cloudprovider.openvessel.io/inject-initialized-conditions=CCEInitial_CustomedInitial.</p> <p>After the node is labeled, its status.Conditions is polled to check whether the type of conditions has a flag name, such as CCEInitial and CustomedInitial. If all input flags exist and their status is True, the node initialization is complete and the initialization taint is removed.</p> <ul style="list-style-type: none"> • Use only letters and digits. Max. characters: 20. • Max. flags: 2.
extendParam	NodeExtendParam object	Extended parameters for creating a node.

Table 4-479 Login

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

Table 4-480 UserPassword

Parameter	Type	Description
username	String	Login account. Defaults to root .

Parameter	Type	Description
password	String	<p>If a username and a 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"> • Contains 8 to 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. The password field must be salted during node creation. For details, see Adding a Salt in the password Field When Creating a Node.

Table 4-481 Volume

Parameter	Type	Description
size	Integer	<p>Disk size, in GB.</p> <ul style="list-style-type: none"> • System disk: 40 to 1024 -Value range for data disks: 100 to 32768
volumetype	String	<p>Disk type. For details about possible values, see the description of the root_volume parameter in the API used to create an ECS in the ECS API reference.</p> <ul style="list-style-type: none"> • SAS: high I/O disk • SSD: ultra-high I/O disk • SATA: common I/O disk SATA disks have been removed from EVS. You can find them attached only on existing nodes.
extendParam	Map<String, Object>	<p>Extended disk parameters, defined in extendparam 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 dssPoolID, 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 dss. This field is used only for DeC clusters.</p>

Parameter	Type	Description
hw:passthrough	Boolean	<ul style="list-style-type: none"> Pay attention to this field if your ECS is SDI-compliant. If the value of this field is true, an SCSI disk will be created. If the node pool type is ElasticBMS, this field must be set to true. If a node specification involves local disks and EVS disks at the same time, set the disk initialization parameters. For details, see Attaching Disks to a Node.
metadata	VolumeMetadata 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 4-482 VolumeMetadata

Parameter	Type	Description
__system__encrypted	String	Whether the EVS disk is encrypted. The value 0 indicates that the EVS disk is not encrypted, and the value 1 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 metadata . This field is used with __system__encrypted .

Table 4-483 Storage

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

Table 4-484 StorageSelectors

Parameter	Type	Description
name	String	Selector name, used as the index of selectorNames in storageGroup . Therefore, the name of each selector must be unique.
storageType	String	Specifies the storage type. Currently, only evs (EVS volumes) and local (local volumes) are supported. The local storage does not support disk selection. All local disks will form a VG. Therefore, only one storageSelector of the local type is allowed.
matchLabels	matchLabels object	Matching field of an EVS volume. The size , volumeType , metadataEncrypted , metadataCmkid and count fields are supported.

Table 4-485 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. Currently, SSD, GPSSD and SAS are supported.
metadataEncrypted	String	Disk encryption identifier. 0 indicates that the disk is not encrypted, and 1 indicates that the disk is encrypted.
metadataCmkid	String	Customer master key ID of an encrypted disk. The value is a 36-byte string.
count	String	Number of disks to be selected. If this parameter is left blank, all disks of this type are selected.

Table 4-486 StorageGroups

Parameter	Type	Description
name	String	Name of a virtual storage group, which must be unique.

Parameter	Type	Description
cceManaged	Boolean	Storage space for Kubernetes and runtime components. Only one group can be set to true . If this parameter is left blank, the default value false is used.
selectorNames	Array of strings	This parameter corresponds to name in storageSelectors . A group can match multiple selectors, but a selector can match only one group.
virtualSpaces	Array of VirtualSpace objects	Detailed management of space configuration in a group.

Table 4-487 VirtualSpace

Parameter	Type	Description
name	String	Name of a virtualSpace. <ul style="list-style-type: none"> • Kubernetes: Kubernetes space configuration. lvmConfig needs to be configured. • runtime: runtime space configuration. runtimeConfig needs to be configured. • user: user space configuration. lvmConfig needs to be configured.
size	String	Size of a virtualSpace. The value must be an integer in percentage. Example: 90%. NOTE The sum of the percentages of all virtualSpaces in a group cannot exceed 100%.
lvmConfig	LVMConfig object	LVM configurations, applicable to kubernetes and user spaces. Note that one virtual space supports only one config.
runtimeConfig	RuntimeConfig object	runtime configurations, applicable to the runtime space. Note that one virtual space supports only one config.

Table 4-488 LVMConfig

Parameter	Type	Description
lvType	String	LVM write mode. linear indicates the linear mode. striped indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.
path	String	Path to which the disk is attached. This parameter takes effect only in user configuration. The value is an absolute path. Digits, letters, periods (.), hyphens (-), and underscores (_) are allowed.

Table 4-489 RuntimeConfig

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

Table 4-490 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. NOTE If ids has been set, you do not need to set count and eip .
count	Integer	Number of EIPs to be dynamically created. NOTE count and eip must be set at the same time.
eip	NodeEIPSpec object	EIP configuration.

Table 4-491 NodeEIPSpec

Parameter	Type	Description
iptype	String	EIP type, specified in publicip.type in the API for assigning an EIP. . .

Parameter	Type	Description
bandwidth	NodeBandwidth object	Bandwidth parameters of the EIP

Table 4-492 NodeBandwidth

Parameter	Type	Description
chargemode	String	Bandwidth billing mode. <ul style="list-style-type: none"> If this field is not specified, the billing is based on bandwidth. - If the field is null, the billing is based on bandwidth. If the field value is traffic, the billing is based on traffic. If the value is out of the preceding options, the cloud server will fail to be created. NOTE <ul style="list-style-type: none"> 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%. 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%.
size	Integer	Bandwidth size, specified in bandwidth.size in the API for assigning an EIP. . .
sharetype	String	Bandwidth sharing type. Value options: PER (exclusive bandwidth)

Table 4-493 NodeNicSpec

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

Table 4-494 NicSpec

Parameter	Type	Description
subnetId	String	ID of the subnet to which the NIC belongs. If subnetId is not specified when creating the primary NIC, the cluster subnet is used. When creating a secondary NIC, you must specify subnetId .
fixedIps	Array of strings	The IP address of the primary NIC is specified by fixedIps . The number of IP addresses cannot be greater than the number of created nodes. fixedIps and ipBlock cannot be specified at the same time.
ipBlock	String	CIDR format of the primary NIC IP range. The IP address of the created node falls in this range. fixedIps and ipBlock cannot be specified at the same time.

Table 4-495 Taint

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

Table 4-496 UserTag

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

Table 4-497 Runtime

Parameter	Type	Description
name	String	Container runtime. Defaults to docker .

Table 4-498 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 Maximum Number of Pods That Can Be Created on a Node.</p>
DockerLVMDockerConfigOverride	String	<p>Docker data disk configuration item. (This parameter has been discarded. Use the storage field instead.) Example default configuration: "DockerLVMDockerConfigOverride":"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 diskType based on the actual drive letter type. The following fields are included:</p> <ul style="list-style-type: none"> • userLV (optional): size of the user space, for example, vgpaas/20%VG. • userPath (optional): mount path of the user space, for example, /home/wqt-test. • diskType: disk type. Currently, only evs, hdd, and ssd are supported. • lvType: type of a logic volume. The value can be linear or striped. • dockerThinpool: Docker space size, for example, vgpaas/60%VG. • kubernetesLV: kubelet space size, for example, vgpaas/20%VG.

Parameter	Type	Description
dockerBaseSize	Integer	<p>Available disk space of a single container on a node, in GB.</p> <p>If this parameter is left blank or is set to 0, the default value is used. In Device Mapper mode, the default value is 10. In OverlayFS mode, the available space of a single container is not limited by default, and the dockerBaseSize setting takes effect only on EulerOS nodes in the cluster of the new version.</p> <p>For details about how to allocate the space for the container runtime, see Data Disk Space Allocation.</p> <p>When Device Mapper is used, you are advised to set dockerBaseSize 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>
publicKey	String	Public key of a node.
alpha.cce/preInstall	String	<p>Pre-installation script.</p> <p>NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)</p>
alpha.cce/postInstall	String	<p>Post-installation script.</p> <p>NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)</p>
alpha.cce/NodeImageID	String	This parameter is required when a custom image is used to create a BMS node.
enterprise_project_id	String	ID of the enterprise project to which the node belongs
chargingMode	Integer	Billing mode of a node. This parameter has been deprecated. Use the billingMode parameter in NodeSpec .
agency_name	String	<p>Name of an agency</p> <p>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.</p>

Table 4-499 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 minNodeCount 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 4-500 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 4-501 SecurityID

Parameter	Type	Description
id	String	Security group ID

Table 4-502 NodePoolStatus

Parameter	Type	Description
currentNode	Integer	Total number of nodes in the current node pool (excluding the nodes that are being deleted)

Parameter	Type	Description
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.
phase	String	<p>Node pool status.</p> <ul style="list-style-type: none"> • Null: available (the number of current nodes in the node pool has reached the expected value, and no node scaling is being performed.) • 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.) • Synchronized: pending scaling (the number of current nodes in the node pool does not reach the expected value, or node scaling is being performed.) • 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.) <p>NOTE 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"> • Node pool scaling status: You can obtain accurate status of the current node pool based on parameters such as currentNode, creatingNode, or deletingNode. • Node pool scale-out: You can use conditions to obtain the detailed status of a node pool. Scalable can replace SoldOut. • Deleting: The object is being deleted. • Error: An error occurs.
jobId	String	ID of a job executed on the node pool
conditions	Array of NodePoolCondition objects	Node pool status details. For details, see the definition of Condition .

Table 4-503 NodePoolCondition

Parameter	Type	Description
type	String	Condition type. The options are as follows: <ul style="list-style-type: none"> • Scalable: whether a node pool can be scaled. If the status is False, node pool scaling will not be triggered again. • QuotaInsufficient: Quotas on which node pool scaling depends are insufficient, affecting the node pool scaling status. • ResourceInsufficient: Resources on which node pool scaling depends are insufficient, affecting the node pool scaling status. • UnexpectedError: The node pool fails to be scaled out due to unexpected reasons, affecting the node pool scaling status. • Error: A node pool error occurs. A common trigger is deletion failure.
status	String	Current status of Condition . The options are as follows: <ul style="list-style-type: none"> • "True" • "False"
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 4-504 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .

Table 4-505 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 true , the default node pool is displayed.

Request Parameters

Table 4-506 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 Obtaining a User Token .

Response Parameters

Status code: 200

Table 4-507 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 NodePool objects	/

Table 4-508 NodePool

Parameter	Type	Description
kind	String	API type. The value is fixed at NodePool .
apiVersion	String	API version. The value is fixed at v3 .
metadata	NodePoolMetadata object	Metadata information of the node pool
spec	NodePoolSpec object	Node pool specifications
status	NodePoolStatus object	Node pool status

Table 4-509 NodePoolMetadata

Parameter	Type	Description
name	String	Node pool name. NOTE 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"> You cannot create node pools named DefaultPool.
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
updateTimestamp	String	Update time.
creationTimestamp	String	Creation time.

Table 4-510 NodePoolSpec

Parameter	Type	Description
type	String	Node pool type. If this parameter is left blank, the value vm is used by default. <ul style="list-style-type: none"> vm: ECS ElasticBMS: C6 general computing-plus BMS. An example specification is c6.22xlarge.2.physical.
nodeTemplate	NodeSpec 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	NodePoolNodeAutoscaling object	Auto scaling parameters
nodeManagement	NodeManagement object	Node management configuration

Parameter	Type	Description
podSecurityGroups	Array of SecurityID objects	Security groups configuration
customSecurityGroups	Array of strings	<p>Custom security group settings for a node pool. New nodes scaled out in a node pool can be bound to a specified security group.</p> <ul style="list-style-type: none"> • Specifying no security group ID will add the new nodes to the default security group of the worker nodes. • Specifying a valid security group ID will put new nodes in that security group. • When specifying a security group, do not modify the rules of the port on which CCE running depends.

Table 4-511 NodeSpec

Parameter	Type	Description
flavor	String	Node specifications. For details about the node specifications supported by CCE, see Node Flavor Description .
az	String	Name of the AZ where the node to be created is located. . .
os	String	<p>Node OS. For details about the supported OSs, see Node OS Description.</p> <p>NOTE</p> <ul style="list-style-type: none"> • 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. • If <code>alpha.cce/NodeImageID</code> in <code>extendParam</code> is specified during node creation, you do not need to set this field.
login	Login object	Node login mode. Either the key pair or password must be used for login.
rootVolume	Volume object	Information about disks on the node

Parameter	Type	Description
dataVolumes	Array of Volume objects	Data disk parameters of the node. Currently, you can add the second data disk for your node on the CCE console. This data disk is used by the container runtime and kubelet. Do not uninstall this disk. Otherwise, the node will become unavailable. For DeC nodes, the parameter description is the same as that for rootVolume .
storage	Storage object	Disk initialization management parameter. This parameter is complex to configure. For details, see Attaching Disks to a Node . If this parameter retains its default, disks are managed based on the DockerLVMConfigOverride (discarded) parameter in extendParam. This parameter is supported by clusters of version 1.15.11 and later. NOTE If a node specification involves local disks and EVS disks at the same time, do not retain the default value of this parameter to prevent unexpected disk partitions.
publicIP	NodePublicIP object	EIP of a node. NOTE This parameter is not supported when you add a node to a node pool.
nodeNicSpec	NodeNicSpec 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"> 0: pay-per-use

Parameter	Type	Description
taints	Array of Taint 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"> • Key: A key must contain 1 to 63 characters starting with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain name can be used as the prefix of a key. • Value: A value must start with a letter or digit and can contain a maximum of 63 characters, including letters, digits, hyphens (-), underscores (_), and periods (.). • Effect: Available options are NoSchedule, PreferNoSchedule, and NoExecute. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In using 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>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"> • Key: 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 and contain a maximum of 253 characters. Example DNS subdomain: example.com/my-key • Value: The value can be left blank or contain 1 to 63 characters that start with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed in the character string. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In using 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>Example:</p> <pre>"k8sTags": { "key": "value" }</pre>
ecsGroupId	String	Cloud server group ID. If this field is specified, the node is created in the specified cloud server group.
dedicatedHostId	String	<p>ID of the DeH to which the node is scheduled.</p> <p>NOTE This parameter is not supported when you add a node during node pool creation.</p>
userTags	Array of UserTag objects	Cloud server tag. The key of a tag must be unique. The maximum number of custom tags supported by CCE depends on the region 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 using 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.
runtime	Runtime object	Container runtime. Defaults to docker .

Parameter	Type	Description
initializedConditions	Array of strings	<p>Custom initialization flag.</p> <p>Before CCE nodes are initialized, they are tainted with node.cloudprovider.kubernetes.io/uninitialized to prevent pods from being scheduled to them.</p> <p>CCE supports custom initialization flags. After receiving the initializedConditions parameter, CCE converts the parameter value into a node label and provisions the label with the node, for example, cloudprovider.openvessel.io/inject-initialized-conditions=CCEInitial_CustomedInitial.</p> <p>After the node is labeled, its status.Conditions is polled to check whether the type of conditions has a flag name, such as CCEInitial and CustomedInitial. If all input flags exist and their status is True, the node initialization is complete and the initialization taint is removed.</p> <ul style="list-style-type: none"> • Use only letters and digits. Max. characters: 20. • Max. flags: 2.
extendParam	NodeExtendParam object	Extended parameters for creating a node.

Table 4-512 Login

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

Table 4-513 UserPassword

Parameter	Type	Description
username	String	Login account. Defaults to root .

Parameter	Type	Description
password	String	<p>If a username and a 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"> • Contains 8 to 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. The password field must be salted during node creation. For details, see Adding a Salt in the password Field When Creating a Node.

Table 4-514 Volume

Parameter	Type	Description
size	Integer	<p>Disk size, in GB.</p> <ul style="list-style-type: none"> • System disk: 40 to 1024 -Value range for data disks: 100 to 32768
volumetype	String	<p>Disk type. For details about possible values, see the description of the root_volume parameter in the API used to create an ECS in the ECS API reference.</p> <ul style="list-style-type: none"> • SAS: high I/O disk • SSD: ultra-high I/O disk • SATA: common I/O disk SATA disks have been removed from EVS. You can find them attached only on existing nodes.
extendParam	Map<String, Object>	<p>Extended disk parameters, defined in extendparam 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 dssPoolID, 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 dss. This field is used only for DeC clusters.</p>

Parameter	Type	Description
hw:passthrough	Boolean	<ul style="list-style-type: none"> Pay attention to this field if your ECS is SDI-compliant. If the value of this field is true, an SCSI disk will be created. If the node pool type is ElasticBMS, this field must be set to true. If a node specification involves local disks and EVS disks at the same time, set the disk initialization parameters. For details, see Attaching Disks to a Node.
metadata	VolumeMetadata 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 4-515 VolumeMetadata

Parameter	Type	Description
__system__encrypted	String	Whether the EVS disk is encrypted. The value 0 indicates that the EVS disk is not encrypted, and the value 1 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 metadata . This field is used with __system__encrypted .

Table 4-516 Storage

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

Table 4-517 StorageSelectors

Parameter	Type	Description
name	String	Selector name, used as the index of selectorNames in storageGroup . Therefore, the name of each selector must be unique.
storageType	String	Specifies the storage type. Currently, only evs (EVS volumes) and local (local volumes) are supported. The local storage does not support disk selection. All local disks will form a VG. Therefore, only one storageSelector of the local type is allowed.
matchLabels	matchLabels object	Matching field of an EVS volume. The size , volumeType , metadataEncrypted , metadataCmkid and count fields are supported.

Table 4-518 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. Currently, SSD, GPSSD and SAS are supported.
metadataEncrypted	String	Disk encryption identifier. 0 indicates that the disk is not encrypted, and 1 indicates that the disk is encrypted.
metadataCmkid	String	Customer master key ID of an encrypted disk. The value is a 36-byte string.
count	String	Number of disks to be selected. If this parameter is left blank, all disks of this type are selected.

Table 4-519 StorageGroups

Parameter	Type	Description
name	String	Name of a virtual storage group, which must be unique.

Parameter	Type	Description
cceManaged	Boolean	Storage space for Kubernetes and runtime components. Only one group can be set to true . If this parameter is left blank, the default value false is used.
selectorNames	Array of strings	This parameter corresponds to name in storageSelectors . A group can match multiple selectors, but a selector can match only one group.
virtualSpaces	Array of VirtualSpace objects	Detailed management of space configuration in a group.

Table 4-520 VirtualSpace

Parameter	Type	Description
name	String	Name of a virtualSpace. <ul style="list-style-type: none"> • Kubernetes: Kubernetes space configuration. lvmConfig needs to be configured. • runtime: runtime space configuration. runtimeConfig needs to be configured. • user: user space configuration. lvmConfig needs to be configured.
size	String	Size of a virtualSpace. The value must be an integer in percentage. Example: 90%. NOTE The sum of the percentages of all virtualSpaces in a group cannot exceed 100%.
lvmConfig	LVMConfig object	LVM configurations, applicable to kubernetes and user spaces. Note that one virtual space supports only one config.
runtimeConfig	RuntimeConfig object	runtime configurations, applicable to the runtime space. Note that one virtual space supports only one config.

Table 4-521 LVMConfig

Parameter	Type	Description
lvType	String	LVM write mode. linear indicates the linear mode. striped indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.
path	String	Path to which the disk is attached. This parameter takes effect only in user configuration. The value is an absolute path. Digits, letters, periods (.), hyphens (-), and underscores (_) are allowed.

Table 4-522 RuntimeConfig

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

Table 4-523 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. NOTE If ids has been set, you do not need to set count and eip .
count	Integer	Number of EIPs to be dynamically created. NOTE count and eip must be set at the same time.
eip	NodeEIPSpec object	EIP configuration.

Table 4-524 NodeEIPSpec

Parameter	Type	Description
iptype	String	EIP type, specified in publicip.type in the API for assigning an EIP. . .

Parameter	Type	Description
bandwidth	NodeBandwidth object	Bandwidth parameters of the EIP

Table 4-525 NodeBandwidth

Parameter	Type	Description
chargemode	String	<p>Bandwidth billing mode.</p> <ul style="list-style-type: none"> If this field is not specified, the billing is based on bandwidth. - If the field is null, the billing is based on bandwidth. If the field value is traffic, the billing is based on traffic. If the value is out of the preceding options, the cloud server will fail to be created. <p>NOTE</p> <ul style="list-style-type: none"> 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%. 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%.
size	Integer	Bandwidth size, specified in bandwidth.size in the API for assigning an EIP. . .
sharetype	String	Bandwidth sharing type. Value options: PER (exclusive bandwidth)

Table 4-526 NodeNicSpec

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

Table 4-527 NicSpec

Parameter	Type	Description
subnetId	String	ID of the subnet to which the NIC belongs. If subnetId is not specified when creating the primary NIC, the cluster subnet is used. When creating a secondary NIC, you must specify subnetId .
fixedIps	Array of strings	The IP address of the primary NIC is specified by fixedIps . The number of IP addresses cannot be greater than the number of created nodes. fixedIps and ipBlock cannot be specified at the same time.
ipBlock	String	CIDR format of the primary NIC IP range. The IP address of the created node falls in this range. fixedIps and ipBlock cannot be specified at the same time.

Table 4-528 Taint

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

Table 4-529 UserTag

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

Table 4-530 Runtime

Parameter	Type	Description
name	String	Container runtime. Defaults to docker .

Table 4-531 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 Maximum Number of Pods That Can Be Created on a Node.</p>
DockerLVMDockerStorageConfigOverride	String	<p>Docker data disk configuration item. (This parameter has been discarded. Use the storage field instead.) Example default configuration: <code>"DockerLVMDockerStorageConfigOverride": "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 diskType based on the actual drive letter type. The following fields are included:</p> <ul style="list-style-type: none"> • userLV (optional): size of the user space, for example, vgpaas/20%VG. • userPath (optional): mount path of the user space, for example, /home/wqt-test. • diskType: disk type. Currently, only evs, hdd, and ssd are supported. • lvType: type of a logic volume. The value can be linear or striped. • dockerThinpool: Docker space size, for example, vgpaas/60%VG. • kubernetesLV: kubelet space size, for example, vgpaas/20%VG.

Parameter	Type	Description
dockerBaseSize	Integer	<p>Available disk space of a single container on a node, in GB.</p> <p>If this parameter is left blank or is set to 0, the default value is used. In Device Mapper mode, the default value is 10. In OverlayFS mode, the available space of a single container is not limited by default, and the dockerBaseSize setting takes effect only on EulerOS nodes in the cluster of the new version.</p> <p>For details about how to allocate the space for the container runtime, see Data Disk Space Allocation.</p> <p>When Device Mapper is used, you are advised to set dockerBaseSize 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>
publicKey	String	Public key of a node.
alpha.cce/preInstall	String	<p>Pre-installation script.</p> <p>NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)</p>
alpha.cce/postInstall	String	<p>Post-installation script.</p> <p>NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)</p>
alpha.cce/NodeImageID	String	This parameter is required when a custom image is used to create a BMS node.
enterprise_project_id	String	ID of the enterprise project to which the node belongs
chargingMode	Integer	Billing mode of a node. This parameter has been deprecated. Use the billingMode parameter in NodeSpec .
agency_name	String	<p>Name of an agency</p> <p>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.</p>

Table 4-532 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 minNodeCount 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 4-533 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 4-534 SecurityID

Parameter	Type	Description
id	String	Security group ID

Table 4-535 NodePoolStatus

Parameter	Type	Description
currentNode	Integer	Total number of nodes in the current node pool (excluding the nodes that are being deleted)

Parameter	Type	Description
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.
phase	String	<p>Node pool status.</p> <ul style="list-style-type: none"> • Null: available (the number of current nodes in the node pool has reached the expected value, and no node scaling is being performed.) • 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.) • Synchronized: pending scaling (the number of current nodes in the node pool does not reach the expected value, or node scaling is being performed.) • 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.) <p>NOTE 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"> • Node pool scaling status: You can obtain accurate status of the current node pool based on parameters such as currentNode, creatingNode, or deletingNode. • Node pool scale-out: You can use conditions to obtain the detailed status of a node pool. Scalable can replace SoldOut. • Deleting: The object is being deleted. • Error: An error occurs.
jobId	String	ID of a job executed on the node pool
conditions	Array of NodePoolCondition objects	Node pool status details. For details, see the definition of Condition .

Table 4-536 NodePoolCondition

Parameter	Type	Description
type	String	Condition type. The options are as follows: <ul style="list-style-type: none"> • Scalable: whether a node pool can be scaled. If the status is False, node pool scaling will not be triggered again. • QuotaInsufficient: Quotas on which node pool scaling depends are insufficient, affecting the node pool scaling status. • ResourceInsufficient: Resources on which node pool scaling depends are insufficient, affecting the node pool scaling status. • UnexpectedError: The node pool fails to be scaled out due to unexpected reasons, affecting the node pool scaling status. • Error: A node pool error occurs. A common trigger is deletion failure.
status	String	Current status of Condition . The options are as follows: <ul style="list-style-type: none"> • "True" • "False"
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, scaling-out, or scaling-in state.

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.

NOTE

- Currently, only the node pool name, **initialNodeCount**, **k8sTags**, **taints**, **login**, and **userTags** under **spec**, and scaling-related fields can be updated. If no value is changed in this update, the initial values are used by default.

URI

PUT /api/v3/projects/{project_id}/clusters/{cluster_id}/nodepools/{nodepool_id}

Table 4-537 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
nodepool_id	Yes	String	Node pool ID.

Request Parameters

Table 4-538 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 Obtaining a User Token .

Table 4-539 Request body parameters

Parameter	Mandatory	Type	Description
metadata	Yes	NodePoolMetadataUpdate object	Metadata information of the node pool.
spec	Yes	NodePoolSpecUpdate object	Node pool specifications.

Table 4-540 NodePoolMetadataUpdate

Parameter	Mandatory	Type	Description
name	Yes	String	Name of the node pool. NOTE Naming rules: <ul style="list-style-type: none"> • Enter 1 to 50 characters, starting with a lowercase letter and not ending with a hyphen (-). Only lowercase letters, digits, and hyphens (-) are allowed. • You cannot create node pools named DefaultPool.

Table 4-541 NodePoolSpecUpdate

Parameter	Mandatory	Type	Description
nodeTemplate	Yes	NodeSpecUpdate 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 0 .
autoscaling	Yes	NodePoolNodeAutoscaling object	Auto scaling parameters. Only pay-per-use node pools support auto scaling.

Table 4-542 NodeSpecUpdate

Parameter	Mandatory	Type	Description
taints	Yes	Array of Taint objects	<p>Taints can be added for anti-affinity when creating nodes. A maximum of 20 taints can be added. This parameter is left empty by default. Each taint contains the following parameters:</p> <ul style="list-style-type: none"> • Key: A key must contain 1 to 63 characters starting with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain name can be used as the prefix of a key. • Value: A value must start with a letter or digit and can contain a maximum of 63 characters, including letters, digits, hyphens (-), underscores (_), and periods (.). • Effect: Available options are NoSchedule, PreferNoSchedule, and NoExecute. <p>Example:</p> <pre>"taints": [{ "key": "status", "value": "unavailable", "effect": "NoSchedule" }, { "key": "looks", "value": "bad", "effect": "NoSchedule" }]</pre>

Parameter	Mandatory	Type	Description
k8sTags	Yes	Map<String,String>	<p>Defined in key-value pairs. A maximum of 20 key-value pairs are allowed. This parameter is left empty by default.</p> <ul style="list-style-type: none"> • Key: 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 and contain a maximum of 253 characters. Example DNS subdomain: example.com/my-key • Value: The value can be left blank or contain 1 to 63 characters that start with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed in the character string. <p>Example:</p> <pre>"k8sTags": { "key": "value" }</pre>
userTags	Yes	Array of UserTag objects	<p>Cloud server tag. The key of a tag must be unique. The maximum number of custom tags supported by CCE depends on the region and cannot exceed 8. This parameter is left empty by default.</p>

Parameter	Mandatory	Type	Description
initializedConditions	No	Array of strings	<p>Custom initialization flag. Before CCE nodes are initialized, they are tainted with node.cloudprovider.kubernetes.io/uninitialized to prevent pods from being scheduled to them.</p> <p>CCE supports custom initialization flags. After receiving the initializedConditions parameter, CCE converts the parameter value into a node label and provisions the label with the node, for example, cloudprovider.openvessel.io/inject-initialized-conditions=CCEInitial_CustomedInitial.</p> <p>After the node is labeled, its status.Conditions is polled to check whether the type of conditions has a flag name, such as CCEInitial and CustomedInitial. If all input flags exist and their status is True, the node initialization is complete and the initialization taint is removed.</p> <p>This parameter is left empty by default.</p> <ul style="list-style-type: none"> • Use only letters and digits. Max. characters: 20. • Max. flags: 2.

Table 4-543 Taint

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

Table 4-544 UserTag

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

Table 4-545 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 minNodeCount 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.

Response Parameters

Status code: 200

Table 4-546 Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at NodePool .
apiVersion	String	API version. The value is fixed at v3 .
metadata	NodePoolMetadata object	Metadata information of the node pool
spec	NodePoolSpec object	Node pool specifications
status	NodePoolStatus object	Node pool status

Table 4-547 NodePoolMetadata

Parameter	Type	Description
name	String	Node pool name. NOTE 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"> You cannot create node pools named DefaultPool.
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
updateTimestamp	String	Update time.
creationTimestamp	String	Creation time.

Table 4-548 NodePoolSpec

Parameter	Type	Description
type	String	Node pool type. If this parameter is left blank, the value vm is used by default. <ul style="list-style-type: none"> • vm: ECS • ElasticBMS: C6 general computing-plus BMS. An example specification is c6.22xlarge.2.physical.
nodeTemplate	NodeSpec 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	NodePoolNodeAutoscaling object	Auto scaling parameters
nodeManagement	NodeManagement object	Node management configuration
podSecurityGroups	Array of SecurityID objects	Security groups configuration
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"> • Specifying no security group ID will add the new nodes to the default security group of the worker nodes. • Specifying a valid security group ID will put new nodes in that security group. • When specifying a security group, do not modify the rules of the port on which CCE running depends.

Table 4-549 NodeSpec

Parameter	Type	Description
flavor	String	Node specifications. For details about the node specifications supported by CCE, see Node Flavor Description .
az	String	Name of the AZ where the node to be created is located. . .

Parameter	Type	Description
os	String	Node OS. For details about the supported OSs, see Node OS Description . NOTE <ul style="list-style-type: none"> 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. If <code>alpha.cce/NodeImageID</code> in <code>extendParam</code> is specified during node creation, you do not need to set this field.
login	Login object	Node login mode. Either the key pair or password must be used for login.
rootVolume	Volume object	Information about disks on the node
dataVolumes	Array of Volume objects	Data disk parameters of the node. Currently, you can add the second data disk for your node on the CCE console. This data disk is used by the container runtime and kubelet. Do not uninstall this disk. Otherwise, the node will become unavailable. For DeC nodes, the parameter description is the same as that for rootVolume .
storage	Storage object	Disk initialization management parameter. This parameter is complex to configure. For details, see Attaching Disks to a Node . If this parameter retains its default, disks are managed based on the <code>DockerLVMConfigOverride</code> (discarded) parameter in <code>extendParam</code> . This parameter is supported by clusters of version 1.15.11 and later. NOTE If a node specification involves local disks and EVS disks at the same time, do not retain the default value of this parameter to prevent unexpected disk partitions.
publicIP	NodePublicIP object	EIP of a node. NOTE This parameter is not supported when you add a node to a node pool.
nodeNicSpec	NodeNicSpec 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.

Parameter	Type	Description
billingMode	Integer	Node billing mode. <ul style="list-style-type: none"> • 0: pay-per-use
taints	Array of Taint 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"> • Key: A key must contain 1 to 63 characters starting with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain name can be used as the prefix of a key. • Value: A value must start with a letter or digit and can contain a maximum of 63 characters, including letters, digits, hyphens (-), underscores (_), and periods (.). • Effect: Available options are NoSchedule, PreferNoSchedule, and NoExecute. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In using 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>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"> • Key: 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 and contain a maximum of 253 characters. Example DNS subdomain: example.com/my-key • Value: The value can be left blank or contain 1 to 63 characters that start with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed in the character string. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In using 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>Example:</p> <pre>"k8sTags": { "key": "value" }</pre>
ecsGroupId	String	Cloud server group ID. If this field is specified, the node is created in the specified cloud server group.
dedicatedHostId	String	<p>ID of the DeH to which the node is scheduled.</p> <p>NOTE This parameter is not supported when you add a node during node pool creation.</p>
userTags	Array of UserTag objects	Cloud server tag. The key of a tag must be unique. The maximum number of custom tags supported by CCE depends on the region 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 using 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.
runtime	Runtime object	Container runtime. Defaults to docker .

Parameter	Type	Description
initializedConditions	Array of strings	<p>Custom initialization flag.</p> <p>Before CCE nodes are initialized, they are tainted with node.cloudprovider.kubernetes.io/uninitialized to prevent pods from being scheduled to them.</p> <p>CCE supports custom initialization flags. After receiving the initializedConditions parameter, CCE converts the parameter value into a node label and provisions the label with the node, for example, cloudprovider.openvessel.io/inject-initialized-conditions=CCEInitial_CustomedInitial.</p> <p>After the node is labeled, its status.Conditions is polled to check whether the type of conditions has a flag name, such as CCEInitial and CustomedInitial. If all input flags exist and their status is True, the node initialization is complete and the initialization taint is removed.</p> <ul style="list-style-type: none"> • Use only letters and digits. Max. characters: 20. • Max. flags: 2.
extendParam	NodeExtendParam object	Extended parameters for creating a node.

Table 4-550 Login

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

Table 4-551 UserPassword

Parameter	Type	Description
username	String	Login account. Defaults to root .

Parameter	Type	Description
password	String	<p>If a username and a 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"> • Contains 8 to 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. The password field must be salted during node creation. For details, see Adding a Salt in the password Field When Creating a Node.

Table 4-552 Volume

Parameter	Type	Description
size	Integer	<p>Disk size, in GB.</p> <ul style="list-style-type: none"> • System disk: 40 to 1024 -Value range for data disks: 100 to 32768
volumetype	String	<p>Disk type. For details about possible values, see the description of the root_volume parameter in the API used to create an ECS in the ECS API reference.</p> <ul style="list-style-type: none"> • SAS: high I/O disk • SSD: ultra-high I/O disk • SATA: common I/O disk SATA disks have been removed from EVS. You can find them attached only on existing nodes.
extendParam	Map<String, Object>	<p>Extended disk parameters, defined in extendparam 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 dssPoolID, 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 dss. This field is used only for DeC clusters.</p>

Parameter	Type	Description
hw:passthrough	Boolean	<ul style="list-style-type: none"> Pay attention to this field if your ECS is SDI-compliant. If the value of this field is true, an SCSI disk will be created. If the node pool type is ElasticBMS, this field must be set to true. If a node specification involves local disks and EVS disks at the same time, set the disk initialization parameters. For details, see Attaching Disks to a Node.
metadata	VolumeMetadata 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 4-553 VolumeMetadata

Parameter	Type	Description
__system__encrypted	String	Whether the EVS disk is encrypted. The value 0 indicates that the EVS disk is not encrypted, and the value 1 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 metadata . This field is used with __system__encrypted .

Table 4-554 Storage

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

Table 4-555 StorageSelectors

Parameter	Type	Description
name	String	Selector name, used as the index of selectorNames in storageGroup . Therefore, the name of each selector must be unique.
storageType	String	Specifies the storage type. Currently, only evs (EVS volumes) and local (local volumes) are supported. The local storage does not support disk selection. All local disks will form a VG. Therefore, only one storageSelector of the local type is allowed.
matchLabels	matchLabels object	Matching field of an EVS volume. The size , volumeType , metadataEncrypted , metadataCmkid and count fields are supported.

Table 4-556 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. Currently, SSD, GPSSD and SAS are supported.
metadataEncrypted	String	Disk encryption identifier. 0 indicates that the disk is not encrypted, and 1 indicates that the disk is encrypted.
metadataCmkid	String	Customer master key ID of an encrypted disk. The value is a 36-byte string.
count	String	Number of disks to be selected. If this parameter is left blank, all disks of this type are selected.

Table 4-557 StorageGroups

Parameter	Type	Description
name	String	Name of a virtual storage group, which must be unique.

Parameter	Type	Description
cceManaged	Boolean	Storage space for Kubernetes and runtime components. Only one group can be set to true . If this parameter is left blank, the default value false is used.
selectorNames	Array of strings	This parameter corresponds to name in storageSelectors . A group can match multiple selectors, but a selector can match only one group.
virtualSpaces	Array of VirtualSpace objects	Detailed management of space configuration in a group.

Table 4-558 VirtualSpace

Parameter	Type	Description
name	String	Name of a virtualSpace. <ul style="list-style-type: none"> • Kubernetes: Kubernetes space configuration. lvmConfig needs to be configured. • runtime: runtime space configuration. runtimeConfig needs to be configured. • user: user space configuration. lvmConfig needs to be configured.
size	String	Size of a virtualSpace. The value must be an integer in percentage. Example: 90%. NOTE The sum of the percentages of all virtualSpaces in a group cannot exceed 100%.
lvmConfig	LVMConfig object	LVM configurations, applicable to kubernetes and user spaces. Note that one virtual space supports only one config.
runtimeConfig	RuntimeConfig object	runtime configurations, applicable to the runtime space. Note that one virtual space supports only one config.

Table 4-559 LVMConfig

Parameter	Type	Description
lvType	String	LVM write mode. linear indicates the linear mode. striped indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.
path	String	Path to which the disk is attached. This parameter takes effect only in user configuration. The value is an absolute path. Digits, letters, periods (.), hyphens (-), and underscores (_) are allowed.

Table 4-560 RuntimeConfig

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

Table 4-561 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. NOTE If ids has been set, you do not need to set count and eip .
count	Integer	Number of EIPs to be dynamically created. NOTE count and eip must be set at the same time.
eip	NodeEIPSpec object	EIP configuration.

Table 4-562 NodeEIPSpec

Parameter	Type	Description
iptype	String	EIP type, specified in publicip.type in the API for assigning an EIP. . .

Parameter	Type	Description
bandwidth	NodeBandwidth object	Bandwidth parameters of the EIP

Table 4-563 NodeBandwidth

Parameter	Type	Description
chargemode	String	Bandwidth billing mode. <ul style="list-style-type: none"> If this field is not specified, the billing is based on bandwidth. - If the field is null, the billing is based on bandwidth. If the field value is traffic, the billing is based on traffic. If the value is out of the preceding options, the cloud server will fail to be created. <p>NOTE</p> <ul style="list-style-type: none"> 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%. 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%.
size	Integer	Bandwidth size, specified in bandwidth.size in the API for assigning an EIP. . .
sharetype	String	Bandwidth sharing type. Value options: PER (exclusive bandwidth)

Table 4-564 NodeNicSpec

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

Table 4-565 NicSpec

Parameter	Type	Description
subnetId	String	ID of the subnet to which the NIC belongs. If subnetId is not specified when creating the primary NIC, the cluster subnet is used. When creating a secondary NIC, you must specify subnetId .
fixedIps	Array of strings	The IP address of the primary NIC is specified by fixedIps . The number of IP addresses cannot be greater than the number of created nodes. fixedIps and ipBlock cannot be specified at the same time.
ipBlock	String	CIDR format of the primary NIC IP range. The IP address of the created node falls in this range. fixedIps and ipBlock cannot be specified at the same time.

Table 4-566 Taint

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

Table 4-567 UserTag

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

Table 4-568 Runtime

Parameter	Type	Description
name	String	Container runtime. Defaults to docker .

Table 4-569 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 Maximum Number of Pods That Can Be Created on a Node.</p>
DockerLVMDockerConfigOverride	String	<p>Docker data disk configuration item. (This parameter has been discarded. Use the storage field instead.) Example default configuration: <code>"DockerLVMDockerConfigOverride":"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 diskType based on the actual drive letter type. The following fields are included:</p> <ul style="list-style-type: none"> • userLV (optional): size of the user space, for example, vgpaas/20%VG. • userPath (optional): mount path of the user space, for example, /home/wqt-test. • diskType: disk type. Currently, only evs, hdd, and ssd are supported. • lvType: type of a logic volume. The value can be linear or striped. • dockerThinpool: Docker space size, for example, vgpaas/60%VG. • kubernetesLV: kubelet space size, for example, vgpaas/20%VG.

Parameter	Type	Description
dockerBaseSize	Integer	<p>Available disk space of a single container on a node, in GB.</p> <p>If this parameter is left blank or is set to 0, the default value is used. In Device Mapper mode, the default value is 10. In OverlayFS mode, the available space of a single container is not limited by default, and the dockerBaseSize setting takes effect only on EulerOS nodes in the cluster of the new version.</p> <p>For details about how to allocate the space for the container runtime, see Data Disk Space Allocation.</p> <p>When Device Mapper is used, you are advised to set dockerBaseSize 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>
publicKey	String	Public key of a node.
alpha.cce/preInstall	String	<p>Pre-installation script.</p> <p>NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)</p>
alpha.cce/postInstall	String	<p>Post-installation script.</p> <p>NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)</p>
alpha.cce/NodeImageID	String	This parameter is required when a custom image is used to create a BMS node.
enterprise_project_id	String	ID of the enterprise project to which the node belongs
chargingMode	Integer	Billing mode of a node. This parameter has been deprecated. Use the billingMode parameter in NodeSpec .
agency_name	String	<p>Name of an agency</p> <p>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.</p>

Table 4-570 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 minNodeCount 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 4-571 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 4-572 SecurityID

Parameter	Type	Description
id	String	Security group ID

Table 4-573 NodePoolStatus

Parameter	Type	Description
currentNode	Integer	Total number of nodes in the current node pool (excluding the nodes that are being deleted)

Parameter	Type	Description
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.
phase	String	<p>Node pool status.</p> <ul style="list-style-type: none"> • Null: available (the number of current nodes in the node pool has reached the expected value, and no node scaling is being performed.) • 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.) • Synchronized: pending scaling (the number of current nodes in the node pool does not reach the expected value, or node scaling is being performed.) • 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.) <p>NOTE 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"> • Node pool scaling status: You can obtain accurate status of the current node pool based on parameters such as currentNode, creatingNode, or deletingNode. • Node pool scale-out: You can use conditions to obtain the detailed status of a node pool. Scalable can replace SoldOut. • Deleting: The object is being deleted. • Error: An error occurs.
jobId	String	ID of a job executed on the node pool
conditions	Array of NodePoolCondition objects	Node pool status details. For details, see the definition of Condition .

Table 4-574 NodePoolCondition

Parameter	Type	Description
type	String	Condition type. The options are as follows: <ul style="list-style-type: none"> • Scalable: whether a node pool can be scaled. If the status is False, node pool scaling will not be triggered again. • QuotaInsufficient: Quotas on which node pool scaling depends are insufficient, affecting the node pool scaling status. • ResourceInsufficient: Resources on which node pool scaling depends are insufficient, affecting the node pool scaling status. • UnexpectedError: The node pool fails to be scaled out due to unexpected reasons, affecting the node pool scaling status. • Error: A node pool error occurs. A common trigger is deletion failure.
status	String	Current status of Condition . The options are as follows: <ul style="list-style-type: none"> • "True" • "False"
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

```
{
  "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
}
}
```

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 4-575 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
cluster_id	Yes	String	Cluster ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .
nodepool_id	Yes	String	Node pool ID.

Request Parameters

Table 4-576 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 Obtaining a User Token .

Response Parameters

Status code: 200

Table 4-577 Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at NodePool .
apiVersion	String	API version. The value is fixed at v3 .
metadata	NodePoolMetadata object	Metadata information of the node pool
spec	NodePoolSpec object	Node pool specifications
status	NodePoolStatus object	Node pool status

Table 4-578 NodePoolMetadata

Parameter	Type	Description
name	String	Node pool name. NOTE 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"> You cannot create node pools named DefaultPool.
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.

Parameter	Type	Description
annotations	Map<String,String>	Annotations of a node pool in key-value pairs
updateTimestamp	String	Update time.
creationTimestamp	String	Creation time.

Table 4-579 NodePoolSpec

Parameter	Type	Description
type	String	Node pool type. If this parameter is left blank, the value vm is used by default. <ul style="list-style-type: none"> • vm: ECS • ElasticBMS: C6 general computing-plus BMS. An example specification is c6.22xlarge.2.physical.
nodeTemplate	NodeSpec 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	NodePoolNodeAutoscaling object	Auto scaling parameters
nodeManagement	NodeManagement object	Node management configuration
podSecurityGroups	Array of SecurityID objects	Security groups configuration
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"> • Specifying no security group ID will add the new nodes to the default security group of the worker nodes. • Specifying a valid security group ID will put new nodes in that security group. • When specifying a security group, do not modify the rules of the port on which CCE running depends.

Table 4-580 NodeSpec

Parameter	Type	Description
flavor	String	Node specifications. For details about the node specifications supported by CCE, see Node Flavor Description .
az	String	Name of the AZ where the node to be created is located. . .
os	String	Node OS. For details about the supported OSs, see Node OS Description . NOTE <ul style="list-style-type: none"> 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. If <code>alpha.cce/NodeImageID</code> in <code>extendParam</code> is specified during node creation, you do not need to set this field.
login	Login object	Node login mode. Either the key pair or password must be used for login.
rootVolume	Volume object	Information about disks on the node
dataVolumes	Array of Volume objects	Data disk parameters of the node. Currently, you can add the second data disk for your node on the CCE console. This data disk is used by the container runtime and kubelet. Do not uninstall this disk. Otherwise, the node will become unavailable. For DeC nodes, the parameter description is the same as that for rootVolume .
storage	Storage object	Disk initialization management parameter. This parameter is complex to configure. For details, see Attaching Disks to a Node . If this parameter retains its default, disks are managed based on the <code>DockerLVMConfigOverride (discarded)</code> parameter in <code>extendParam</code> . This parameter is supported by clusters of version 1.15.11 and later. NOTE If a node specification involves local disks and EVS disks at the same time, do not retain the default value of this parameter to prevent unexpected disk partitions.
publicIP	NodePublicIP object	EIP of a node. NOTE This parameter is not supported when you add a node to a node pool.

Parameter	Type	Description
nodeNicSpec	NodeNicSpec 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"> • 0: pay-per-use
taints	Array of Taint 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"> • Key: A key must contain 1 to 63 characters starting with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed. A DNS subdomain name can be used as the prefix of a key. • Value: A value must start with a letter or digit and can contain a maximum of 63 characters, including letters, digits, hyphens (-), underscores (_), and periods (.). • Effect: Available options are NoSchedule, PreferNoSchedule, and NoExecute. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In using 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>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"> • Key: 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 and contain a maximum of 253 characters. Example DNS subdomain: example.com/my-key • Value: The value can be left blank or contain 1 to 63 characters that start with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed in the character string. When creating a node, you can specify the initial value for this parameter and this field is not returned during query. In using 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>Example:</p> <pre>"k8sTags": { "key": "value" }</pre>
ecsGroupId	String	Cloud server group ID. If this field is specified, the node is created in the specified cloud server group.
dedicatedHostId	String	<p>ID of the DeH to which the node is scheduled.</p> <p>NOTE This parameter is not supported when you add a node during node pool creation.</p>
userTags	Array of UserTag objects	Cloud server tag. The key of a tag must be unique. The maximum number of custom tags supported by CCE depends on the region 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 using 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.
runtime	Runtime object	Container runtime. Defaults to docker .

Parameter	Type	Description
initializedConditions	Array of strings	<p>Custom initialization flag.</p> <p>Before CCE nodes are initialized, they are tainted with node.cloudprovider.kubernetes.io/uninitialized to prevent pods from being scheduled to them.</p> <p>CCE supports custom initialization flags. After receiving the initializedConditions parameter, CCE converts the parameter value into a node label and provisions the label with the node, for example, cloudprovider.openvessel.io/inject-initialized-conditions=CCEInitial_CustomedInitial.</p> <p>After the node is labeled, its status.Conditions is polled to check whether the type of conditions has a flag name, such as CCEInitial and CustomedInitial. If all input flags exist and their status is True, the node initialization is complete and the initialization taint is removed.</p> <ul style="list-style-type: none"> • Use only letters and digits. Max. characters: 20. • Max. flags: 2.
extendParam	NodeExtendParam object	Extended parameters for creating a node.

Table 4-581 Login

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

Table 4-582 UserPassword

Parameter	Type	Description
username	String	Login account. Defaults to root .

Parameter	Type	Description
password	String	<p>If a username and a 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"> • Contains 8 to 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. The password field must be salted during node creation. For details, see Adding a Salt in the password Field When Creating a Node.

Table 4-583 Volume

Parameter	Type	Description
size	Integer	<p>Disk size, in GB.</p> <ul style="list-style-type: none"> • System disk: 40 to 1024 -Value range for data disks: 100 to 32768
volumetype	String	<p>Disk type. For details about possible values, see the description of the root_volume parameter in the API used to create an ECS in the ECS API reference.</p> <ul style="list-style-type: none"> • SAS: high I/O disk • SSD: ultra-high I/O disk • SATA: common I/O disk SATA disks have been removed from EVS. You can find them attached only on existing nodes.
extendParam	Map<String, Object>	<p>Extended disk parameters, defined in extendparam 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 dssPoolID, 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 dss. This field is used only for DeC clusters.</p>

Parameter	Type	Description
hw:passthrough	Boolean	<ul style="list-style-type: none"> Pay attention to this field if your ECS is SDI-compliant. If the value of this field is true, an SCSI disk will be created. If the node pool type is ElasticBMS, this field must be set to true. If a node specification involves local disks and EVS disks at the same time, set the disk initialization parameters. For details, see Attaching Disks to a Node.
metadata	VolumeMetadata 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 4-584 VolumeMetadata

Parameter	Type	Description
__system__encrypted	String	Whether the EVS disk is encrypted. The value 0 indicates that the EVS disk is not encrypted, and the value 1 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 metadata . This field is used with __system__encrypted .

Table 4-585 Storage

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

Table 4-586 StorageSelectors

Parameter	Type	Description
name	String	Selector name, used as the index of selectorNames in storageGroup . Therefore, the name of each selector must be unique.
storageType	String	Specifies the storage type. Currently, only evs (EVS volumes) and local (local volumes) are supported. The local storage does not support disk selection. All local disks will form a VG. Therefore, only one storageSelector of the local type is allowed.
matchLabels	matchLabels object	Matching field of an EVS volume. The size , volumeType , metadataEncrypted , metadataCmkid and count fields are supported.

Table 4-587 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. Currently, SSD, GPSSD and SAS are supported.
metadataEncrypted	String	Disk encryption identifier. 0 indicates that the disk is not encrypted, and 1 indicates that the disk is encrypted.
metadataCmkid	String	Customer master key ID of an encrypted disk. The value is a 36-byte string.
count	String	Number of disks to be selected. If this parameter is left blank, all disks of this type are selected.

Table 4-588 StorageGroups

Parameter	Type	Description
name	String	Name of a virtual storage group, which must be unique.

Parameter	Type	Description
cceManaged	Boolean	Storage space for Kubernetes and runtime components. Only one group can be set to true . If this parameter is left blank, the default value false is used.
selectorNames	Array of strings	This parameter corresponds to name in storageSelectors . A group can match multiple selectors, but a selector can match only one group.
virtualSpaces	Array of VirtualSpace objects	Detailed management of space configuration in a group.

Table 4-589 VirtualSpace

Parameter	Type	Description
name	String	Name of a virtualSpace. <ul style="list-style-type: none"> • Kubernetes: Kubernetes space configuration. lvmConfig needs to be configured. • runtime: runtime space configuration. runtimeConfig needs to be configured. • user: user space configuration. lvmConfig needs to be configured.
size	String	Size of a virtualSpace. The value must be an integer in percentage. Example: 90%. NOTE The sum of the percentages of all virtualSpaces in a group cannot exceed 100%.
lvmConfig	LVMConfig object	LVM configurations, applicable to kubernetes and user spaces. Note that one virtual space supports only one config.
runtimeConfig	RuntimeConfig object	runtime configurations, applicable to the runtime space. Note that one virtual space supports only one config.

Table 4-590 LVMConfig

Parameter	Type	Description
lvType	String	LVM write mode. linear indicates the linear mode. striped indicates the striped mode, in which multiple disks are used to form a strip to improve disk performance.
path	String	Path to which the disk is attached. This parameter takes effect only in user configuration. The value is an absolute path. Digits, letters, periods (.), hyphens (-), and underscores (_) are allowed.

Table 4-591 RuntimeConfig

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

Table 4-592 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. NOTE If ids has been set, you do not need to set count and eip .
count	Integer	Number of EIPs to be dynamically created. NOTE count and eip must be set at the same time.
eip	NodeEIPSpec object	EIP configuration.

Table 4-593 NodeEIPSpec

Parameter	Type	Description
iptype	String	EIP type, specified in publicip.type in the API for assigning an EIP. . .

Parameter	Type	Description
bandwidth	NodeBandwidth object	Bandwidth parameters of the EIP

Table 4-594 NodeBandwidth

Parameter	Type	Description
chargemode	String	Bandwidth billing mode. <ul style="list-style-type: none"> If this field is not specified, the billing is based on bandwidth. - If the field is null, the billing is based on bandwidth. If the field value is traffic, the billing is based on traffic. If the value is out of the preceding options, the cloud server will fail to be created. <p>NOTE</p> <ul style="list-style-type: none"> 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%. 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%.
size	Integer	Bandwidth size, specified in bandwidth.size in the API for assigning an EIP. . .
sharetype	String	Bandwidth sharing type. Value options: PER (exclusive bandwidth)

Table 4-595 NodeNicSpec

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

Table 4-596 NicSpec

Parameter	Type	Description
subnetId	String	ID of the subnet to which the NIC belongs. If subnetId is not specified when creating the primary NIC, the cluster subnet is used. When creating a secondary NIC, you must specify subnetId .
fixedIps	Array of strings	The IP address of the primary NIC is specified by fixedIps . The number of IP addresses cannot be greater than the number of created nodes. fixedIps and ipBlock cannot be specified at the same time.
ipBlock	String	CIDR format of the primary NIC IP range. The IP address of the created node falls in this range. fixedIps and ipBlock cannot be specified at the same time.

Table 4-597 Taint

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

Table 4-598 UserTag

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

Table 4-599 Runtime

Parameter	Type	Description
name	String	Container runtime. Defaults to docker .

Table 4-600 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 Maximum Number of Pods That Can Be Created on a Node.</p>
DockerLVMDockerConfigOverride	String	<p>Docker data disk configuration item. (This parameter has been discarded. Use the storage field instead.) Example default configuration: <code>"DockerLVMDockerConfigOverride": "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 diskType based on the actual drive letter type. The following fields are included:</p> <ul style="list-style-type: none"> • userLV (optional): size of the user space, for example, vgpaas/20%VG. • userPath (optional): mount path of the user space, for example, /home/wqt-test. • diskType: disk type. Currently, only evs, hdd, and ssd are supported. • lvType: type of a logic volume. The value can be linear or striped. • dockerThinpool: Docker space size, for example, vgpaas/60%VG. • kubernetesLV: kubelet space size, for example, vgpaas/20%VG.

Parameter	Type	Description
dockerBaseSize	Integer	<p>Available disk space of a single container on a node, in GB.</p> <p>If this parameter is left blank or is set to 0, the default value is used. In Device Mapper mode, the default value is 10. In OverlayFS mode, the available space of a single container is not limited by default, and the dockerBaseSize setting takes effect only on EulerOS nodes in the cluster of the new version.</p> <p>For details about how to allocate the space for the container runtime, see Data Disk Space Allocation.</p> <p>When Device Mapper is used, you are advised to set dockerBaseSize 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>
publicKey	String	Public key of a node.
alpha.cce/preInstall	String	<p>Pre-installation script.</p> <p>NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)</p>
alpha.cce/postInstall	String	<p>Post-installation script.</p> <p>NOTE The input value must be Base64-encoded. (Command: echo -n "Content to be encoded" base64)</p>
alpha.cce/NodeImageID	String	This parameter is required when a custom image is used to create a BMS node.
enterprise_project_id	String	ID of the enterprise project to which the node belongs
chargingMode	Integer	Billing mode of a node. This parameter has been deprecated. Use the billingMode parameter in NodeSpec .
agency_name	String	<p>Name of an agency</p> <p>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.</p>

Table 4-601 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 minNodeCount 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 4-602 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 4-603 SecurityID

Parameter	Type	Description
id	String	Security group ID

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)

Parameter	Type	Description
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.
phase	String	<p>Node pool status.</p> <ul style="list-style-type: none"> • Null: available (the number of current nodes in the node pool has reached the expected value, and no node scaling is being performed.) • 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.) • Synchronized: pending scaling (the number of current nodes in the node pool does not reach the expected value, or node scaling is being performed.) • 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.) <p>NOTE 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"> • Node pool scaling status: You can obtain accurate status of the current node pool based on parameters such as currentNode, creatingNode, or deletingNode. • Node pool scale-out: You can use conditions to obtain the detailed status of a node pool. Scalable can replace SoldOut. • Deleting: The object is being deleted. • Error: An error occurs.
jobId	String	ID of a job executed on the node pool
conditions	Array of NodePoolCondition objects	Node pool status details. For details, see the definition of Condition .

Table 4-605 NodePoolCondition

Parameter	Type	Description
type	String	Condition type. The options are as follows: <ul style="list-style-type: none"> • Scalable: whether a node pool can be scaled. If the status is False, node pool scaling will not be triggered again. • QuotaInsufficient: Quotas on which node pool scaling depends are insufficient, affecting the node pool scaling status. • ResourceInsufficient: Resources on which node pool scaling depends are insufficient, affecting the node pool scaling status. • UnexpectedError: The node pool fails to be scaled out due to unexpected reasons, affecting the node pool scaling status. • Error: A node pool error occurs. A common trigger is deletion failure.
status	String	Current status of Condition . The options are as follows: <ul style="list-style-type: none"> • "True" • "False"
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](#).

4.5 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 4-606 Path Parameters

Parameter	Mandatory	Type	Description
namespace	Yes	String	Namespace where the PVC is located. Constraints: <ul style="list-style-type: none"> • Custom namespace. You must create a namespace in the cluster before using it. • The namespace provided by the system is namespace default. • Namespaces kube-system and kube-public cannot be used.

Request Parameters

Table 4-607 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 Obtaining a User Token .
X-Cluster-ID	No	String	Cluster ID. This field is mandatory when the URL format https://Endpoint/uri is used. For details about how to obtain the value, see How to Obtain Parameters in the API URI .

Table 4-608 Request body parameters

Parameter	Mandatory	Type	Description
apiVersion	Yes	String	API version. The value is fixed at v1 .
kind	Yes	String	API type. The value is fixed at PersistentVolumeClaim .
metadata	Yes	PersistentVolumeClaimMetadata object	Cluster object metadata, which is a collection of attributes.
spec	Yes	PersistentVolumeClaimSpec object	Detailed description of the cluster object. CCE creates or updates objects by defining or updating spec .
status	No	PersistentVolumeClaimStatus object	PVC status. The request to create a PVC does not need to carry this parameter.

Table 4-609 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"> • Key: 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 and contain a maximum of 253 characters. Example DNS subdomain: example.com/my-key • Value: The value can be left blank or contain 1 to 63 characters that start with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed in the character string.

Table 4-610 PersistentVolumeClaimSpec

Parameter	Mandatory	Type	Description
volumeID	Yes	String	ID of an existing storage volume. <ul style="list-style-type: none"> • If an SFS, EVS, or SFS Turbo volume is used, set this parameter to the ID of the volume. • If an OBS bucket is used, set this parameter to the OBS bucket name.

Parameter	Mandatory	Type	Description
storageType	Yes	String	Cloud storage class. This parameter is used together with volumeID . That is, volumeID and storageType must be configured at the same time. <ul style="list-style-type: none"> • bs: EVS • nfs: SFS • obs: OBS • efs: SFS Turbo
accessModes	Yes	Array of strings	Access mode of the volume. Only the first value in all selected options is valid. <ul style="list-style-type: none"> • ReadWriteOnce: The volume can be mounted as read-write by a single node. <p>NOTE 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"> • ReadOnlyMany (default): The volume can be mounted as read-only by many nodes. • ReadWriteMany: The volume can be mounted as read-write by many nodes.
storageClassName	No	String	Storage class name of the PVC.
volumeName	No	String	Name of the PV bound to the PVC.
resources	No	ResourceRequirements object	Resource requests and limits.
volumeMode	No	String	PV type specified by the PVC.

Table 4-611 ResourceRequirements

Parameter	Mandatory	Type	Description
limits	No	Map<String,String>	Resource limits. This parameter is invalid during creation.

Parameter	Mandatory	Type	Description
requests	No	Map<String,String>	Resource requests. This parameter is invalid during creation.

Table 4-612 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 4-613 Response body parameters

Parameter	Type	Description
apiVersion	String	API version. The value is fixed at v1 .
kind	String	API type. The value is fixed at PersistentVolumeClaim .
metadata	PersistentVolumeClaimMetadata object	Cluster object metadata, which is a collection of attributes.
spec	PersistentVolumeClaimSpec object	Detailed description of the cluster object. CCE creates or updates objects by defining or updating spec .
status	PersistentVolumeClaimStatus object	PVC status. The request to create a PVC does not need to carry this parameter.

Table 4-614 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"> • Key: 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 and contain a maximum of 253 characters. Example DNS subdomain: example.com/my-key • Value: The value can be left blank or contain 1 to 63 characters that start with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed in the character string.

Table 4-615 PersistentVolumeClaimSpec

Parameter	Type	Description
volumeID	String	ID of an existing storage volume. <ul style="list-style-type: none"> • If an SFS, EVS, or SFS Turbo volume is used, set this parameter to the ID of the volume. • If an OBS bucket is used, set this parameter to the OBS bucket name.
storageType	String	Cloud storage class. This parameter is used together with volumeID . That is, volumeID and storageType must be configured at the same time. <ul style="list-style-type: none"> • bs: EVS • nfs: SFS • obs: OBS • efs: SFS Turbo

Parameter	Type	Description
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"> • ReadWriteOnce: The volume can be mounted as read-write by a single node. <p>NOTE 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"> • ReadOnlyMany (default): The volume can be mounted as read-only by many nodes. • ReadWriteMany: The volume can be mounted as read-write by many nodes.
storageClassName	String	Storage class name of the PVC.
volumeName	String	Name of the PV bound to the PVC.
resources	ResourceRequirements object	Resource requests and limits.
volumeMode	String	PV type specified by the PVC.

Table 4-616 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 4-617 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 4-618 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 4-619 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 true (delete) or false (not delete). The value defaults to false .
storageType	No	String	Cloud storage type. This parameter is used together with deleteVolume. That is, deleteVolume and storageType must be configured at the same time. <ul style="list-style-type: none"> • bs: EVS • nfs: SFS • obs: OBS • efs: SFS Turbo

Request Parameters

Table 4-620 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 Obtaining a User Token .
X-Cluster-ID	No	String	Cluster ID. This field is mandatory when the URL format https://Endpoint/uri is used. For details about how to obtain the value, see How to Obtain Parameters in the API URI .

Response Parameters

Status code: 200

Table 4-621 Response body parameters

Parameter	Type	Description
apiVersion	String	API version. The value is fixed at v1 .
kind	String	API type. The value is fixed at PersistentVolumeClaim .
metadata	PersistentVolumeClaimMetadata object	Cluster object metadata, which is a collection of attributes.
spec	PersistentVolumeClaimSpec object	Detailed description of the cluster object. CCE creates or updates objects by defining or updating spec .
status	PersistentVolumeClaimStatus object	PVC status. The request to create a PVC does not need to carry this parameter.

Table 4-622 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"> • Key: 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 and contain a maximum of 253 characters. Example DNS subdomain: example.com/my-key • Value: The value can be left blank or contain 1 to 63 characters that start with a letter or digit. Only letters, digits, hyphens (-), underscores (_), and periods (.) are allowed in the character string.

Table 4-623 PersistentVolumeClaimSpec

Parameter	Type	Description
volumeID	String	ID of an existing storage volume. <ul style="list-style-type: none"> • If an SFS, EVS, or SFS Turbo volume is used, set this parameter to the ID of the volume. • If an OBS bucket is used, set this parameter to the OBS bucket name.
storageType	String	Cloud storage class. This parameter is used together with volumeID . That is, volumeID and storageType must be configured at the same time. <ul style="list-style-type: none"> • bs: EVS • nfs: SFS • obs: OBS • efs: SFS Turbo

Parameter	Type	Description
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"> • ReadWriteOnce: The volume can be mounted as read-write by a single node. <p>NOTE 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"> • ReadOnlyMany (default): The volume can be mounted as read-only by many nodes. • ReadWriteMany: The volume can be mounted as read-write by many nodes.
storageClassName	String	Storage class name of the PVC.
volumeName	String	Name of the PV bound to the PVC.
resources	ResourceRequirements object	Resource requests and limits.
volumeMode	String	PV type specified by the PVC.

Table 4-624 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 4-625 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 4-626 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 Obtaining a User Token .

Table 4-627 Request body parameters

Parameter	Mandatory	Type	Description
kind	Yes	String	API type. The value is fixed at Addon and cannot be changed. Any user-defined value is invalid.
apiVersion	Yes	String	API version. The value is fixed at v3 and cannot be changed. Any user-defined value is invalid.
metadata	Yes	Metadata object	Basic information about the object. Metadata is a collection of attributes.
spec	Yes	InstanceRequestSpec object	Detailed description of add-on installation or upgrade.

Table 4-628 Metadata

Parameter	Mandatory	Type	Description
uid	No	String	Unique ID.
name	No	String	Add-on name.
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"> For add-on installation, the value is fixed at {"addon.install/type":"install"}. For add-on upgrade, the value is fixed at {"addon.upgrade/type":"upgrade"}.
updateTimestamp	No	String	Update time.
creationTimestamp	No	String	Creation time.

Table 4-629 InstanceRequestSpec

Parameter	Mandatory	Type	Description
version	No	String	Version of the add-on to be installed or upgraded, for example, 1.0.0 . If not specified, the stable version supported by the cluster is used.
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, coredns .

Response Parameters

Status code: 201

Table 4-630 Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at Addon and cannot be changed.
apiVersion	String	API version. The value is fixed at v3 and cannot be changed.
metadata	Metadata object	Basic information about the object. Metadata is a collection of attributes.
spec	InstanceSpec object	Detailed description of the add-on instance.

Parameter	Type	Description
status	AddonInstanceStatus object	Add-on instance status.

Table 4-631 Metadata

Parameter	Type	Description
uid	String	Unique ID.
name	String	Add-on name.
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"> For add-on installation, the value is fixed at {"addon.install/type":"install"}. For add-on upgrade, the value is fixed at {"addon.upgrade/type":"upgrade"}.
updateTimestamp	String	Update time.
creationTimestamp	String	Creation time.

Table 4-632 InstanceSpec

Parameter	Type	Description
clusterID	String	Cluster ID.
version	String	Add-on template version, for example, 1.0.0 .
addonTemplateName	String	Add-on template name, for example, coredns .
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.

Parameter	Type	Description
values	Map<String, Object>	Add-on template installation parameters (varying depending on the add-on). Set the parameters accordingly.

Table 4-633 AddonInstanceStatus

Parameter	Type	Description
status	String	Add-on instance status.
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	Versions object	Information about the current add-on version.

Table 4-634 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 SupportVersions objects	Cluster versions that support the add-on.
creationTimestamp	String	Creation time.
updateTimestamp	String	Update time.

Table 4-635 SupportVersions

Parameter	Type	Description
clusterType	String	Cluster type that supports the add-on.

Parameter	Type	Description
clusterVersion	Array of strings	Cluster versions that support the add-on. The value is a regular expression.

Example Requests

```
{
  "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",
    "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.",

```


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 4-636 Query Parameters

Parameter	Mandatory	Type	Description
addon_template_name	No	String	Name of the specified template. If this parameter is left unspecified, all templates are queried.

Request Parameters

Table 4-637 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 Obtaining a User Token .

Response Parameters

Status code: 200

Table 4-638 Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at Addon and cannot be changed.
apiVersion	String	API version. The value is fixed at v3 and cannot be changed.
items	Array of AddonTemplate objects	List of add-on templates.

Table 4-639 AddonTemplate

Parameter	Type	Description
kind	String	API type. The value is fixed at Addon and cannot be changed.
apiVersion	String	API version. The value is fixed at v3 and cannot be changed.
metadata	Metadata object	Basic information about the object. Metadata is a collection of attributes.
spec	TemplateSpec object	Detailed description of the add-on template.

Table 4-640 Metadata

Parameter	Type	Description
uid	String	Unique ID.
name	String	Add-on name.
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"> For add-on installation, the value is fixed at {"addon.install/type":"install"}. For add-on upgrade, the value is fixed at {"addon.upgrade/type":"upgrade"}.
updateTimestamp	String	Update time.

Parameter	Type	Description
creationTimes tamp	String	Creation time.

Table 4-641 Templatespec

Parameter	Type	Description
type	String	Template type. The value is helm or static .
require	Boolean	Whether the add-on is mandatory.
labels	Array of strings	Group to which the template belongs.
logoURL	String	URL of the logo image.
readmeURL	String	Add-on details and usage description.
description	String	Description
versions	Array of Versions objects	Template version details.

Table 4-642 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 SupportVersions objects	Cluster versions that support the add-on.
creationTimes tamp	String	Creation time.
updateTimestamp	String	Update time.

Table 4-643 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",
      "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 4-644 Path Parameters

Parameter	Mandatory	Type	Description
id	Yes	String	Add-on instance ID.

Request Parameters

Table 4-645 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 Obtaining a User Token .

Table 4-646 Request body parameters

Parameter	Mandatory	Type	Description
kind	Yes	String	API type. The value is fixed at Addon and cannot be changed. Any user-defined value is invalid.
apiVersion	Yes	String	API version. The value is fixed at v3 and cannot be changed. Any user-defined value is invalid.
metadata	Yes	Metadata object	Basic information about the object. Metadata is a collection of attributes.
spec	Yes	InstanceRequestSpec object	Detailed description of add-on installation or upgrade.

Table 4-647 Metadata

Parameter	Mandatory	Type	Description
uid	No	String	Unique ID.
name	No	String	Add-on name.
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"> For add-on installation, the value is fixed at {"addon.install/type":"install"}. For add-on upgrade, the value is fixed at {"addon.upgrade/type":"upgrade"}.
updateTimestamp	No	String	Update time.
creationTimestamp	No	String	Creation time.

Table 4-648 InstanceRequestSpec

Parameter	Mandatory	Type	Description
version	No	String	Version of the add-on to be installed or upgraded, for example, 1.0.0 . If not specified, the stable version supported by the cluster is used.
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, coredns .

Response Parameters

Status code: 200

Table 4-649 Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at Addon and cannot be changed.
apiVersion	String	API version. The value is fixed at v3 and cannot be changed.
metadata	Metadata object	Basic information about the object. Metadata is a collection of attributes.
spec	InstanceSpec object	Detailed description of the add-on instance.

Parameter	Type	Description
status	AddonInstanceStatus object	Add-on instance status.

Table 4-650 Metadata

Parameter	Type	Description
uid	String	Unique ID.
name	String	Add-on name.
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"> For add-on installation, the value is fixed at {"addon.install/type":"install"}. For add-on upgrade, the value is fixed at {"addon.upgrade/type":"upgrade"}.
updateTimestamp	String	Update time.
creationTimestamp	String	Creation time.

Table 4-651 InstanceSpec

Parameter	Type	Description
clusterID	String	Cluster ID.
version	String	Add-on template version, for example, 1.0.0 .
addonTemplateName	String	Add-on template name, for example, coredns .
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.

Parameter	Type	Description
values	Map<String, Object>	Add-on template installation parameters (varying depending on the add-on). Set the parameters accordingly.

Table 4-652 AddonInstanceStatus

Parameter	Type	Description
status	String	Add-on instance status.
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	Versions object	Information about the current add-on version.

Table 4-653 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 SupportVersions objects	Cluster versions that support the add-on.
creationTimestamp	String	Creation time.
updateTimestamp	String	Update time.

Table 4-654 SupportVersions

Parameter	Type	Description
clusterType	String	Cluster type that supports the add-on.

Parameter	Type	Description
clusterVersion	Array of strings	Cluster versions that support the add-on. The value is a regular expression.

Example Requests

```
{
  "kind": "Addon",
  "apiVersion": "v3",
  "metadata": {
    "annotations": {
      "addon.upgrade/type": "upgrade"
    }
  },
  "spec": {
    "clusterID": "1b2ec02d-a3b2-11ec-b0d0-0255ac100099",
    "version": "1.2.10",
    "addonTemplateName": "gpu-beta",
    "values": {
      "basic": {
        "device_version": "1.2.10",
        "driver_version": "1.2.10",
        "obs_url": "*****",
        "region": "*****",
        "swr_addr": "*****",
        "swr_user": "hwofficial",
        "rbac_enabled": true
      }
    },
    "custom": {
      "is_driver_from_nvidia": true,
      "nvidia_driver_download_url": "https://us.download.nvidia.com/tesla/396.37/NVIDIA-Linux-
x86_64-396.37.run"
    }
  }
}
```

Example Responses

Status code: 200

OK

```
{
  "kind": "Addon",
  "apiVersion": "v3",
  "metadata": {
    "uid": "684fa9b2-a987-11ec-ba79-0255ac100096",
    "name": "gpu-beta",
    "creationTimestamp": "2022-03-22T02:25:57Z",
    "updateTimestamp": "2022-03-22T02:25:57Z"
  },
  "spec": {
    "clusterID": "1b2ec02d-a3b2-11ec-b0d0-0255ac100099",
    "version": "1.2.10",
    "addonTemplateName": "gpu-beta",
    "addonTemplateType": "helm",
    "addonTemplateLogo": "*****",
    "addonTemplateLabels": [ "Accelerator" ],
    "description": "A device plugin for nvidia.com/gpu resource on nvidia driver",
    "values": {
      "basic": {
        "device_version": "1.2.10",

```

```

"driver_version": "1.2.10",
"obs_url": "*****",
"rbac_enabled": true,
"region": "*****",
"swr_addr": "*****",
"swr_user": "hwofficial"
},
"custom": {
  "is_driver_from_nvidia": true,
  "nvidia_driver_download_url": "https://us.download.nvidia.com/tesla/396.37/NVIDIA-Linux-
x86_64-396.37.run"
}
},
"status": {
  "status": "upgrading",
  "Reason": "addon upgrading",
  "message": "",
  "targetVersions": null,
  "currentVersion": {
    "version": "1.2.10",
    "input": {
      "basic": {
        "device_version": "1.2.10",
        "driver_version": "1.2.10",
        "obs_url": "*****",
        "region": "*****",
        "swr_addr": "*****",
        "swr_user": "hwofficial"
      }
    }
  },
  "parameters": {
    "custom": {
      "is_driver_from_nvidia": true,
      "nvidia_driver_download_url": ""
    }
  }
},
"stable": true,
"translate": {
  "en_US": {
    "addon": {
      "changeLog": "Supported GPU driver of a new version for CentOS.",
      "description": "A device plugin for nvidia.com/gpu resource on nvidia driver"
    },
    "description": {
      "Parameters.custom.drivers_info.cuda": "CUDA Toolkit",
      "Parameters.custom.drivers_info.product": "Product",
      "Parameters.custom.drivers_info.product_series": "Product Series",
      "Parameters.custom.drivers_info.product_type": "Product Type",
      "Parameters.custom.nvidia_driver_download_url": "Download the nvidia driver accroding to the
input link"
    },
    "key": {
      "Parameters.custom.nvidia_driver_download_url": "Nvidia Driver"
    }
  },
  "fr_FR": {
    "addon": {
      "changeLog": "Pilote GPU pris en charge d'une nouvelle version pour CentOS.",
      "description": "Un plug-in de dispositif pour les ressources GPU sur un pilote NVIDIA."
    },
    "description": {
      "Parameters.custom.drivers_info.cuda": "Boîte à outils CUDA",
      "Parameters.custom.drivers_info.product": "Produit",
      "Parameters.custom.drivers_info.product_series": "Serie de produits",
      "Parameters.custom.drivers_info.product_type": "type de produit",
      "Parameters.custom.nvidia_driver_download_url": "Téléchargez le pilote nvidia accroding sur le lien
d'entrée"
    }
  },

```



```

    "key" : {
      "Parameters.custom.nvidia_driver_download_url" : "Nvidia Driver"
    }
  },
  "zh_CN" : {
    "addon" : {
      "changeLog" : "",
      "description" : ""
    },
    "description" : {
      "Parameters.custom.drivers_info.cuda" : "CUDA Toolkit",
      "Parameters.custom.drivers_info.product" : "",
      "Parameters.custom.drivers_info.product_series" : "",
      "Parameters.custom.drivers_info.product_type" : "",
      "Parameters.custom.nvidia_driver_download_url" : ""
    },
    "key" : {
      "Parameters.custom.nvidia_driver_download_url" : ""
    }
  }
},
"supportVersions" : null,
"creationTimestamp" : "2022-01-11T14:57:44Z",
"updateTimestamp" : "2022-01-11T15:04:37Z"
}
}
}

```

Status Codes

Status Code	Description
200	OK

Error Codes

See [Error Codes](#).

4.6.4 Deleting an Add-on Instance

Function

This API is used to delete an add-on instance.

URI

DELETE /api/v3/addons/{id}

Table 4-655 Path Parameters

Parameter	Mandatory	Type	Description
id	Yes	String	Add-on instance ID.

Table 4-656 Query Parameters

Parameter	Mandatory	Type	Description
cluster_id	No	String	Cluster ID. For details about how to obtain a cluster ID, see How to Obtain Parameters in the API URI .

Request Parameters

Table 4-657 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 Obtaining a User Token .

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](#).

4.6.5 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 4-658 Path Parameters

Parameter	Mandatory	Type	Description
id	Yes	String	Add-on instance ID.

Table 4-659 Query Parameters

Parameter	Mandatory	Type	Description
cluster_id	No	String	Cluster ID. For details about how to obtain a cluster ID, see How to Obtain Parameters in the API URI .

Request Parameters

Table 4-660 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 Obtaining a User Token .

Response Parameters

Status code: 200

Table 4-661 Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at Addon and cannot be changed.
apiVersion	String	API version. The value is fixed at v3 and cannot be changed.
metadata	Metadata object	Basic information about the object. Metadata is a collection of attributes.
spec	InstanceSpec object	Detailed description of the add-on instance.
status	AddonInstanceStatus object	Add-on instance status.

Table 4-662 Metadata

Parameter	Type	Description
uid	String	Unique ID.
name	String	Add-on name.
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"> For add-on installation, the value is fixed at {"addon.install/type":"install"}. For add-on upgrade, the value is fixed at {"addon.upgrade/type":"upgrade"}.
updateTimestamp	String	Update time.
creationTimestamp	String	Creation time.

Table 4-663 InstanceSpec

Parameter	Type	Description
clusterID	String	Cluster ID.
version	String	Add-on template version, for example, 1.0.0 .
addonTemplate	String	Add-on template name, for example, coredns .
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 4-664 AddonInstanceStatus

Parameter	Type	Description
status	String	Add-on instance status.
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	Versions object	Information about the current add-on version.

Table 4-665 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.

Parameter	Type	Description
supportVersions	Array of SupportVersions objects	Cluster versions that support the add-on.
creationTimestamp	String	Creation time.
updateTimestamp	String	Update time.

Table 4-666 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",
    "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](#).

4.6.6 Listing Add-on Instances

Function

This API is used to list all add-on instances in the cluster.

URI

GET /api/v3/addons

Table 4-667 Query Parameters

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	Cluster ID. For details about how to obtain a cluster ID, see How to Obtain Parameters in the API URI .

Request Parameters

Table 4-668 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 Obtaining a User Token .

Response Parameters

Status code: 200

Table 4-669 Response body parameters

Parameter	Type	Description
kind	String	API type. The value is fixed at Addon and cannot be changed.
apiVersion	String	API version. The value is fixed at v3 and cannot be changed.
items	Array of AddonInstance objects	Add-on instance list.

Table 4-670 AddonInstance

Parameter	Type	Description
kind	String	API type. The value is fixed at Addon and cannot be changed.
apiVersion	String	API version. The value is fixed at v3 and cannot be changed.
metadata	Metadata object	Basic information about the object. Metadata is a collection of attributes.
spec	InstanceSpec object	Detailed description of the add-on instance.
status	AddonInstanceStatus object	Add-on instance status.

Table 4-671 Metadata

Parameter	Type	Description
uid	String	Unique ID.
name	String	Add-on name.
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"> For add-on installation, the value is fixed at {"addon.install/type":"install"}. For add-on upgrade, the value is fixed at {"addon.upgrade/type":"upgrade"}.
updateTimestamp	String	Update time.
creationTimestamp	String	Creation time.

Table 4-672 InstanceSpec

Parameter	Type	Description
clusterID	String	Cluster ID.
version	String	Add-on template version, for example, 1.0.0 .

Parameter	Type	Description
addonTemplate	String	Add-on template name, for example, coredns .
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 4-673 AddonInstanceStatus

Parameter	Type	Description
status	String	Add-on instance status.
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	Versions object	Information about the current add-on version.

Table 4-674 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 SupportVersions objects	Cluster versions that support the add-on.

Parameter	Type	Description
creationTimes tamp	String	Creation time.
updateTimest amp	String	Update time.

Table 4-675 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",
      "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
        }
      }
    }
  } ],
}
```


4.7.1 Querying Resource Quotas

Function

This API is used to query CCE resource quotas.

URI

GET /api/v3/projects/{project_id}/quotas

Table 4-676 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the value, see How to Obtain Parameters in the API URI .

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 Obtaining a User Token .

Response Parameters

Status code: 200

Table 4-678 Response body parameters

Parameter	Type	Description
quotas	Array of QuotaResource objects	Resources

Table 4-679 QuotaResource

Parameter	Type	Description
quotaKey	String	Resource type
quotaLimit	Integer	Quota
used	Integer	Number of created resources
unit	String	Unit
regionId	String	Region ID. This parameter is not returned if not involved.
availabilityZoneId	String	AZ ID. This parameter is not returned if not involved.

Example Requests

None

Example Responses

Status code: 200

The resource quotas are obtained successfully.

```
{
  "quotas": [ {
    "quotaKey": "cluster",
    "quotaLimit": 20,
    "used": 13,
    "unit": "count"
  } ]
}
```

Status Codes

Status Code	Description
200	The resource quotas are obtained successfully.

Error Codes

See [Error Codes](#).

4.8 API Versions

4.8.1 Listing API Versions

Function

This API is used to query the list of API versions supported by CCE.

URI

GET /

Request Parameters

None

Response Parameters

Status code: 200

Table 4-680 Response body parameters

Parameter	Type	Description
versions	Array of APIVersionDetail objects	API version list.

Table 4-681 APIVersionDetail

Parameter	Type	Description
id	String	API version ID. Example: v3
links	Array of APIVersionLink objects	URL of the API version.
min_version	String	Microversion. If the APIs of this version support microversions, set this parameter to the supported minimum microversion. If no microversion is supported, leave this parameter blank.
status	String	API version status. Possible values: <ul style="list-style-type: none"> ● CURRENT: preferred API version ● SUPPORTED: old API version that is still supported ● DEPRECATED: discarded API version that will be deleted

Parameter	Type	Description
updated	String	Time when the API was released (UTC). For example, if the API version is v3, the value is '2018-09-15 00:00:00Z'.
version	String	Microversion. If the APIs of this version support microversions, set this parameter to the supported maximum microversion. If no microversion is supported, leave this parameter blank.

Table 4-682 APIVersionLink

Parameter	Type	Description
href	String	URL of the API version.
rel	String	Link attributes. self : A self link contains a versioned link to the resource. Use these links immediately after linking.

Example Requests

None

Example Responses

Status code: 200

The API version list is queried successfully.

```
{
  "versions": [ {
    "id": "v3",
    "links": [ {
      "href": "https://cce.region.***.com/v3",
      "rel": "self"
    } ],
    "min_version": "",
    "status": "CURRENT",
    "updated": "2018-09-15 00:00:00Z",
    "version": ""
  } ]
}
```

Status Codes

Status Code	Description
200	The API version list is queried successfully.

Error Codes

See [Error Codes](#).

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 API Gateway](#)
- [Calling Kubernetes APIs Through the Cluster API Server](#)

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 API for obtaining a cluster in a specified project to obtain the cluster ID.
Endpoint	Entry (URL) for a web service, which can be obtained from Endpoints.
uri	Path in which the resource requested by the API is located. Obtain the value from the URI of the API. For details, see Kubernetes API .

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: Use the API for obtaining cluster information to query the cluster UID.
- Method 2: Query the cluster on the **Cluster Information** 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 **UAE-Abu Dhabi** region is **cce.ae-ad-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.ae-ad-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 CPU and 200 MiB memory.

```
curl --location --request POST 'https://07da5****.cce.ae-ad-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.

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 on how to obtain a user token, see Authentication .

The content of the **nginx.json** file is as follows:

```
{
  "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

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 certificate and save it as client.crt.
cat ./kubeconfig.json |grep client-certificate-data | awk -F "" '{print $4}' | base64 -d > ./client.crt
# Obtain the 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: Query the API server address (private or public network address) on the **Cluster Information** page of the CCE console and download the certificate (`client.crt` and `client.key` files).

Step 2 Call native Kubernetes APIs using the cluster certificate.

For example, run the curl command to call an API to view pod information. In the following example, `192.168.0.198:5443` is the IP address of the cluster API server.

```
curl --cert ./client.crt --key ./client.key https://192.168.0.198:5443/api/v1/namespaces/default/pods/
```

For more cluster APIs, see [Kubernetes APIs](#).

----End

Related Documents

- [Kubernetes official SDK](#) (including Go, Python, and Java)
- [Accessing a Cluster Using Kubernetes APIs](#)

6 Permissions Policies and Supported Actions

This chapter describes fine-grained permissions management for your CCE resources. If your account does not need individual IAM users, you can skip the operations in this chapter.

By default, new IAM users do not have permissions assigned. You need to add a user to one or more groups, and attach permissions policies or roles to these groups. Users inherit permissions from the groups to which they are added and can perform specified operations on cloud services based on the permissions.

You can grant users permissions by using `iam:attach-policy-user`. Roles are a type of service-based, coarse-grained authorization mechanism that defines permissions related to user 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 the permissions required to call all APIs, but IAM users must be assigned the permissions to call the required APIs. The permissions required for calling an API are determined by the actions supported by the API. Only users who have been granted permissions 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 in IAM

IAM provides system-defined policies that can be directly used. 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:

- **Permission:** a statement in a policy that allows or denies certain operations.
- **APIs:** REST APIs that can be called in a custom policy.

- **Actions:** added to a custom policy to control permissions for specific operations.
- **Dependent actions:** actions on which a specific action depends to take effect. When assigning permissions for the action to a user, you also need to assign permissions for the dependent actions.
- **IAM project/Enterprise project:** An action in a custom policy can be applied to either or both of these projects. Policies that contain actions supporting both IAM and enterprise projects can be assigned to user groups and take effect in both IAM and Enterprise Management. Policies that only contain actions supporting IAM projects can be assigned to user groups and only take effect for IAM. Such policies will not take effect if they are assigned to user groups 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 that can be defined in custom policies:

Table 6-1 Cluster

Permission	APIs	Action	IAM Project	Enterprise Project
Listing clusters in a specified project	GET /api/v3/projects/{project_id}/clusters	cce:cluster:list	√	√
Obtaining information about a specified 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 information about a specified 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	√	√

Permission	APIs	Action	IAM Project	Enterprise Project
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	√	√
Changing the specifications of a cluster	POST /api/v2/projects/{project_id}/clusters/:clusterid/resize	cce:cluster:resize	√	√
Obtaining a cluster certificate	POST /api/v3/projects/{project_id}/clusters/{cluster_id}/clustercert	cce:cluster:get	√	√

Table 6-2 Node

Permission	API	Action	IAM Project	Enterprise Project
Obtaining information about all nodes in a cluster	GET /api/v3/projects/{project_id}/clusters/{cluster_id}/nodes	cce:node:list	√	√
Obtaining information about a specified 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	√	√
Updating information about a specified 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

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 Nodepool

Permission	API	Action	IAM Project	Enterprise Project
Obtaining information about all node pools in a specified cluster	GET /api/v3/projects/{project_id}/clusters/{cluster_id}/nodepools	cce:nodepool:list	√	√
Obtaining information about 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 information about 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

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 chart information	GET /v2/charts/{id}	cce:chart:get	√	x
Deleting a chart	DELETE /v2/charts/{id}	cce:chart:delete	√	x

Table 6-6 Release

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

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 Addon

Permission	API	Action	IAM Project	Enterprise Project
Creating an add-on instance	POST /api/v3/addons	cce:addonInstance:create	√	√
Obtaining information about 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

Permission	API	Action	IAM Project	Enterprise Project
Querying 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.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.
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.

Status Code	Error Code	Error Message	Description	Troubleshooting
403	CCE.01403001	Forbidden.	Access denied.	Rectify the fault based on the returned message or contact technical support.
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.
409	CCE.03409001	Addon instance has installed.	The add-on instance has been installed.	Delete the add-on instance and install it again.

Status Code	Error Code	Error Message	Description	Troubleshooting
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.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 an API is called. Therefore, you need to 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 the Account ID

An account 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 **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 storage-driver add-ons are installed by default for clusters of Kubernetes 1.13 and earlier versions, and the coredns, everest, and npd add-ons are installed by default for clusters of Kubernetes 1.15 and later versions. For details, see Table 7-3 .

Table 7-3 Data structure of the AddonTemplate field

Parameter	Mandatory	Type	Description
addonTemplateName	Yes	String	Add-on name. The value can be coredns , autoscaler , gpu-beta , or storage-driver .
version	No	String	Add-on version. To view the add-on version information, log in to the CCE console, click Add-ons in the navigation pane. Click the target add-on name. On the Version History tab page, you can view the versions of the add-on. If this parameter is left blank, the latest version is used by default.
values	No	Json Map	Parameters to be set for installing an add-on. <ul style="list-style-type: none"> • coredns: For details, see Table 7-4. • autoscaler: For details, see Table 7-8. • gpu-beta: For details, see Table 7-10. <p>NOTE You do not need to set this parameter when installing the storage-driver add-on.</p>

Table 7-4 Special key-value pairs of the coredns values parameters

Parameter	Mandatory	Type	Description
flavor	Yes	flavor object	Add-on specifications.
custom	Yes	custom object	Custom coredns parameters.

Table 7-5 Data structure of the flavor field

Parameter	Mandatory	Type	Description
description	Yes	String	Specification description.
name	Yes	String	Flavor name.
replicas	Yes	Integer	Number of pods.
resources	Yes	resource object	Add-on resources.

Table 7-6 Data structure of the resource field

Parameter	Mandatory	Type	Description
name	Yes	String	Resource name.
limitsCpu	Yes	String	CPU quota limit.
limitsMem	Yes	String	Memory quota limit.
requestsCpu	Yes	String	Requested CPU.
requestsMem	Yes	String	Requested memory.

Table 7-7 Data structure of the custom field

Parameter	Mandatory	Type	Description
stub_domains	No	Json Map	JSON map that uses the DNS suffix key (such as acme.local) and the value composed of the JSON array of the DNS IP address.

Parameter	Mandatory	Type	Description
upstream_nameservers	No	Json Array	The value of this parameter replaces the name server value obtained from <code>/etc/resolv.conf</code> of the node by default. Restriction: A maximum of three upstream name servers can be specified.
tenant_id	Yes	String	Project ID.

Table 7-8 Special key-value pairs of the autoscaler values parameters

Parameter	Mandatory	Type	Description
flavor	Yes	flavor Object	Add-on specifications.
custom	Yes	custom object	Custom autoscaler parameters.

Table 7-9 Data structure of the custom field

Parameter	Mandatory	Type	Description
coresTotal	Yes	Integer	Maximum and minimum number of cores of a cluster. The format is <code><min>:<max></code> . The autoscaler add-on performs cluster scaling within this range.
maxEmptyBulkDeleteFlag	Yes	Integer	Maximum number of empty nodes that can be deleted at a time.
maxNodesTotal	Yes	Integer	Maximum number of nodes in all node groups. The autoscaler add-on performs cluster scaling within this range.
memoryTotal	Yes	Integer	Maximum and minimum size of memory of a cluster. The format is <code><min>:<max></code> . The autoscaler add-on performs cluster scaling within this range.
scaleDownDelayAfterAdd	Yes	Integer	Interval for performing scale-down evaluation after a scale-up.

Parameter	Mandatory	Type	Description
scaleDownDelayAfterDelete	Yes	Integer	Interval for performing scale-down evaluation after nodes are deleted. The default value is scanInterval .
scaleDownDelayAfterFailure	Yes	Integer	Interval for performing scale-down evaluation after a scale-down failure.
scaleDownEnabled	Yes	Boolean	Whether to enable cluster scaling down.
scaleDownUnneededTime	Yes	Integer	Time waited to start a scale-down after a node has reached the scale-down threshold.
scaleDownUtilizationThreshold	Yes	Float	Threshold percentage of the total CPU or memory occupied by all pods running on the node to the allocatable resources on the node. When the actual percentage is lower than this threshold, node scale-down is triggered.
scaleUpCpuUtilizationThreshold	Yes	Float	CPU usage threshold for triggering scale-up.
scaleUpMemUtilizationThreshold	Yes	Float	Memory usage threshold for triggering scale-up.
scaleUpUnscheduledPodEnabled	Yes	Boolean	Scale-up is triggered when there are unscheduled pods.
scaleUpUtilizationEnabled	Yes	Boolean	Scale-up is triggered when the resource usage reaches the threshold.
tenant_id	Yes	String	Project ID.
unremovableNodeRecheckTimeout	Yes	Integer	Interval for checking whether a node can be removed.

Table 7-10 Special Key-Value in values of AddonTemplate (GPU-beta)

Parameter	Mandatory	Type	Description
custom	Yes	custom object	Customer parameters of gpu-beta.

Table 7-11 Data structure of the custom field

Parameter	Mandatory	Type	Description
is_driver_from_nvidia	Yes	Boolean	Whether to use the NVIDIA driver. Set this parameter to true .
nvidia_driver_download_url	Yes	String	Driver downloading URL.

Table 7-12 Data structure of the spec field

Parameter	Mandatory	Type	Description
type	Yes	String	<p>Cluster type. Possible values:</p> <ul style="list-style-type: none"> VirtualMachine: The cluster is a hybrid cluster. A hybrid cluster manages a group of node resources based on Kubernetes. It can manage VMs, bare-metal machines, or a combination of both. Kubernetes automatically schedules containers onto available nodes. Before creating a containerized workload, you must have an available cluster. ARM64: The cluster is a Kunpeng cluster. Containers in CCE's Kunpeng clusters can run on Kunpeng servers that use ARM architecture and Kunpeng processors. Kunpeng-accelerated cloud servers are easy to deploy and provide comparable scaling and scheduling performance as x86-based cloud servers.

Parameter	Mandatory	Type	Description
flavor	Yes	String	<p>Cluster flavor, which cannot be changed after the cluster is created.</p> <ul style="list-style-type: none"> • cce.s1.small: small-scale, single-master hybrid cluster (≤ 50 nodes). • cce.s1.medium: medium-scale, single-master hybrid cluster (≤ 200 nodes). • cce.s1.large: large-scale, single-master hybrid cluster ($\leq 1,000$ nodes). • cce.s2.small: small-scale, high-availability hybrid cluster (≤ 50 nodes). • cce.s2.medium: medium-scale, high-availability hybrid cluster (≤ 200 nodes). • cce.s2.large: large-scale, high-availability hybrid cluster ($\leq 1,000$ nodes). <p>NOTE</p> <ul style="list-style-type: none"> • s1: single-master hybrid cluster. • s2: high-availability hybrid cluster. • dec: dedicated hybrid cluster specifications. For example, cce.dec.s1.small is a small-scale, single-master, dedicated hybrid cluster (≤ 50 nodes). • For example, ≤ 50 nodes indicates that the maximum number of nodes that can be managed by the cluster is 50. • A single-master cluster has only one master node. If the master node is down, the cluster will become unavailable and stop serving new workloads. However, existing workloads in the cluster are not affected. • A high-availability cluster has multiple master nodes. Faults in a single master node will not take the cluster down.

Parameter	Mandatory	Type	Description
version	No	String	<p>Cluster's baseline Kubernetes version. The latest version is recommended.</p> <ul style="list-style-type: none"> You can create clusters of two latest versions on the CCE console. To learn which cluster versions are available, choose Dashboard > Buy Cluster on the CCE console and check the Version parameter. You can call APIs to create clusters of other versions. However, these clusters will be gradually brought offline. For details about the offline policy, see the official CCE announcement. <p>NOTE If this parameter is not set, the cluster of the latest version is created by default.</p>
description	No	String	<p>Cluster description, for example, which purpose the cluster is intended to serve. By default, this parameter is left unspecified. To modify cluster description after the cluster is created, call the API used to update information about a specified cluster or go to the cluster details page on the CCE console.</p>
ipv6enable	No	Boolean	<p>Whether the cluster supports IPv6. Clusters of version 1.15 and later support IPv6.</p>
hostNetwork	Yes	HostNetwork object	<p>Node networking parameters, including VPC and subnet ID. hostNetwork is mandatory because nodes in a cluster communicate with each other by using a VPC.</p>
containerNetwork	Yes	ContainerNetwork object	<p>Container networking parameters, including the container network model and container CIDR block.</p>
authentication	No	Authentication object	<p>Configurations of the cluster authentication mode.</p>
kubernetesServiceIpRange	No	String	<p>Service CIDR block or the IP address range which the kubernetes clusterIp must fall within. This parameter is available only for clusters of v1.11.7 and later.</p>

Parameter	Mandatory	Type	Description
billingMode	No	Integer	Billing mode of a cluster. Currently, only pay-per-use clusters can be created. Value 0 indicates pay-per-use. If this parameter is left unspecified, the default value 0 is used.
masters	No	MasterSpec object	Advanced configurations of the master node
extendParam	No	Map<String, String>	Extended fields in the format of key-value pairs. You can configure multi-AZ clusters and dedicated hybrid clusters, and create clusters in specific enterprise projects.

Request Example

```
{
  "kind": "Cluster",
  "apiVersion": "v3",
  "metadata": {
    "name": "mycluster",
    "labels": {
      "foo": "bar"
    },
    "annotations": {
      "cluster.install.addons/install": "[{\"addonTemplateName\":\"autoscaler\",\"version\":\"1.15.11\",
      \"values\":{\"flavor\":{\"description\":\"Has only one instance\",\"name\":\"Single\",\"replicas\":1,\"resources\":
      [{\"limitsCpu\":\"90m\",\"limitsMem\":\"200Mi\",\"name\":\"autoscaler\",\"requestsCpu\":\"50m\",
      \"requestsMem\":\"100Mi\"}],\"custom\":{\"coresTotal\":3200,\"maxEmptyBulkDeleteFlag\":
      10,\"maxNodesTotal\":1000,\"memoryTotal\":128000,\"scaleDownDelayAfterAdd\":
      10,\"scaleDownDelayAfterDelete\":10,\"scaleDownDelayAfterFailure\":3,\"scaleDownEnabled\":false,
      \"scaleDownUnneededTime\":10,\"scaleDownUtilizationThreshold\":0.5,\"scaleUpCpuUtilizationThreshold\":
      1,\"scaleUpMemUtilizationThreshold\":1,\"scaleUpUnscheduledPodEnabled\":true,
      \"scaleUpUtilizationEnabled\":true,\"tenant_id\":\"22a8a02394794b908d256103a5b63c65\",
      \"unremovableNodeRecheckTimeout\":5}}},{\"addonTemplateName\":\"coredns\",\"version\":\"1.15.3\",
      \"values\":{\"flavor\":{\"description\":\"Has two instances\",\"name\":\"2500\",\"replicas\":2,\"resources\":
      [{\"limitsCpu\":\"500m\",\"limitsMem\":\"512Mi\",\"name\":\"coredns\",\"requestsCpu\":\"500m\",
      \"requestsMem\":\"512Mi\"}],\"custom\":{\"stub_domains\":{\"},\"upstream_nameservers\":[],\"tenant_id\":
      \"22a8a02394794b908d256103a5b63c65\"}},{\"addonTemplateName\":\"gpu-beta\",\"version\":\"1.1.7\",
      \"values\":{\"custom\":{\"is_driver_from_nvidia\":true,\"nvidia_driver_download_url\":\"https://
      us.download.nvidia.com/tesla/396.37/NVIDIA-Linux-x86_64-396.37.run\"}},{\"addonTemplateName\":
      \"storage-driver\",\"version\":\"\"}]]"
    }
  },
  "spec": {
    "type": "VirtualMachine",
    "flavor": "cce.s1.small",
    "version": "v1.15.11",
    "description": "this is a demo cluster",
    "hostNetwork": {
      "vpc": "1cb74d47-8e09-4d14-a065-75f4fc03e5eb",
      "subnet": "d44c089c-fbdc-49b3-ae8f-7b42c64219cf"
    },
    "containerNetwork": {
      "mode": "overlay_l2",

```

```
"cidr" : "172.16.0.0/16"  
},  
"extendParam" : {  
  "foo" : "bar"  
}  
}  
}
```

7.6 How to Obtain Parameters in the API URI

project_id

project_id indicates the project ID, which can be obtained from the console or APIs. For details, see [Obtaining a Project 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

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, select **Nodes** on the left, and move the cursor to the node name to view the node ID.

----End

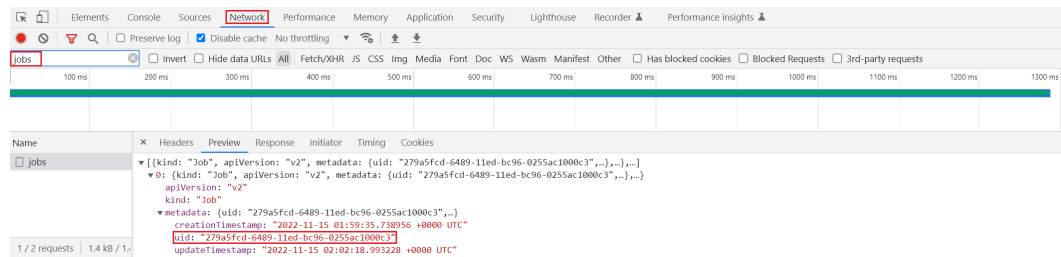
job_id

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 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 ECSs).

Python

To salt a password in the Python 3.7.7 environment, perform the following steps:

NOTE

The **python crypt** package has compatibility issues in macOS. If the package cannot be executed, run it in Linux.

1. Add \ before \$ in the salt. Generate a ciphertext password based on the updated salt.

```
python3 -c "import crypt;print(crypt.crypt('*****', crypt.mksalt()))"
```

2. Encode the value of the **password** field using Base64.

```
echo -n '*****' | base64 | tr "\n" " " | sed s/[[:space:]]//g
```

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 (StringUtils.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:

- For a cluster using the container tunnel network model, the value depends only on **the maximum number of pods on a node**.
- For clusters using the VPC network model, the value depends on **the maximum number of pods on a node** and **the minimum number of container IP addresses that can be allocated to a node**. It is recommended that the maximum number of pods on a node be less than or equal to the number of container IP addresses that can be allocated to the node. Otherwise, pods may fail to be scheduled.
- For CCE Turbo clusters using the Cloud Native Network 2.0 model, the value depends on **the maximum number of pods on a node** and **the minimum number of ENIs on a CCE Turbo cluster node**. It is recommended that the maximum number of pods on a node be less than or equal to the number of ENIs on the node. Otherwise, pods may fail to be scheduled.

Number of Container IP Addresses That Can Be Allocated on a Node

If you select **VPC network** for **Network Model** when creating a CCE cluster, you also need to set the number of container IP addresses that can be allocated to each node (`alpha.cce/fixPoolMask`). If the pod uses the host network (**hostNetwork: true**), the pod does not occupy the IP address of the allocatable container network. For details, see [Container Network vs. Host Network](#).

This parameter affects the maximum number of pods that can be created on a node. Each pod occupies an IP address (when the **container network** is used). If the number of available IP addresses is insufficient, pods cannot be created. If the pod uses the host network (**hostNetwork: true**), the pod does not occupy the IP address of the allocatable container 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. For example, in the preceding figure, **the number of container IP addresses that can be allocated to a node is 125 (128 - 3)**.

Maximum Number of Pods on a Node

When creating a node, you can configure the maximum number of pods that can be created on the node (`maxPods`). 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 7-13 lists the default maximum number of pods on a node based on node specifications.

Table 7-13 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

Number of Node ENIs (CCE Turbo Clusters)

In a CCE Turbo cluster, ECS nodes use sub-ENIs. The maximum number of pods that can be created on a node depends on the number of ENIs that can be used by the node.

Container Network vs. 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: The pod uses the host network (**hostNetwork: true** needs to be configured for the pod) and occupies the host port. The pod IP address is the host IP address. The pod does not occupy the IP addresses of the container network. To use the host network, you must confirm whether the container ports conflict with the host ports. Do not use the host network unless you know exactly which host port is used by which container.

7.12 Node OS

Mappings Between Cluster Versions and OS Versions

The following table lists the mappings between released cluster versions and OS versions.

Table 7-14 Mappings between cluster versions and OSs

OS	Cluster Version
CentOS 7.6	v1.23
	v1.21

7.13 Data Disk Space Allocation

This section describes how to allocate data disk space to nodes so that you can configure the data disk space accordingly.

Allocating Data Disk Space

When creating a node, you need to configure a data disk for the node and ensure that the data disk capacity is greater than or equal to 100 GB. You can click **Expand** to customize the data disk space allocation of the node.

- **Allocate Disk Space:**

CCE divides the data disk space for container engines and pods. The container engine space stores the **Docker/containerd** working directories, container images, 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.
- **Allocate Pod Basesize:** indicates the base size of a container. 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 be smaller 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.

Allocating Disk Space

For a non-shared data disk, 100 GB for example, can be divided as follows (depending on the container storage Rootfs): For details about the container storage Rootfs corresponding to different operating systems, 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



Allocating Basesize 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 GB.
- In OverlayFS mode, the pod container space is not limited by default.

NOTE

In the case of using Docker on EulerOS 2.9 nodes, `basesize` will not take effect if `CAP_SYS_RESOURCE` or `privileged` is configured for a container.

When configuring `basesize`, consider the maximum number of pods on a node. The container engine space should be greater than the total disk space used by containers. Formula: **the 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 GB 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 GB, the thin pool space occupied by files keeps increasing until 10 GB 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 7-15 Node OSs and container engines in CCE clusters

OS	Container Storage Rootfs	Customized 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 container engine is Docker. The default value is 10G. Not supported when Rootfs is set to OverlayFS.
EulerOS 2.5	Device Mapper	Supported only when the container engine is Docker. The default value is 10G.
EulerOS 2.9	OverlayFS	Supported only by clusters of v1.19.16, v1.21.3, v1.23.3, and later. The container basesize is not limited by default. Not supported when the cluster versions are earlier than v1.19.16, v1.21.3, and v1.23.3.

Table 7-16 Node OSs and container engines in CCE Turbo clusters

OS	Container Storage Rootfs	Customized Basesize
CentOS 7.x	OverlayFS	Not supported.
EulerOS 2.9	ECS VMs use OverlayFS. ECS PMs use Device Mapper.	Supported only when Rootfs is set to OverlayFS and the container engine is Docker. The container basesize is not limited by default. Supported when Rootfs is set to Device Mapper and the container engine is Docker. The default value is 10G.

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 above the high threshold (default: 85%) will trigger garbage collection. The garbage collection will delete least recently used images until the low threshold (default: 80%) has been 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.

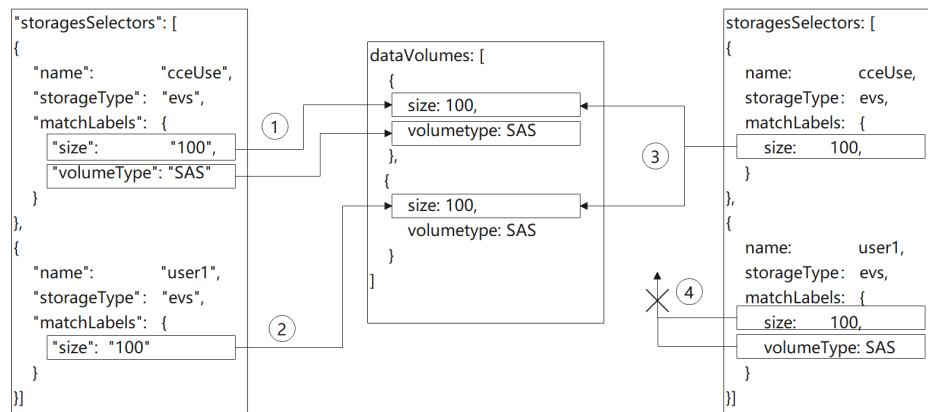
The **storage** field consists of **storageSelectors** and **storageGroups**. The **storageSelectors** field is responsible for disk selection, and the **storageGroups** field is responsible for disk processing.

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:



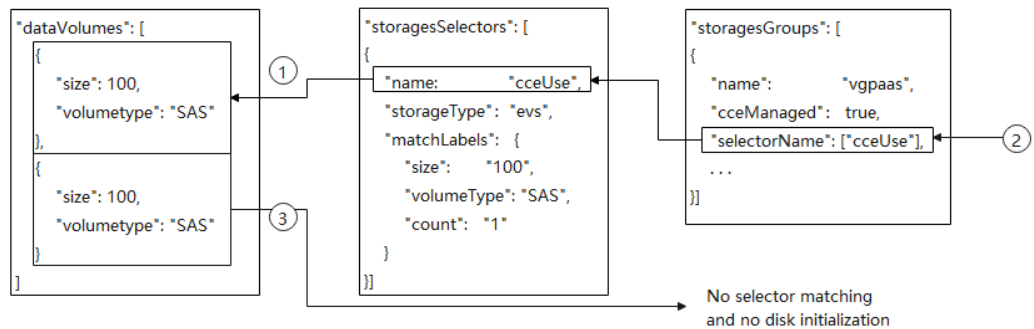
- a. In step 1, the first disk in **dataVolumes** is matched by the EVS disk whose size is 100 GB and storage class is SAS. In step 2, the second disk in **dataVolumes** is matched by the EVS disk whose size is 100 GB because the first disk has been selected.

- 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 GB 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 GB data disk.
2. The selected disk is managed by CCE and used as a data disk.
3. The other 100 GB 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 GB 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-253:0-786433-7cb37dc21202bfe2fc78dd1d33b70571e7e1982e56a4118f6acdc630cbc8b38 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-253:0-786433-5ecc4420da9a58fb66108db599a8267af3e8856da86b9c3d7fb82099a8781ae8 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-253:0-786433-7cb37dc21202bfe2fc78dd1d33b70571e7e1982e56a4118f6acdc630cbc8b38 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-253:0-786433-5ecc4420da9a58fb66108db599a8267af3e8856da86b9c3d7fb82099a8781ae8 253:10 0 10G 0 dm  /var/lib/docker/devicemapper/
└─sdc                                253:4  0 10G 0 lvm  /mnt/paas/kubernetes/kubelet
sde                                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": "ae-ad-1a",
  "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%"
        }
      ]
    }
  ]
}

```

```

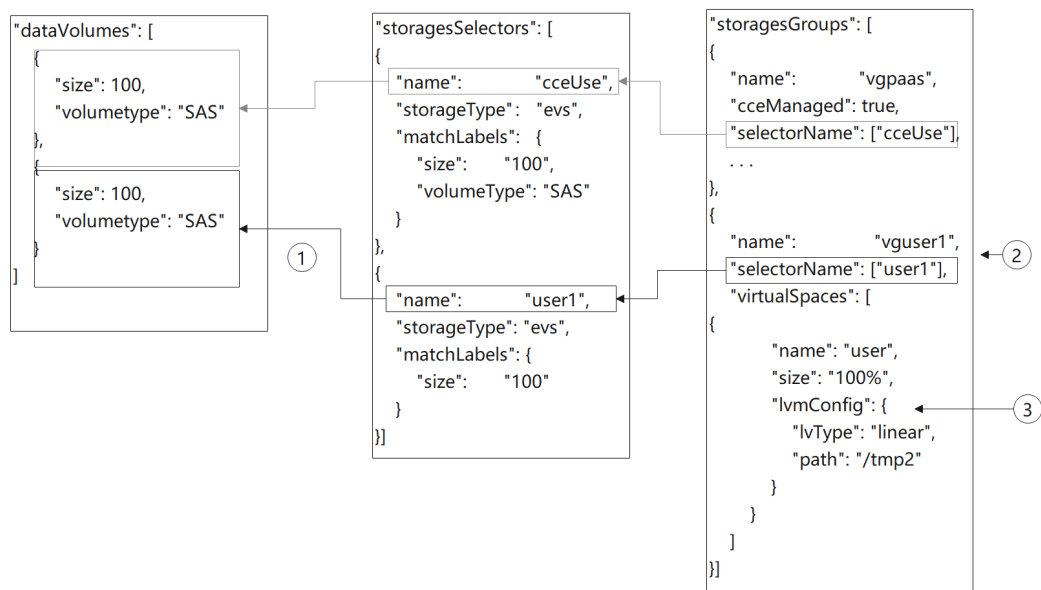
    }
  ],
  "count": 1
}
}

```

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 GB 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 GB 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-a6ac0d3ae8bffd57a92e6812079e503db49942619d5bbc69bb9516b31e15e67 253:10 0 10G 0 dm  /var/lib/docker/devicemapper/
            └─docker-253:0-917505-f9dfa31cdc3eb514a797c98311372ac8497d9a99581acdfefff0114bdf80e525 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-a6ac0d3ae8bffd57a92e6812079e503db49942619d5bbc69bb9516b31e15e67 253:10 0 10G 0 dm  /var/lib/docker/devicemapper/
         └─docker-253:0-917505-f9dfa31cdc3eb514a797c98311372ac8497d9a99581acdfefff0114bdf80e525 253:11 0 10G 0 dm  /var/lib/docker/devicemapper/
   └─vgpaas-kubernetes                    253:4  0 20G 0 lvm  /mnt/paas/kubernetes/kubelet
      └─vguser1-user                       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": "ae-ad-1a",
    "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 striping fails.

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": "ae-ad-1a",
    "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
}

```

8 Change History

Released On	Description
2023-03-31	CCE Turbo clusters can be created.
2023-01-31	Added Kubernetes APIs .
2022-08-24	Added the X-Cluster-ID parameter to the API for creating or deleting a PersistentVolumeClaim (PVC).
2021-08-26	Updated Permissions Policies and Supported Actions .
2020-08-31	This issue is the first official release.