

Alibaba Cloud Alibaba Cloud Container Service for Kubernetes

API Reference (Old version)

Issue: 20200701

Legal disclaimer

Alibaba Cloud reminds you to carefully read and fully understand the terms and conditions of this legal disclaimer before you read or use this document. If you have read or used this document, it shall be deemed as your total acceptance of this legal disclaimer.

1. You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloud-authorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.
2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company, or individual in any form or by any means without the prior written consent of Alibaba Cloud.
3. The content of this document may be changed due to product version upgrades, adjustments, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and the updated versions of this document will be occasionally released through Alibaba Cloud-authorized channels. You shall pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.
4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides the document in the context that Alibaba Cloud products and services are provided on an "as is", "with all faults" and "as available" basis. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not bear any liability for any errors or financial losses incurred by any organizations, companies, or individuals arising from their download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, bear responsibility for any indirect, consequential, exemplary, incidental, special, or punitive damages, including lost profits arising from the use or trust in this document, even if Alibaba Cloud has been notified of the possibility of such a loss.

5. By law, all the contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of this document shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates.
6. Please contact Alibaba Cloud directly if you discover any errors in this document.

Document conventions

Style	Description	Example
	A danger notice indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	 Danger: Resetting will result in the loss of user configuration data.
	A warning notice indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	 Warning: Restarting will cause business interruption. About 10 minutes are required to restart an instance.
	A caution notice indicates warning information, supplementary instructions, and other content that the user must understand.	 Notice: If the weight is set to 0, the server no longer receives new requests.
	A note indicates supplemental instructions, best practices, tips, and other content.	 Note: You can use Ctrl + A to select all files.
>	Closing angle brackets are used to indicate a multi-level menu cascade.	Click Settings > Network > Set network type .
Bold	Bold formatting is used for buttons, menus, page names, and other UI elements.	Click OK .
Courier font	Courier font is used for commands.	Run the <code>cd /d C:/window</code> command to enter the Windows system folder.
Italic	Italic formatting is used for parameters and variables.	<code>bae log list --instanceid Instance_ID</code>
[] or [a b]	This format is used for an optional value, where only one item can be selected.	<code>ipconfig [-all -t]</code>

Style	Description	Example
{ } or {a b}	This format is used for a required value, where only one item can be selected.	switch {active stand}

Contents

Legal disclaimer.....	I
Document conventions.....	I
1 Use the Kubernetes API.....	1
2 Make API requests.....	5
2.1 Overview.....	5
2.2 Common parameters.....	5
2.3 Request structure.....	7
2.4 Return results.....	8
2.5 Signature.....	8
3 Cluster API list.....	13
3.1 Create a Kubernetes cluster.....	13
3.2 Create a managed Kubernetes cluster.....	23
3.3 Create a serverless Kubernetes cluster.....	36
3.4 Create a managed edge Kubernetes cluster.....	39
3.5 Create a Kubernetes cluster that supports sandboxed-containers.....	46
3.6 Scale out a multi-zone cluster.....	54
3.7 Scale out a managed edge cluster.....	57
3.8 Expand a Kubernetes cluster that supports sandboxed-containers.....	61
3.9 Obtain the cluster kubeconfig file.....	64
3.10 Change the cluster tag.....	66
3.11 Query resource tags.....	67
3.12 Add an ECS instance to a cluster.....	69
3.13 Add a node to a managed edge cluster.....	73
3.14 Add existing ENS instances to a managed edge cluster.....	76
3.15 View a cluster.....	78
3.16 View all clusters.....	81
3.17 Remove a node.....	83
3.18 Delete a cluster.....	85

1 Use the Kubernetes API

This topic describes how to access the Kubernetes API by using REST calls based on curl commands. You can easily perform operations on Kubernetes clusters through the HTTPS protocol. The following text provides examples on how to create and delete pods, and how to create and modify deployments.

Obtain cluster credentials

1. Log on to the . On the Clusters page, find the target cluster and click **Manage** in the Actions column.

Cluster Name/ID	Labels	Type	Region (All)	Network Type	Cluster Status	Nodes	Created At	Version	Actions
haij... c1b...	...	Managed Kubernetes	China (Hangzhou)	VPC Network vpc- bp1ha32nxbx...	Running	2	May 29, 2020, 08:51:07 UTC	1.14.8- alyun.1	Manage View Logs Node Pools Upgrade Cluster More

2. On the Basic Information page, you can find cluster credentials in the **kubeconfig** file. Copy the content of the **kubeconfig** file and save it to your local computer.

Connect to a Kubernetes cluster using kubectl

[Open Cloud Shell](#)

1. Download the latest kubectl client. For more information, see [Kubernetes version changelog](#).
2. Install and set up the kubectl client. For more information, see [Install and set up kubectl](#).
3. Configure the cluster credentials.

Public Access

Internal Access

Revoke KubeConfig

Copy the following content to \$HOME/.kube/config on your local computer.

```
apiVersion: v1
clusters:
- cluster:
  server: https://...
  certificate-authority-data: LS0tLS1CRUdJTiBDRV...
  E1HTnNiM1ZrTUE4ROExVUVDaE1JVUdGdVozcG9iM1V4RXpBUkFnTlZCQU1UQ210MVltVnlibVYwW1hNdwpIaGNOTWpBd05USTVNRGcxTVRJeFdoY05NekF3T1RJM01EZz
```

Copy

3. Run the following commands to retrieve information about the certificate authority (CA), the private key, and API server.

```
# cat ./kubeconfig |grep client-certificate-data | awk -F ' ' '{print $2}' |base64 -d > client-cert.pem
# cat ./kubeconfig |grep client-key-data | awk -F ' ' '{print $2}' |base64 -d > client-key.pem
```

```
# APISERVER=`cat ./kubeconfig |grep server | awk -F ' ' '{print $2}'`
```

Access the Kubernetes API through curl commands

List all namespaces in the current cluster:

```
# curl --cert client-cert.pem --key client-key.pem -k $APISERVER/api/v1/namespaces
```

- Commonly used commands for pod management

List all pods in the default namespace:

```
# curl --cert client-cert.pem --key client-key.pem -k $APISERVER/api/v1/namespaces/default/pods
```

Create a pod (in JSON format):

```
# cat nginx-pod.json
{
  "apiVersion": "v1",
  "kind": "Pod",
  "metadata": {
    "name": "nginx",
    "namespace": "default"
  },
  "spec": {
    "containers": [
      {
        "name": "nginx",
        "image": "nginx:alpine",
        "ports": [
          {
            "containerPort": 80
          }
        ]
      }
    ]
  }
}
```

```
# curl --cert client-cert.pem --key client-key.pem -k $APISERVER/api/v1/namespaces/default/pods -X POST --header 'content-type: application/json' -d@nginx-pod.json
```

Create a pod (in YAML format):

```
# cat nginx-pod.yaml
apiVersion: v1
kind: Pod
metadata:
  name: nginx
  namespace: default
spec:
  containers:
  - name: nginx
    image: nginx:alpine
    ports:
    - containerPort: 80
```

```
# curl --cert client-cert.pem --key client-key.pem -k $APISERVER/api/v1/namespaces/default/pods -X POST --header 'content-type: application/yaml' --data-binary @nginx-pod.yaml
```

Query pod status:

```
# curl --cert client-cert.pem --key client-key.pem -k $APISERVER/api/v1/namespaces/default/pods/nginx
```

Query pod logs:

```
# curl --cert client-cert.pem --key client-key.pem -k $APISERVER/api/v1/namespaces/default/pods/nginx/log
```

Query metrics about a pod (through the metric-server API):

```
# curl --cert client-cert.pem --key client-key.pem -k $APISERVER/apis/metrics.k8s.io/v1beta1/namespaces/default/pods/nginx
```

Delete a pod:

```
# curl --cert client-cert.pem --key client-key.pem -k $APISERVER/api/v1/namespaces/default/pods/nginx -X DELETE
```

- Commonly used commands for deployment management

Create a deployment:

```
# cat nginx-deploy.yaml
apiVersion: extensions/v1beta1
kind: Deployment
metadata:
  name: nginx-deploy
  labels:
    app: nginx
spec:
  replicas: 2
  selector:
    matchLabels:
      app: nginx
  template:
    metadata:
      labels:
        app: nginx
    spec:
      containers:
      - name: nginx
        image: nginx:alpine
        ports:
        - containerPort: 80
      resources:
        requests:
          cpu: "2"
          memory: "4Gi"
```

```
# curl --cert client-cert.pem --key client-key.pem -k $APISERVER/apis/extensions/v1beta1/namespaces/default/deployments -X POST --header 'content-type: application/yaml' --data-binary @nginx-deploy.yaml
```

Query deployments:

```
# curl --cert client-cert.pem --key client-key.pem -k $APISERVER/apis/extensions/v1beta1/namespaces/default/deployments
```

Modify a deployment (change the number of replicas):

```
# curl --cert client-cert.pem --key client-key.pem -k $APISERVER/apis/extensions/v1beta1/namespaces/default/deployments/nginx-deploy -X PATCH -H 'Content-Type: application/strategic-merge-patch+json' -d '{"spec": {"replicas": 4}}'
```

Modify a deployment (change the container image):

```
# curl --cert client-cert.pem --key client-key.pem -k $APISERVER/apis/extensions/v1beta1/namespaces/default/deployments/nginx-deploy -X PATCH -H 'Content-Type: application/strategic-merge-patch+json' -d '{"spec": {"template": {"spec": {"containers": [{"name": "nginx", "image": "nginx:1.7.9"}]}}}}'
```

Related topics

[Officially supported Kubernetes client libraries](#)

[Kubernetes API specifications](#)

[Other methods to access the Kubernetes clusters](#)

We recommend that you use officially maintained client libraries to use the Kubernetes API.

2 Make API requests

2.1 Overview

The call to Container Service API interfaces is performed by sending HTTP requests to the server address of the Container Service APIs and adding corresponding request parameters to the requests according to the interface instructions. The system returns the results according to the processing results of the requests.

1. [Common parameters](#)
2. [Request structure](#)
3. [Return results](#)
4. [Signature](#)

2.2 Common parameters

Public request headers

Public request parameters are the request parameters that each interface must use.

Parameter	Description	Options
Authorization	The authentication information used to verify the validity of a request. Format: AccessKeyId: Signature.	Required
Content-Length	The content length of an HTTP request, which is defined in RFC 2616.	Required
Content-Type	The content type of an HTTP request, which is defined in RFC 2616.	Required

Parameter	Description	Options
Content-MD5	The Base64-encoded results converted from 128-bit MD5 hash value of the HTTP message body. We recommend that you add this message to all requests to prevent requests from being tampered.	Required
Date	The construction time of a request. Currently, only the GMT format is supported. If the difference between the construction time and the MNS server time exceeds 15 minutes, invalid request is returned.	Required
Host	The host access value, for example, diku.aliyuncs.com.	Required
Accept	The return type required by the client. application/json and application/xml are supported.	Required
x-acis-version	The API version. The current version is 2015-12-15.	Required
x-acis-region-id	A region indicates the physical location of an Elastic Compute Service (ECS) instance.	Required
x-acis-signature-nonce	The unique random number used to prevent network replay attacks. Different random numbers must be used for different requests.	Required
x-acis-signature-method	The method of user signature. Currently only HMAC-SHA1 is supported.	Required

Examples

```
GET /clusters HTTP/1.1
Host: cs.aliyuncs.com
```

```
Accept: application/json
User-Agent: cs-sdk-python/0.0.1 (Darwin/15.2.0/x86_64;2.7.10)
x-acs-signature-nonce: f63659d4-10ac-483b-99da-ea8fde61eae3
Authorization: acs <yourAccessKeyId>:<yourSignature>
x-acs-signature-version: 1.0
Date: Wed, 16 Dec 2015 11:18:47 GMT
x-acs-signature-method: HMAC-SHA1
Content-Type: application/json; charset=utf-8
X-Acs-Region-Id: cn-beijing
Content-Length: 0
```

Public response headers

Each time you send a request to call an interface, the system returns a unique identifier (RequestId), no matter the request is successful or not.

Examples

XML example:

```
<? xml version="1.0" encoding="UTF-8"? >
<!--Result Root Node-->
<Interface name+Response>
| <!--Return request tag-->
| <RequestId>4C467B38-3910-447D-87BC-AC049166F216</RequestId>
| <!--Return result data-->
</Interface name+Response>
```

JSON example:

```
{
  "RequestId": "4C467B38-3910-447D-87BC-AC049166F216"
  /* Response data */
}
```

2.3 Request structure

Service address

The access address of Alibaba Cloud Container Service API is cs.aliyuncs.com.

Communication protocol

The system supports request communication by using the HTTP or HTTPS channel. We recommend that you use the HTTPS channel to send requests for more security.

Request methods

Use HTTP methods such as PUT, POST, GET, and DELETE to send different requests.

Request parameters

Each request must contain the public request parameters and the request parameters unique to specified operations.

Request encoding

Both requests and returned results are encoded by using the UTF-8 character set.

2.4 Return results

After the API service is called, data is returned in a unified format. The returned HTTP status code `2xx` indicates that the call is successful. The returned HTTP status code `4xx` or `5xx` indicates that the call fails. When the call is successful, data can be returned mainly in two formats: XML and JSON. When a request is sent, an external system can pass in a parameter to define the format of the returned data. The default format is XML.

Examples of returned results in this document are formatted for ease of viewing. The actual results returned are not formatted with line breaks or indentation.

2.5 Signature

Introduction

The Access Key ID and Access Key Secret are officially issued to you by Alibaba Cloud (you can apply for and manage them on the Alibaba Cloud official website). The Access Key ID is used to identify your identity. The Access Key Secret is the key used to encrypt the signature string and verify the signature string on the server side. You must keep the Access Key Secret confidential. Only you and Alibaba Cloud can know it.

Container Service verifies each access request it receives. Therefore, all requests sent to Container Service must contain signature information. Container Service performs symmetric encryption by using the Access Key ID and Access Key Secret to verify the identity of request senders. If the calculated verification code is the same as the one provided, the request is considered as valid. Otherwise, Container Service rejects the request and returns the HTTP 403 error.

You can add the authorization header in the HTTP request to contain the signature information, indicating that the message has been authorized.

Container Service requires to contain the signature in the HTTP header in the format of `Authorization: acs [Access Key ID]:[Signature]`.

The **Signature** calculation method is as follows:

```
Signature = base64(hmac-sha1(VERB + "\n"  
    + ACCEPT + "\n"  
    + Content-MD5 + "\n"  
    + Content-Type + "\n"  
    + Date + "\n"  
    + CanonicalizedHeaders + "\n"  
    + CanonicalizedResource))
```

- **VERB** indicates the HTTP method, For example, **PUT**.
- **Accept** indicates the return type required by the client, which can be application/json or application/xml.
- **Content-MD5** indicates the MD5 value of the requested content.
- **Content-Type** indicates the type of the requested content.
- **Date** indicates the operation time, which cannot be null. Currently, only the GMT format is supported. If the difference between the request time and the CAS server time exceeds 15 minutes, CAS considers the request as invalid and returns error 400. For more information, see the 5th section. For example, Thu, 17 Mar 2012 18:49:58 GMT.
- **CanonicalizedHeaders** indicates a combination of fields started with **x-acsc-** in the HTTP request.
- **CanonicalizedResource** indicates the uniform resource identifier (URI) of the resource in the HTTP request. For example, /clusters? name=my-clusters&resource=new.

**Note:**

Conform to the following specifications for **CanonicalizedHeaders** (headers started with **x-acsc-**) before signature verification:

1. Convert the names of all HTTP request headers started with **x-acsc-** to lowercase letters. For example, convert **X-ACS-Meta-Name: TaoBao** to **x-acsc-meta-name: TaoBao**. The names of request headers are case-insensitive according to Alibaba Cloud specifications. However, we recommend that you use the lowercase letters.
2. If the value part of a public request header is too long, replace the **\t**, **\n**, **\r**, and **\f** separators with spaces.
3. Sort all HTTP request headers that are obtained from the preceding step and compliant with Alibaba Cloud specifications in the lexicographic ascending order.
4. Delete any space at either side of a separator between request header and content. For example, convert **x-acsc-meta-name: TaoBao, Alipay** to **x-acsc-meta-name:TaoBao,Alipay**.

5. Separate all headers and contents with the `\n` separator to form the final CanonicalizedHeaders.

**Note:**

The format specification for CanonicalizedResource: CanonicalizedResource indicates the standard description of the resource you want to access. Sort sub-resources and query in the lexicographically ascending order and separate them by using the `&` separator to generate a sub-resource string (all parameters after `?`).

```
http://cs.aliyuncs.com/clusters?name=my-clusters&resource=new
```

The CanonicalizedResource format is:

```
/clusters? name=my-clusters&resource=new
```

Signature example**Overview**

The following example shows the signature process.

In the example, the Access Key ID and Access Key Secret are `access_key_id` and `access_key_secret` respectively. We recommend that you use your own API call program to calculate the signature string in the following example. Then, compare your signature string with the example result.

The request example is as follows:

```
POST http://cs.aliyuncs.com/clusters?param1=value1&param2=value2 HTTP/1.1
Accept-Encoding: identity
Content-Length: 210
Content-MD5: 6U4ALMkKSj0PYbeQSHqgmA==
x-acs-version: 2015-12-15
Accept: application/json
User-Agent: cs-sdk-python/0.0.1 (Darwin/15.2.0/x86_64;2.7.10)
x-acs-signature-nonce: fbf6909a-93a5-45d3-8b1c-3e03a7916799
x-acs-signature-version: 1.0
Date: Wed, 16 Dec 2015 12:20:18 GMT
x-acs-signature-method: HMAC-SHA1
Content-Type: application/json;charset=utf-8
X-Acs-Region-Id: cn-beijing
Authorization: acs <yourAccessKeyId>:<yourSignature>
{"password": "Just$****", "instance_type": "ecs.m2.medium", "name": "my-test-cluster-9708****", "size": 1, "network_mode": "vpc", "data_disk_category": "cloud", "data_disk_size": 10, "ecs_image_id": "m-253l****l"}
```

Request construction process

Calculate Content-Length and Content-MD5

Assemble `SignatureString`. In the example, the length of the signature string is 307. An `\n` line break is at the end of all lines except the last line.

```
POST
application/json
6U4ALMkKSj0PYbeQSHqgmA==
application/json;charset=utf-8
Wed, 16 Dec 2015 12:20:18 GMT
x-accs-region-id:cn-beijing
x-accs-signature-method:HMAC-SHA1
x-accs-signature-nonce:fbf6909a-93a5-45d3-8b1c-3e03a7916799
x-accs-signature-version:1.0
x-accs-version:2015-12-15
/clusters? param1=value1&param2=value2
```

Calculate Signature

```
# Use Access Key Secret to encrypt the signature string. In the example, the accessKeySecret is access_key_secret.
hmac-sha1(SignatureString): fee03d405e421ebaf514adec881038c4b313584d
# Convert the encrypted string to a byte array, similar to the Content-MD5 calculation method.
# Convert the byte array into a Base64 string to get the final signature string.
base64(bytes(hmac-sha1(SignatureString))): ZmVlMDNkNDA1ZTQyMWVlYWY1MTRhZGVjODgxMDM4YzRlMzEzNTg0ZA==
Signature: ZmVlMDNkNDA1ZTQyMWVlYWY1MTRhZGVjODgxMDM4YzRlMzEzNTg0ZA==
```

Finish

After the preceding processing, add some other header information to construct the final HTTP request as follows:

```
POST http://cs.aliyuncs.com/clusters?param1=value1&param2=value2 HTTP/1.1
Accept-Encoding: identity
Content-Length: 210
Content-MD5: 6U4ALMkKSj0PYbeQSHqgmA==
x-accs-version: 2015-12-15
Accept: application/json
User-Agent: cs-sdk-python/0.0.1 (Darwin/15.2.0/x86_64;2.7.10)
x-accs-signature-nonce: fbf6909a-93a5-45d3-8b1c-3e03a7916799
x-accs-signature-version: 1.0
Date: Wed, 16 Dec 2015 12:20:18 GMT
x-accs-signature-method: HMAC-SHA1
Content-Type: application/json;charset=utf-8
X-Acs-Region-Id: cn-beijing
Authorization: acs <yourAccessKeyId>:<yourSignature>
{"password": "Just$****", "instance_type": "ecs.m2.medium", "name": "my-test-cluster-9708****", "size": 1, "network_mode": "vpc", "data_disk_category": "cloud", "data_disk_size": 10, "ecs_image_id": "m-253ll****"}
```

3 Cluster API list

3.1 Create a Kubernetes cluster

You can call the CreateCluster operation to create a Kubernetes cluster and add nodes to the cluster.

Request information

Request line

```
POST /clusters HTTP/1.1
```

Custom request headers

None. For more information, see [Public request headers](#).

Request body



Note:

You must set parameters correctly. Otherwise, errors may occur in the cluster creation process. When you create a cluster in the console, you can click **Generate API Request Parameters** on the Create Kubernetes Cluster page to view all parameters. For more information, see [#unique_12](#).

```
{ "disable_rollback": "whether to roll back when the operation fails", "name": "cluster name", "timeout_mins": "cluster creation timeout", "cluster_type": "cluster type. Kubernetes", "region_id": "region", "vpcid": "VPC ID", "master_vswitch_ids": "VSwitch IDs of master nodes. Specify three VSwitches in different zones to ensure high availability.", "master_instance_type": "instance types of master nodes. Specify three instance types.", "master_count": "number of master nodes. Valid values: 3 and 5", "container_cidr": "pod CIDR block", "service_cidr": "service CIDR block", "ssh_flags": "whether to allow Internet access through SSH", "cloud_monitor_flags": "whether to install the CloudMonitor agent", "login_password": "SSH logon password. Select between SSH authentication and key pair authentication", "key_pair": "key pair name. Select between SSH authentication and key pair authentication", "master_instance_charge_type": "billing method of master nodes. PostPaid or PrePaid", "master_period_unit": "unit of the subscription duration of master nodes. Month or Year. This parameter takes effect only when parameter master_instance_charge_type is set to PrePaid", "master_period": "subscription duration of master nodes. This parameter takes effect only when parameter master_instance_charge_type is set to PrePaid", "master_auto_renew": "whether to automatically renew the subscription of master nodes", "master_auto_renew_period": "subscription duration when master nodes are renewed", "master_system_disk_category": "system disk type of master nodes", "master_system_disk_size": "system disk size of master nodes", "master_data_disk": "whether to mount data disks to master nodes.", "master_data_disks": "data disk configurations of master nodes", "master_system_disk_snapshot_policy_id": "snapshot policy ID on the system disk of master nodes", "worker_instance_charge_type": "billing method of worker nodes. PrePaid or PostPaid", "worker_period_unit": "unit of the subscription duration of worker nodes. Month or Year. This parameter takes effect only when parameter worker_instance_charge_type is set to PrePaid", "worker_period": "
```

subscription duration of worker nodes. This parameter takes effect only when parameter `worker_instance_charge_type` is set to `PrePaid`, `"worker_auto_renew"`: "whether to automatically renew the subscription of worker nodes. Valid values: true and false", `"worker_auto_renew_period"`: "subscription duration when worker nodes are renewed", `"worker_instance_types"`: "instance types of worker nodes", `"worker_vswitch_ids"`: "VSwitch IDs of worker nodes. Specify one to five VSwitches", `"worker_system_disk_category"`: "system disk type of worker nodes", `"worker_system_disk_size"`: "system disk size of worker nodes", `"worker_data_disk"`: "whether to mount data disks to worker nodes.", `"worker_data_disks"`: "data disk configurations of worker nodes", `"worker_system_disk_snapshot_policy_id"`: "snapshot policy ID on the system disk of worker nodes", `"num_of_nodes"`: "number of worker nodes", `"snat_entry"`: "whether to configure SNAT entries", `"endpoint_public_access"`: "whether to enable Internet access to the cluster endpoint", `"cpu_policy"`: "CPU policy. Valid values: static and none", `"node_port_range"`: "port range of nodes. Valid values: 30000 to 65535", `"proxy_mode"`: "network mode. Valid values: iptables and ipvs", `"addons"`: "optional add-ons in an array", `"tags"`: "tags in an array", `"security_group_id"`: "security group ID", `"auto_snapshot_policy_id"`: "automatic snapshot policy ID", `"taint"`: "taints in an array" `"deletion_protection"`: "whether to enable cluster deletion protection. If this option is enabled, the cluster cannot be deleted through the console or API operations." `"os_type"`: "operating system of the nodes that run pods. For example, Linux and Windows", `"platform"`: "architecture of the nodes that run pods", `"runtime"`: "container runtime. Default is Docker", }

Table 3-1: Request body parameters

Parameter	Type	Required	Description
cluster_type	string	Yes	The type of the cluster.
key_pair	string	Yes	The name of the key pair. Select between <code>login_password</code> and <code>key_pair</code> .
login_password	string	Yes	The SSH logon password. The password must be 8 to 30 characters in length and contain three of the following four types of characters: uppercase letters, lowercase letters, digits, and special characters. Select between <code>login_password</code> and <code>key_pair</code> .
master_instance_charge_type	string	No	The billing method of master nodes. Valid values: - PrePaid: subscription - PostPaid: pay-as-you-go Default value: PostPaid.
master_instance_types	list	Yes	The instance types of master nodes. For more information, see #unique_13 .
master_system_disk_category	string	Yes	The system disk type of master nodes. Valid values: - cloud_efficiency: ultra disk - cloud_ssd: standard SSD

Parameter	Type	Required	Description
master_system_disk_size	int	Yes	The system disk size of master nodes. Unit: GiB.
master_vswitch_ids	list	Yes	The VSwitch IDs of master nodes. You must specify 1 to 3 VSwitch IDs. We recommend that you specify three VSwitches in different zones to ensure high availability.
num_of_nodes	int	Yes	The number of worker nodes. Valid values: [0, 300]
name	string	Yes	The name of the cluster. The name can contain uppercase letters, lowercase letters, Chinese characters, digits, and hyphens (-).
region_id	string	Yes	The ID of the region where the cluster is deployed.
snat_entry	bool	Yes	Whether to configure SNAT. - If the VPC network has Internet access, set this parameter to false. - If the VPC network has no Internet access: - Set this parameter to true to configure SNAT. This enables the cluster to access the Internet. - Set this parameter to false to not configure SNAT. The cluster will have no Internet access.
worker_instance_types	list	Yes	The instance types of worker nodes. For more information, see #unique_13 .
worker_system_disk_category	string	Yes	The system disk type of worker nodes.
worker_system_disk_size	int	Yes	The system disk size of worker nodes. Unit: GiB.
worker_vswitch_ids	list	Yes	The VSwitch IDs of worker nodes.

Parameter	Type	Required	Description
addons	list	No	The add-ons to the Kubernetes cluster. - Add-ons parameters: - name: Required. The name of the add-on. - version: Optional. Default is the latest version. - config: Optional. If this parameter is not set, it indicates that this parameter is not required. - Network plug-in: Select between Flannel and Terway. - Log Service: Optional. If Log Service is not enabled, the cluster audit function is not available. - Ingress: The Ingress component is installed by default.
container_cidr	string	No	The CIDR block of containers. This CIDR block must not overlap with that of the VPC network. If the VPC network is automatically created by the system, the container CIDR block is set to 172.16.0.0/16 by default.
cloud_monitor_flags	bool	No	Whether to install the CloudMonitor agent. - A value of true indicates yes. - A value of false indicates no.
cpu_policy	string	No	The CPU policy. For Kubernetes 1.12.6 and later versions, static and none are supported. Default value: none.
disable_rollback	bool	No	Whether to roll back when the operation fails. - A value of true indicates no. - A value of false indicates yes. Default value: true. If you choose to roll back when the operation fails, resources created during the operation are released. We recommend that you set this parameter to true.

Parameter	Type	Required	Description
master_auto_renew	bool	No	Whether to automatically renew the subscription of master nodes. This parameter takes effect only when parameter master_instance_charge_type is set to PrePaid . Valid values: - A value of true indicates that auto renewal is enabled. - A value of false indicates that auto renewal is disabled.
master_auto_renew_period	int	No	The duration of the subscription. This parameter is required when parameter master_instance_charge_type is set to PrePaid and parameter master_auto_renew is set to true. When PeriodUnit is set to Month, valid values are as follows: 1, 2, 3, 6, and 12.
master_count	int	No	The number of master nodes. Valid values: 3 and 5. Default value: 3.
master_data_disk	bool	No	Whether to mount data disks to master nodes. - A value of true indicates yes. - A value of false indicates no. Default is false.
master_data_disks	list	No	The data disk configurations of master nodes, such as the disk type and disk size. This parameter takes effect only when parameter master_data_disk is set to true. - category: The type of the data disk. Valid values: - cloud: basic disk - cloud_efficiency: ultra disk - cloud_ssd: standard SSD - size: The size of the data disk. Unit: GiB. - encrypted:

Parameter	Type	Required	Description
master_period	int	No	The duration of the subscription. This parameter is required when parameter master_instance_charge_type is set to PrePaid . When PeriodUnit is set to Month, valid values are as follows: 1, 2, 3, 6, and 12.
master_period_unit	string	No	The unit of the subscription duration. This parameter is required when parameter master_instance_charge_type is set to PrePaid . Valid values: A value of Month indicates that the subscription duration is measured in months.
node_port_range	string	No	The port range of nodes. Valid values: 30000 to 65535.
proxy_mode	string	No	The kube-proxy mode. Valid values: iptables and IPVS. Default value: iptables.
endpoint_public_access	bool	No	Whether to allow Internet access to the API Server. - A value of true indicates that Internet access to the API Server is allowed. This is the default setting. - A value of false indicates that Internet access to the API Server is denied.
security_group_id	string	No	The ID of the security group to which the ECS instances in the cluster belong.
service_cidr	string	No	The service CIDR block. This CIDR block must not overlap with that of the VPC network or container. If the VPC network is automatically created by the system, the service CIDR block is set to 172.19.0.0/20 by default.

Parameter	Type	Required	Description
ssh_flags	bool	No	Whether to enable SSH logon. - A value of true indicates that SSH logon is enabled. - A value of false indicates that SSH logon is disabled.
tags	list	No	The tags of the cluster. - key: The name of the tag. - value: The value of the tag.
taints	list	No	Taints are applied to nodes to make sure that pods are not scheduled to unsuitable nodes. Tolerations are applied to pods. If a pod has a toleration that matches the taint on a node, this pod can be scheduled to the node.
kubernetes_version	string	No	The Kubernetes version. Default is the latest version.
timeout_mins	int	No	The maximum time in minutes during which a resource creation operation must complete. Default: 60.
vpcid	string	No	Optional. The ID of the VPC network where the cluster is deployed. If this parameter is not set, the system automatically creates a VPC network with CIDR block 192.168.0.0/16.  Note: The vpcid and vswitchid parameters must be used together.
worker_auto_renew	bool	No	Whether to automatically renew the subscription of worker nodes. Valid values: - A value of true indicates that auto renewal is enabled. - A value of false indicates that auto renewal is disabled.

Parameter	Type	Required	Description
worker_auto_renew_period	int	No	The duration of the subscription. This parameter is required when parameter worker_instance_charge_type is set to PrePaid and parameter worker_auto_renew is set to true. When PeriodUnit is set to Month, valid values are as follows: 1, 2, 3, 6, and 12.
worker_data_disk	string	No	Whether to mount data disks to worker nodes. Valid values: - A value of true indicates yes. - A value of false indicates no.
worker_data_disks	list	No	The data disk configurations of worker nodes, such as the disk type and disk size. This parameter takes effect only when parameter worker_data_disk is set to true. - category: The type of the data disk. Valid values: - cloud: basic disk - cloud_efficiency: ultra disk - cloud_ssd: standard SSD - size: The size of the data disk. Unit: GiB. - encrypted: whether to encrypt the data disk. Valid values: true and false.
worker_instance_charge_type	string	No	The billing method of worker nodes. Valid values: - PrePaid: subscription - PostPaid: pay-as-you-go Default value: PostPaid.
worker_period	int	No	The duration of the subscription. This parameter is required when parameter worker_instance_charge_type is set to PrePaid. When PeriodUnit is set to Month, valid values are as follows: 1, 2, 3, 6, and 12.

Parameter	Type	Required	Description
worker_per iod_unit	string	No	The unit of the subscription duration. This parameter is required when parameter <code>worker_instance_charge_type</code> is set to <code>PrePaid</code> . Valid values: • A value of <code>Month</code> indicates that the subscription duration is measured in months.
deletion_p rotection	bool	No	Whether to enable cluster deletion protection. If this option is enabled, the cluster cannot be deleted through the console or API operations.
os_type	string	No	The operating system of the nodes that run pods. For example, Linux and Windows.
platform	string	No	The architecture of the nodes that run pods.
runtime	json	No	The container runtime. Default is Docker. You need to specify the name and version.
node_cidr_ mask	string	No	The prefix length of the node CIDR block.

Response information

Response line

```
HTTP/1.1 202 Accepted
```

Custom response headers

None. For more information, see [Public response headers](#).

Table 3-2: Response body

Parameter	Type	Description
cluster_id	string	The ID of the cluster.
request_id	string	The ID of the request.
task_id	string	The ID of the task. It is system-assigned and can be used to query the task status.

Examples

Sample requests

```
POST /clusters HTTP/1.1
<Common request headers> { "cluster_type": "Kubernetes",
  "name": "my-test-Kubernetes-cluster",
  "region_id": "cn-beijing",
  "disable_rollback": true,
  "timeout_mins": 60, "kubernetes_version": "1.12.6-aliyun.1",
  "snat_entry": true,
  "endpoint_public_access": false,
  "cloud_monitor_flags": false,
  "node_cidr_mask": "25",
  "proxy_mode": "iptables",
  "tags": [],
  "addons": [{"name": "flannel"}, {"name": "aliyun-log-controller", "config": "{\"sls_project_name\": \"k8s-log-c64f6eab6a1764d3dbec3dc2b9e41****\"}"}, {"name": "nginx-ingress-controller", "config": "{\"IngressSlbNetworkType\": \"internet\"}"}, {"name": "node_port_range": "30000-32767",
  "login_password": "test****",
  "cpu_policy": "none",
  "master_count": 3,
  "master_vswitch_ids": ["vsw-2ze48rkq464rsdts****", "vsw-2ze48rkq464rsdts1****", "vsw-2ze48rkq464rsdts1****"],
  "master_instance_types": ["ecs.sn1.medium", "ecs.sn1.medium", "ecs.sn1.medium"],
  "master_system_disk_category": "cloud_efficiency",
  "master_system_disk_size": 40,
  "worker_instance_types": ["ecs.sn2.3xlarge"],
  "num_of_nodes": 3,
  "worker_system_disk_category": "cloud_efficiency",
  "worker_system_disk_size": 120,
  "vpcid": "vpc-2zegvl5etah5requ0****",
  "worker_vswitch_ids": ["vsw-2ze48rkq464rsdts****"],
  "container_cidr": "172.20.XX.XX/16",
  "service_cidr": "172.21.XX.XX/20",
  "worker_data_disks": [{"category": "cloud_ssd", "size": 500}],
  "master_data_disks": [{"category": "cloud_ssd", "size": 500}],
  "taints": [{"key": "special", "value": "true", "effect": "NoSchedule"}]
}
```

Sample responses

```
HTTP/1.1 202 Accepted
<Common response headers> { "cluster_id": "cb95aa626a47740afbf6aa099b650****",
  "request_id": "687C5BAA-D103-4993-884B-C35E4314A1E1",
  "task_id": "T-5a54309c80282e39ea00002f"
```

```
}
```

3.2 Create a managed Kubernetes cluster

You can call **CreateCluster** to describe how to create a managed Kubernetes cluster that consists of multiple nodes.

Request information

Request line

```
POST /clusters HTTP/1.1
```

Custom request headers

None. For more information, see [Public request headers](#).

Request body



Note:

You must set parameters correctly. Otherwise, errors may occur in the cluster creation process. When you create a cluster in the console, you can click **Generate API Request Parameters** on the Create Kubernetes Cluster page to view all parameters. For more information, see [#unique_12](#).

```
{ "disable_rollback": "whether to roll back when the operation fails", "name": "cluster name", "timeout_mins": "cluster creation timeout", "cluster_type": "cluster type, ManagedKubernetes", "region_id": "region", "vpcid": "VPC ID", "worker_vswitch_ids": "VSwitch IDs. Specify one to three VSwitches", "container_cidr": "pod CIDR block", "service_cidr": "service CIDR block", "cloud_monitor_flags": "whether to install the CloudMonitor agent", "login_password": "SSH logon password. Select between SSH authentication and key pair authentication", "key_pair": "key pair name. Select between SSH authentication and key pair authentication", "worker_instance_charge_type": "billing method of worker nodes. PrePaid or PostPaid", "worker_period_unit": "unit of the subscription duration of worker nodes. Month or Year. This parameter only takes effect when parameter worker_instance_charge_type is set to PrePaid", "worker_period": "subscription duration of worker nodes. This parameter only takes effect when parameter worker_instance_charge_type is set to PrePaid", "worker_auto_renew": "whether to automatically renew the subscription of worker nodes. Valid values: true and false", "worker_auto_renew_period": "subscription duration when worker nodes are renewed", "worker_instance_types": "instance types of worker nodes", "worker_system_disk_category": "system disk type of worker nodes", "worker_system_disk_size": "system disk size of worker nodes", "worker_data_disk": "whether to mount data disks. Valid values: true and false", "worker_data_disk_size": "data disk configuration of worker nodes", "num_of_nodes": "number of worker nodes", "snat_entry": "whether to configure SNAT entries", "endpoint_public_access": "whether to enable Internet access to the cluster endpoint", "proxy_mode": "network mode. Valid values: iptables and ipvs", "addons": "an array of add-ons", "tags": "an array of tags", "security_group_id": "security group ID", "taint": "an array of taints", "cpu_policy": "CPU policy. Valid values: static and none", "runtime": "container runtime. Valid value: docker", "platform": "platform architecture of the nodes that run pods", "os_type": "operating system of the nodes that run pods. For example, Linux and Windows",
```

"node_cidr_mask": "prefix length of the node IP address", "kubernetes_version": "Kubernetes version. Default is the latest version.", "deletion_protection": "whether to enable cluster deletion protection. After this option is enabled, the cluster cannot be deleted through the console or API operations." }

Table 3-3: Request body parameters

Parameter	Type	Required	Description
cluster_type	string	Yes	The type of the cluster.
key_pair	string	Yes	The name of the key pair. Select between login_password and key_pair.
login_password	string	Yes	The SSH logon password. The password must be 8 to 30 characters in length and contain three of the following four types of characters : uppercase letters , lowercase letters, digits, and special characters. Select between login_password and key_pair.
name	string	Yes	The name of the cluster. The name can contain uppercase letters, lowercase letters, Chinese characters, digits, and hyphens (-).
num_of_nodes	int	Yes	The number of worker nodes. Valid values: [0, 100]
region_id	string	Yes	The ID of the region where the cluster is deployed.

Parameter	Type	Required	Description
snat_entry	bool	Yes	Whether to configure SNAT. If the VPC network is automatically created, the parameter must be set to true. If an existing VPC network is specified and it has Internet access, set the parameter to false. Otherwise, set the parameter to true .
vswitch_ids	list	Yes	The VSwitch IDs. The list must contain one to three elements.
worker_system_disk_category	string	Yes	The system disk type of worker nodes.
worker_system_disk_size	int	Yes	The system disk size of worker nodes.

Parameter	Type	Required	Description
addons	list	No	The add-ons of the Kubernetes cluster. • Add-ons parameters: - name: Required. The name of the add-on. - version: Optional. Default is the latest version. - config: Optional. If this parameter is not set, it indicates that this parameter is not required. • Network plug-in : Select between Flannel and Terway. • Log Service: Optional. If Log Service is not enabled, the cluster audit function is unavailable. • Ingress: The ingress component is installed by default.

Parameter	Type	Required	Description
container_cidr	string	No	The CIDR block of pods. This CIDR block must not overlap with that of the VPC network. If the VPC network is automatically created by the system, the pod CIDR block is set to 172.16.0.0/16 by default.
cloud_monitor_flags	bool	No	Whether to install the CloudMonitor agent.
disable_rollback	bool	No	Whether to roll back when the operation fails. • A value of true indicates no. • A value of false indicates yes. If you choose to roll back when the operation fails, resources created during the operation are released. We recommend that you set this parameter to true.

Parameter	Type	Required	Description
public_slb	bool	No	Whether to allow Internet access to the API Server. - A value of true indicates that Internet access to the API Server is allowed. This is the default setting. - A value of false indicates that Internet access to the API Server is denied.  Note: The parameter is obsolete and replaced by parameter endpoint_public_access.
proxy_mode	string	No	The kube-proxy mode. Valid values : iptables and IPVS. Default is iptables.

Parameter	Type	Required	Description
endpoint_public_access	bool	No	Whether to allow Internet access to the API Server. - A value of true indicates that Internet access to the API Server is allowed. This is the default setting. - A value of false indicates that Internet access to the API Server is denied.
security_group_id	string	No	The ID of the security group to which the ECS instances in the cluster belong.
service_cidr	string	No	The service CIDR block. This CIDR block must not overlap with that of the VPC network or containers. If the VPC network is automatically created by the system, the service CIDR block is set to 172.19.0.0/20 by default.
tags	list	No	The tags of the cluster. - key: The name of the tag. - value: The value of the tag.

Parameter	Type	Required	Description
taints	list	No	Taints are applied to nodes to make sure that pods are not scheduled to unsuitable nodes . Tolerations are applied to pods. If a pod has a toleration that matches the taint on a node, this pod can be scheduled to the node.
timeout_mins	int	No	The maximum time in minutes during which a resource creation operation must be complete. Default: 60.
vpcid	string	Yes	Required. The ID of the VPC network where the cluster is deployed. If this parameter is not set, the system automatically creates a VPC network with CIDR block 192.168.0.0/16. The vpcid and vswitchid parameters must be used together.

Parameter	Type	Required	Description
worker_auto_renew	bool	No	Whether to automatically renew the subscription of worker nodes. • A value of true indicates that auto renewal is enabled. • A value of false indicates that auto renewal is disabled.
worker_auto_renew_period	int	No	The duration of the subscription. This parameter is required when parameter worker_instance_charge_type is set to PrePaid . When PeriodUnit is set to Month, valid values are as follows: 1, 2, 3, 6, and 12.
worker_data_disk	string	No	Whether to mount data disks to worker nodes. • A value of true indicates yes. • A value of false indicates no.

Parameter	Type	Required	Description
worker_data_disks	list	No	The data disk configurations of worker nodes, such as the disk type and disk size. This parameter only takes effect when parameter <code>worker_data_disk</code> is set to true. - category: The type of the data disk. Valid values: - cloud: basic disk - cloud_efficiency: ultra disk - cloud_ssd: standard SSD - size: The size of the data disk. Unit : GiB.
worker_data_disk_category	int	No	The type of the data disk.  Note: This parameter is obsolete and replaced by category of parameter worker_data_disks.

Parameter	Type	Required	Description
worker_data_disk_size	string	No	The size of the data disk.  Note: This parameter is obsolete and replaced by size of parameter worker_data_disks.
worker_instance_charge_type	string	No	The billing method of worker nodes. Valid values: - PrePaid: subscription - PostPaid: pay-as-you-go
worker_period	int	No	The duration of the subscription. This parameter is required when parameter worker_instance_charge_type is set to PrePaid. When PeriodUnit is set to Month, valid values are as follows: 1, 2, 3, 6, and 12.

Parameter	Type	Required	Description
worker_period_unit	string	No	The unit of the subscription duration. This parameter is required when parameter <code>worker_instance_charge_type</code> is set to PrePaid.  Note: A value of Month indicates that the subscription duration is measured in months.
worker_instance_types	list	Yes	The instance types of worker nodes. For more information, see #unique_13 .
cpu_policy	string	No	The CPU policy. Kubernetes 1.12.6 and later support static and none. Default is none.
runtime	json	No	The container runtime. Currently, only Docker is supported. - name: The name of the container runtime. - version: The version of the container runtime.
platform	string	No	The platform architecture of the nodes that run pods.

Parameter	Type	Required	Description
os_type	string	No	The operating system of the nodes that run pods. For example, Linux and Windows.
node_cidr_mask	int	No	The prefix length of the node IP address.
kubernetes_version	string	No	The Kubernetes version. Default is the latest version.
deletion_protection	bool	No	Whether to enable cluster deletion protection. After this option is enabled, the cluster cannot be deleted through the console or API operations.

Response information

Response line

```
HTTP/1.1 202 Accepted
```

Custom response headers

None. For more information, see [Public response headers](#).

Response body

```
{ "cluster_id": "string",  
  "request_id": "string",  
  "task_id": "string"  
}
```

Examples

Sample requests

```
POST /clusters HTTP/1.1  
<Common request headers> { "name": "test",  
  "cluster_type": "my-test-Kubernetes-cluster",  
  "disable_rollback": true,  
  "timeout_mins": 60,  
  "kubernetes_version": "1.12.6-aliyun.1",  
  "region_id": "cn-beijing",
```

```

"snat_entry":true,
"cloud_monitor_flags":false,
"endpoint_public_access":false,
"node_cidr_mask":"25",
"proxy_mode":"ipvs",
"tags":[],
"addons":[{"name": "flannel"}, {"name": "aliyun-log-controller", "config": "{ \"sls_project_name\": \"k8s-log-c64f6eab6a1764d3db3dc2b9e41****\" }"}, {"name": "nginx-ingress-controller", "config": "{ \"IngressSlbNetworkType\": \"internet\" }"}],
"worker_instance_types":["ecs.hfc5.xlarge"],
"num_of_nodes":3,
"worker_system_disk_category":"cloud_efficiency",
"worker_system_disk_size":120,
"worker_instance_charge_type":"PostPaid",
"vpcid":"vpc-2zegvl5etah5requ09nec",
"container_cidr":"172.20.0.0/16",
"service_cidr":"172.21.0.0/20",
"vswitch_ids":["vsw-2ze48rkq464rsdts1****"],
"worker_data_disks":[{"category": "cloud_ssd", "size": 500}],
"login_password":"test@19****",
"taint":[{"key": "special", "value": "true", "effect": "NoSchedule"}]
}

```

Sample responses

```

HTTP/1.1 202 Accepted
<Common response headers> {
  "cluster_id": "cb95aa626a47740afbf6aa099b65****",
  "request_id": "687C5BAA-D103-4993-884B-C35E4314A1E1",
  "task_id": "T-5a54309c80282e39ea00002f"
}

```

3.3 Create a serverless Kubernetes cluster

You can call **CreateCluster** to describe how to create a serverless Kubernetes cluster.

Request information

Request line

```
POST /clusters HTTP/1.1
```

Custom request headers

None. For more information, see [Public request headers](#).

Request body

```

{
  "cluster_type": "Ask",
  "name": "cluster name",
  "region_id": "region",
  "zoneid": "zone",
  "nat_gateway": "true", //whether to create a NAT gateway
  "private_zone": "true", //whether to enable privateZone for service discovery
  "vpc_id": "VPC ID", //If this parameter is not specified, a VPC network is automatically
  created.
  "tags": "an array of tags",
}

```

```
"vswitchid": "VSwitch ID",
}
```

Table 3-4: Request body parameters

Name	Type	Required	Description
cluster_type	string	Yes	The type of the cluster.
name	string	Yes	The name of the cluster. The name can contain uppercase and lowercase letters, digits, and hyphens (-).
private_zone	bool	Yes	Whether to enable PrivateZone for service discovery. Valid values: true and false. For more information, see #unique_16 .
region_id	string	Yes	The ID of the region where the cluster is deployed.
zoneid	string	Yes	The zone of the region.
vswitch_id	string	No	The ID of the VSwitch. Optional. If this parameter is not specified, the system automatically creates a VSwitch with CIDR block 192.168.0.0/16.
tags	list	No	The tags of the cluster. - key: The name of the tag. - value: The value of the tag.
nat_gateway	bool	No	Whether to create a NAT Gateway. Valid values: true and false. Default is false.
vpc_id	string	No	The ID of the VPC network where the cluster is deployed. Optional. If this parameter is not specified, the system automatically creates a VPC network with CIDR block 192.168.0.0/16.

**Note:**

The vpc_id and vswitch_id parameters must be used together.

Response information

Response line

```
HTTP/1.1 202 Accepted
```

Custom response headers

None. For more information, see [Public response headers](#).

Table 3-5: Response body

Name	Type	Description
cluster_id	string	The ID of the cluster.
request_id	string	The ID of the request.
task_id	string	The ID of the task. The task ID is automatically assigned by the system and used to query task status.

Examples

Sample request

```
POST /clusters HTTP/1.1
<Common request headers>
{
  "cluster_type": "ASK",
  "name": "my-test-Kubernetes-cluster",
  "region_id": "cn-beijing",
  "tags": [],
  "vpcid": "vpc-2zegvl5etah5requ0****",
  "vswitch_ids": ["vsw-2ze48rkq464rsdts1****"]
}
```

Sample response

```
HTTP/1.1 202 Accepted
<Common response headers>
{
  "cluster_id": "cb95aa626a47740afbf6aa099b650****",
  "request_id": "687C5BAA-D103-4993-884B-C35E4314A1E1",
  "task_id": "T-5a54309c80282e39ea00002f"
```

```
}
```

3.4 Create a managed edge Kubernetes cluster

You can call **CreateCluster** to create a managed edge Kubernetes cluster instance.

Request

Request format

```
POST /clusters HTTP/1.1
```

Custom request headers

None. See [Public request headers](#).

Request body

```
{
  "disable_rollback": "whether to roll back when the operation fails",
  "name": "cluster name",
  "timeout_mins": "cluster creation timeout",
  "cluster_type": "cluster type. ManagedKubernetes",
  "profile": "profile of the edge cluster. Edge",
  "region_id": "region",
  "vpcid": "VPC ID",
  "worker_vswitch_ids": "one or multiple VSwitch IDs. The number of VSwitch IDs can be 1, 2, 3, 4, or 5",
  "container_cidr": "POD CIDR Block",
  "service_cidr": "Service CIDR Block",
  "cloud_monitor_flags": "whether to install the CloudMonitor agent",
  "login_password": "SSH logon password. Select one between SSH authentication and key pair authentication",
  "key_pair": "key pair name. Select one between SSH authentication and key pair authentication",
  "worker_instance_types": "multiple instance types from which Worker instances are to be created",
  "worker_system_disk_category": "system disk type of Worker nodes",
  "worker_system_disk_size": "system disk size of Worker nodes",
  "worker_data_disk": "whether data disks are attached to Worker nodes. Valid values: true or false",
  "worker_data_disk_category": "type of the data disks",
  "worker_data_disk_size": "size of the data disks",
  "num_of_nodes": "number of Worker nodes",
  "snat_entry": "whether to configure SNAT entries",
  "endpoint_public_access": "whether the public network exposes the cluster endpoint",
  "tags": "cluster tags in an array object"
}
```

Table 3-6: Request body parameters

Parameter	Type	Required	Description
cluster_type	string	Yse	The type of the cluster.

Parameter	Type	Required	Description
key_pair	string	Yse	The keypair name. Select one between login_password and key_pair.
login_password	string	Yse	The SSH logon password. The password must be 8 to 30 characters in length and contain letters, numbers, and special characters. Select one between login_password and key_pair.
name	string	Yse	The cluster name . The name can contain uppercase and lowercase letters , Chinese characters , digits, and hyphens (-).
num_of_nodes	Integer	Yes	The number of Worker nodes. Valid values: [0, 300]
profile	string	Yse	The profile of the edge cluster. Default value: Edge.
region_id	string	Yse	The ID of the region where the cluster is located.

Parameter	Type	Required	Description
snat_entry	bool	Yes	Whether to configure SNAT. If the VPC network is automatically created, this parameter must be set to true. If an existing VPC network is used, configure this parameter as follows:
vswitch_ids	list	Yes	The IDs of the VSwitch. The list can contain 1 to 3 list elements.
worker_system_disk_category	string	Yes	The system disk type of Worker nodes.
worker_system_disk_size	Integer	Yes	The system disk size of Worker nodes.
container_cidr	string	No	The CIDR block associated with the container. This CIDR block must not overlap with that of the VPC network. If the VPC network is automatically created by the system, the container is associated with CIDR block 172.16.0.0 /16 by default.

Parameter	Type	Required	Description
cloud_monitor_flags	bool	No	Whether to install the CloudMonitor agent.  Note: If you choose to use the log service or monitoring service, a cloud ECS node is automatically purchased to deploy the corresponding services. In this case, you must configure the logon information and resource specifications of the Worker node.
disable_rollback	bool	No	Whether to roll back when the operation fails. • true: A value of true indicates no. • false: A value of true indicates yes. If you choose to roll back when the operation fails, resources created during the operation are released. We recommend that you choose true
proxy_mode	string	No	The kube-proxy mode. The iptables and IPVS modes are supported. The default mode is iptables.

Parameter	Type	Required	Description
endpoint_public_access	bool	No	Whether to allow public access to the API Server. - A value of true indicates that public access to the API Server is allowed. Default value: true. - A value of false indicates that public access to the API Server is denied.  Note: In the edge cluster scenario, the edge nodes interact with the cloud through the public network. Therefore, the edge cluster must enable the public network access.
service_cidr	string	No	The service CIDR block. This CIDR block must not overlap with that of the VPC network or container. If the VPC network is automatically created by the system, the service CIDR block is set to 172.19.0.0/20 by default.

Parameter	Type	Required	Description
tags	list	No	Tag the cluster: - key: The name of the tag. - value: The value of the tag.
timeout_mins	int	No	The maximum time in minutes during which the cluster creation operation must be completed. Default: 60.
vpcid	string	No	The ID of the VPC network where the cluster is deployed . Optional. If this parameter is not specified, the system automatically creates a VPC network with CIDR block 192.168.0.0/16. Parameter vpcid and vswitchid must be used together.
worker_data_disk	string	No	Whether data disks are attached to Worker nodes. Valid values: - A value of true indicates yes. - A value of false indicates no.
worker_data_disk_category	int	No	The type of the data disk.
worker_data_disk_size	string	No	The size of the data disk.

Response

Response format

```
HTTP/1.1 202 Accepted
```

Custom response headers

None. See [Public response headers](#).

Response body

```
{
  "cluster_id": "string",
  "request_id": "string",
  "task_id": "string"
}
```

Samples

Sample requests

```
POST /clusters HTTP/1.1
<Common request header>
{
  "name": "test",
  "cluster_type": "ManagedKubernetes",
  "profile": "Edge",
  "disable_rollback": true,
  "timeout_mins": 60,
  "region_id": "cn-beijing",
  "snat_entry": true,
  "cloud_monitor_flags": true,
  "endpoint_public_access": true,
  "tags": [],
  "worker_instance_types": ["ecs.hfc5.xlarge"],
  "num_of_nodes": 1,
  "worker_system_disk_category": "cloud_efficiency",
  "vpcid": "vpc-2zegvl5eta*****",
  "container_cidr": "172.20.0.0/16",
  "service_cidr": "172.21.0.0/20",
  "vswitch_ids": ["vsw-2ze48rkq464rsdts1****"],
  "login_password": "test@19****"
}
```

Sample responses

```
HTTP/1.1 202 Accepted
<Common response header>
{
  "cluster_id": "cb95aa626a47740afbf6aa099b65****",
  "request_id": "687C5BAA-D103-4993-884B-C35E4314****",
  "task_id": "T-5a54309c80282e39ea00****"
}
```

```
}
```

3.5 Create a Kubernetes cluster that supports sandboxed-containers

You can call **CreateCluster** to describe how to create a Kubernetes cluster that supports sandboxed-containers.

Request information

Request line

```
POST /clusters HTTP/1.1
```

Custom request headers

None. For more information, see [Public request headers](#).

Request body

```
{
  "name": "cluster name",
  "cluster_type": "cluster type: ManagedKubernetes",
  "disable_rollback": "true",
  "timeout_mins": "60",
  "region_id": "region",
  "snat_entry": "whether to configure SNAT entries",
  "cloud_monitor_flags": "whether to install the CloudMonitor agent",
  "public_slb": "whether to expose the cluster endpoint to the Internet",
  "worker_instance_type": "instance type of worker nodes",
  "num_of_nodes": "number of worker nodes",
  "vpcid": "vpc id",
  "vswitch_ids": "VSwitch IDs of worker nodes",
  "service_cidr": "Service CIDR Block",
  "login_password": "SSH logon password. Select one between SSH authentication and key pair authentication",
  "key_pair": "key pair name. Select one between SSH authentication and key pair authentication",
  "worker_system_disk_category": "system disk type of worker nodes",
  "worker_system_disk_size": "system disk size of worker nodes",
  "worker_data_disk_category": "data disk type of worker nodes",
  "worker_data_disk_size": "data disk size of worker nodes",
  "worker_data_disk": "whether to mount data disks. true|false",
  "worker_instance_charge_type": "billing method of worker nodes. PrePaid|PostPaid",
  "worker_period_unit": "subscription duration unit of worker nodes. Month|Year. This parameter only takes effect when parameter worker_instance_charge_type is set to PrePaid",
  "worker_period": "subscription duration of worker nodes. This parameter only takes effect when parameter worker_instance_charge_type is set to PrePaid",
  "worker_auto_renew": "whether to enable auto renewal for worker nodes. true|false",
  "worker_auto_renew_period": "renewal cycle of worker nodes",
  "kubernetes_version": "Kubernetes version",
  "addons": [{"name": "terway-eniip"}],
  "runtime": {"name": "Sandboxed-Container.runv", "version": "1.0.0"},
  "pod_vswitch_ids": "VSwitch IDs of Pods. Currently, only one VSwitch is supported."
}
```

}

Table 3-7: Request body parameters

Parameter	Type	Required	Description
addons	list	Yes	The add-ons to the Kubernetes cluster. - Add-on parameters: - name: Required. The name of the add-on. - version: Optional. If this parameter is not specified, the latest version is used by default. - config: Optional. If this parameter is not specified, it indicates that this parameter is not needed. - Network plug-in: Select one between Flannel and Terway. - Log Service: Optional. - If Log Service is not enabled , you cannot use cluster auditing features. - To enable Log Service, add <code>{"name":"aliyun-log-controller"}</code> to the add-on array.
cluster_type	string	Yes	The type of the cluster. Currently, only ManagedKubernetes clusters support sandboxed containers.
key_pair	string	Yes	The keypair name. Select one between login_password and key_pair.
kubernetes_version	string	Yes	Sandboxed-containers are only available in 1.14.6-aliyun.1 and later versions.

Parameter	Type	Required	Description
login_password	string	Yes	The SSH logon password. The password must be 8 to 30 characters in length and contain three of the following four types of characters: uppercase letters, lowercase letters, digits, and special characters. Select one between login_password and key_pair .
name	string	Yes	The cluster name. The name can contain uppercase letters , lowercase letters, Chinese characters, digits, and hyphens (-).
num_of_nodes	int	Yes	The number of worker nodes. Valid values: [0, 300]
pod_vswitch_ids	list	Yes	The ID of the VSwitch used by the Pods.
region_id	string	Yes	The ID of the region where the cluster is deployed.
runtime	json	Yes	- name: The container runtime. Currently, Sandboxed-Container .runv is supported. - version: The current version is 1 .0.0.
snat_entry	bool	Yes	Whether to enable SNAT. - If the existing VPC network has Internet access, set this parameter to false. - If the existing VPC network has no Internet access: - Set this parameter to true to enable SNAT. This enables the cluster to access the Internet. - Set this parameter to false to not enable SNAT. The cluster will have no Internet access.

Parameter	Type	Required	Description
vpcid	string	Yes	The ID of the VPC network. If this parameter is not specified, the system automatically creates a VPC network with CIDR block 192.168.0.0/16.  Note: The vpcid and vswitchid parameters must be used together.
vswitch_ids	list	Yes	The ID of the VSwitch. The list must contain 1 to 3 elements.
worker_data_disk	bool	Yes	Whether to mount data disks. This parameter must be set to true to enable sandboxed-containers.
worker_data_disk_size	string	Yes	The size of the data disk. Unit: GiB . This parameter must be set to a value no less than 200 to enable sandboxed-containers.
worker_instance_type	string	Yes	The type of instance that supports sandboxed-containers. Currently, the following instance types are supported: ecs.ebmg5s.24xlarge and ecs.ebmc5s.24xlarge .
worker_system_disk_category	string	Yes	The system disk type of worker nodes.
worker_system_disk_size	int	Yes	The system disk size of worker nodes. Unit: GiB.
cloud_monitor_flags	bool	No	Whether to install the CloudMonitor agent. - A value of true indicates yes. - A value of false indicates no.

Parameter	Type	Required	Description
disable_rollback	bool	No	Whether to roll back when the operation fails. - A value of true indicates no. - A value of false indicates yes. Default value: true. If you choose to roll back when the operation fails, resources created during the operation are released. We recommend that you set the value to true.
endpoint_public_access	bool	No	Whether to allow public access to the API Server. - A value of true indicates that public access to the API Server is allowed. Default value: true. - A value of false indicates that public access to the API Server is denied. The API Server only allows internal access.
service_cidr	string	No	The service CIDR block. This CIDR block must not overlap with that of the VPC network or container. If the VPC network is automatically created by the system, the service CIDR block is set to 172.19.0.0/20 by default.
tags	list	No	The labels of the cluster. - key: The name of the label. - value: The value of the label.
timeout_mins	int	No	The maximum time in minutes during which a resource creation operation must complete. Default: 60.
worker_auto_renew	bool	No	Whether to enable auto renewal for worker nodes. - A value of true indicate yes. - A value of false indicate no.

Parameter	Type	Required	Description
worker_auto_renew_period	int	No	The cycle of auto renewal. This parameter is required when parameter <code>worker_instance_charge_type</code> is set to <code>PrePaid</code> and parameter <code>worker_auto_renew</code> is set to <code>true</code>. - When <code>PeriodUnit=Week</code>, valid values are as follows: 1, 2, and 3. - When <code>PeriodUnit=Month</code>, valid values are as follows: 1, 2, 3, 6, and 12.
worker_data_disk_category	string	No	The data disk type of worker nodes. This parameter only takes effect when parameter <code>worker_data_disk</code> is set to <code>true</code>. Valid values: <code>cloud</code>: basic disk <code>cloud_efficiency</code>: ultra disk <code>cloud_ssd</code>: SSD disk
worker_instance_charge_type	string	No	The billing method of worker nodes. Valid values: - A value of <code>PrePaid</code> indicates that the subscription billing method is used. - A value of <code>PostPaid</code> indicates that the pay-as-you-go billing method is used. Default value: <code>PostPaid</code>.

Parameter	Type	Required	Description
worker_period	int	No	The duration of the subscription. This parameter is required when parameter worker_instance_charge_type is set to PrePaid . - When PeriodUnit=Week, valid values are as follows: 1, 2, 3, and 4. - When PeriodUnit=Month, valid values are as follows: 1, 2, 3, 4, 5, 6, 7, 8, 9, 12, 24, 36, 48, and 60.
worker_period_unit	string	No	The subscription duration unit. This parameter is required when parameter worker_instance_charge_type is set to PrePaid. Valid values: - A value of Week indicates that the subscription duration is measured in weeks. - A value of Month indicates that the subscription duration is measured in months.

Response information

Response line

HTTP/1.1 202 Accepted

Custom response headers

None. See [Public response headers](#).

Table 3-8: Response body

Parameter	Type	Description
cluster_id	string	The ID of the cluster.
request_id	string	The ID of the request.

Parameter	Type	Description
task_id	string	The ID of the task. The task ID is automatically assigned by the system and used to query task status.

Example

Sample request

```
POST /clusters HTTP/1.1
<Common request headers>
{
  "name": "test-sandbox",
  "cluster_type": "ManagedKubernetes",
  "disable_rollback": true,
  "timeout_mins": 60,
  "region_id": "cn-hangzhou",
  "snat_entry": true,
  "cloud_monitor_flags": false,
  "endpoint_public_access": true,
  "worker_instance_type": "ecs.ebmg5s.24xlarge",
  "num_of_nodes": 2,
  "vpcid": "vpc-bp1iybm49v9jgb50xxxxx",
  "vswitch_ids": ["vsw-bp1ue9z93i9zpcblxxxxx"],
  "service_cidr": "192.168.0.0/16",
  "login_password": "xxxxxxxx",
  "worker_system_disk_category": "cloud_efficiency",
  "worker_system_disk_size": 200,
  "worker_data_disk_category": "cloud_efficiency",
  "worker_data_disk_size": 200,
  "worker_data_disk": true,
  "worker_instance_charge_type": "PostPaid",
  "kubernetes_version": "1.14.6-aliyun.1",
  "addons": [{"name": "terway-eniip"}],
  "runtime": {"name": "Sandboxed-Container.runv", "version": "1.0.0"},
  "pod_vswitch_ids": ["vsw-bp18nj6xxfoc2ci2xxxxx"]
}
```

Sample response

```
HTTP/1.1 202 Accepted
<Common response headers>
{
  "cluster_id": "cb95aa626a47740afbf6aa099b65****",
  "request_id": "687C5BAA-D103-4993-884B-C35E4314A1E1",
  "task_id": "T-5a54309c80282e39ea00002f"
```

```
}
```

3.6 Scale out a multi-zone cluster

You can call the ScaleOutCluster operation to add worker nodes to a multi-zone cluster.

Request information

Request line

```
POST /api/v2/clusters/{cluster_id} HTTP/1.1
```

Table 3-9: Request parameters

Parameter	Type	Required	Description
cluster_id	string	Yes	The ID of the cluster that you want to add nodes to.

Custom request headers

None. For more information, see [Public request headers](#).

Request body

```
{ "key_pair": "key pair name. Select between SSH authentication and key pair authentication", "worker_vswitch_ids": "VSwitch IDs. Specify one to three VSwitches", "worker_instance_types": "instance type of worker nodes", "worker_system_disk_category": "system disk type of worker nodes", "worker_system_disk_size": "system disk size of worker nodes", "worker_data_disk": "whether to mount data disks to worker nodes. Valid values: true and false", "tags": "tags in an array", "count": "number of worker nodes to be added", "worker_data_disks": "data disk configurations of worker nodes", }
```

Table 3-10: Request body parameters

Parameter	Type	Required	Description
count	int	Yes	The number of nodes that you want to add.
vswitch_ids	list	No	The VSwitch IDs of worker nodes.

Parameter	Type	Required	Description
worker_data_disk	bool	Yes	Whether to mount data disk to worker nodes. • A value of true indicates yes. • A value of false indicates no.
worker_instance_types	list	Yes	The instance types of worker nodes. For more information, see #unique_13 .
key_pair	string	Yes	The name of the key pair. Select between login_password and key_pair.
login_password	string	Yes	The passwords of the worker nodes to be added. The password must be 8 to 30 characters in length and contain three of the following four character types: uppercase letters, lowercase letters, digits, and special characters. Select between login_password and key_pair.

Parameter	Type	Required	Description
worker_data_disks	list	No	The data disk configurations of worker nodes, such as the disk type and disk size. This parameter takes effect only when parameter <code>worker_data_disk</code> is set to true. - category: The type of the data disk. Valid values: - cloud: basic disk - cloud_efficiency: ultra disk - cloud_ssd: standard SSD - size: The size of the data disk. Unit : GiB.
worker_system_disk_category	string	No	The system disk type of worker nodes.
worker_system_disk_size	int	No	The system disk size of worker nodes. Unit : GiB.

Response information

Response line

```
HTTP/1.1 202 Accepted
```

Custom response headers

None. For more information, see [Public response headers](#).

Response body

```
{ "cluster_id":"string",
  "request_id":"string",
```

```
"task_id":"string",
"instanceId":"string"
}
```

Examples

Sample requests

```
POST /api/v2/clusters/Cccfd68c474454665ace07efce924**** HTTP/1.1
<Common request headers> {  "key_pair":"common", // KeyPair
  "vswitch_ids":["vsw-uf684tfrpwup8gcs****"],
  "worker_instance_types": ["ecs.c5.xlarge"],
  "worker_system_disk_category":"cloud_efficiency",
  "worker_system_disk_size":120,
  "worker_data_disk":false,
  "worker_data_disks": [{"category": "cloud_ssd", "size": 500}],
  "tags":[],
  "count":1
}
```

Sample responses

```
HTTP/1.1 202 Accepted
<Common response headers> {  "cluster_id": "Cccfd68c474454665ace07efce924****",
  "request_id": "687C5BAA-D103-4993-884B-C35E4314A1E1",
  "task_id": "T-5a54309c80282e39ea00002f",
  "instanceId": "Cccfd68c474454665ace07efce924****"
}
```

3.7 Scale out a managed edge cluster

You can call the **ScaleOutCluster** operation to add worker nodes to an edge cluster.

Currently, you can only add Edge Node Service (ENS) instances to edge clusters.

Request information

Request line

```
POST /api/v2/clusters/{cluster_id} HTTP/1.1
```

Table 3-11: Request parameters

Parameter	Type	Required	Description
cluster_id	string	Yes	The ID of the cluster that you want to scale out.

Custom request headers

None. For more information, see [Public request headers](#).

Request body

```
{ "timeout_mins": cluster creation timeout,
  "worker_instance_type": ENS instance type,
  "ens_region_id": ENS region ID, "worker_image_id": operating system image used
  by worker nodes, "ens_internet_charge_type": billing method of worker nodes, "
  worker_period": subscription duration of worker nodes, "worker_auto_renew": whether
  to automatically renew worker nodes, "worker_auto_renew_period": subscription
  duration when worker nodes are renewed, "login_password": worker node logon
  password,
  "count": "number of worker nodes to be added" "worker_system_disk_size": system
  disk size of worker nodes, "worker_data_disk_size": "data disk size of worker nodes", "
  is_edge_worker": whether the worker nodes are edge nodes }
```

Table 3-12: Request body parameters

Parameter	Type	Required	Description
count	int	Yes	The number of worker nodes to be added.
worker_instance_type	string	Yes	The ENS instance type.
ens_region_id	string	Yes	The region ID of ENS instances.
worker_image_id	string	Yes	The operating system image used by worker nodes.
ens_internet_charge_type	string	Yes	This parameter is required at first purchase. If you have existing ENS instances, the current billing method is used. - BandwidthByDay: Daily Peak 95BandwidthByMonth: Monthly 95th Percentile

Parameter	Type	Required	Description
worker_period	int	Yes	The subscription duration of worker nodes. Unit: months . Valid values: 1, 2, 3 , 4, 5, 6, 7, 8, 9, and 12.
login_password	string	Yes	The password of worker nodes.
worker_system_disk_size	int	Yes	The system disk size of worker nodes. Unit : GiB. Valid values: 20 to 100. The value must be a multiple of 10. The system disk size must be greater than that of the operating system image.
is_edge_worker	bool	Yes	Set the value to true.
worker_data_disk_size	int	No	The data disk size of worker nodes. A value of 0 indicates that no data disk is required. Otherwise , specify a value that is a multiple of 10 in the range of 20 to 200.
worker_auto_renew	bool	No	Whether to automatically renew worker nodes. Valid values: - A value of True indicates yes. - A value of False indicates no. Default value: False.

Parameter	Type	Required	Description
worker_auto_renew_period	int	No	The subscription duration when worker nodes are renewed. This parameter is required when parameter worker_auto_renew is set to True. Valid values: Valid values: 1 to 12.
timeout_mins	int	No	The timeout period.

Response information

Response line

```
HTTP/1.1 202 Accepted
```

Custom response headers

None. For more information, see [Public response headers](#).

Response body

```
{
  "cluster_id": "string",
  "request_id": "string",
  "task_id": "string"
}
```

Examples

Sample requests

```
POST /api/v2/clusters/Cccfd68c474454665ace07efce924**** HTTP/1.1
<Common request headers> { "timeout_mins": 60,
  "worker_instance_type": "ens.sn1.tiny",
  "ens_region_id": "cn-beijing-telecom",
  "worker_image_id": "m-2QVLO2T8NYgm8CNQVg15gF",
  "ens_internet_charge_type": "BandwidthByDay",
  "worker_period": 1,
  "worker_auto_renew": true,
  "worker_auto_renew_period": 1,
  "login_password": "Hello1234!",
  "count": 1,
  "worker_system_disk_size": 20,
  "worker_data_disk_size": 0,
  "is_edge_worker": true
```

```
}
```

Sample responses

```
HTTP/1.1 202 Accepted
<Common response headers> {  "cluster_id": "Cccfd68c474454665ace07efce924****",
  "request_id": "687C5BAA-D103-4993-884B-C35E4314A1E1",
  "task_id": "T-5a54309c80282e39ea00002f"
}
```

3.8 Expand a Kubernetes cluster that supports sandboxed-containers

You can call **ScaleOutCluster** to describe how to add worker nodes to a Kubernetes cluster that supports sandboxed-containers. The cluster can be deployed across multiple zones.

Request information

Request line

```
POST /api/v2/clusters/{cluster_id} HTTP/1.1
```

Table 3-13: Request parameters

Parameter	Type	Required	Description
cluster_id	string	Yes	The ID of the cluster.

Custom request headers

None. For more information, see [Public request headers](#).

Request body

```
{
  "login_password": "logon password. Select one between password authentication and key pair authentication",
  "worker_instance_charge_type": "PostPaid",
  "worker_vswitch_ids": "VSwitch IDs. Specify one to three VSwitches",
  "worker_instance_types": "instance type of worker nodes",
  "worker_system_disk_category": "system disk type of worker nodes",
  "worker_system_disk_size": "system disk size of worker nodes",
  "worker_data_disk": "whether to mount data disks. true/false",
  "worker_data_disk_category": "data disk type of worker nodes",
  "worker_data_disk_size": "data disk size of worker nodes",
  "tags": "an array of labels",
  "count": "number of worker nodes to be added"
```

}

Table 3-14: Request body parameters

Parameter	Type	Required	Description
count	int	Yes	The number of worker nodes to be added to the cluster.
login_password	string	Yes	The passwords of the worker nodes to be added. The password must be 8 to 30 characters in length and contain three of the following four types of characters : uppercase letters , lowercase letters, digits, and special characters. Select one between login_password and key_pair.
key_pair	string	Yes	The keypair name. Select one between login_password and key_pair.
vswitch_ids	list	Yes	The VSwitch IDs of the worker nodes.
worker_data_disk	bool	Yes	Whether to mount data disk to worker nodes. - A value of true indicates yes. - A value of false indicates no. To enable sandboxed-containers, you must set the value to true.

Parameter	Type	Required	Description
worker_data_disk_size	string	Yes	The data disk size in GiB. To enable sandboxed-containers, you must mount a data disk of 200 GiB at least.
worker_instance_types	list	Yes	The type of instance that supports sandboxed-containers. Valid values: - ecs.ebmg5s.24xlarge - ecs.ebmc5s.24xlarge - ecs.ebmgn6i.24xlarge
worker_data_disk_category	string	No	The data disk type of worker nodes.
worker_system_disk_category	string	No	The system disk type of worker nodes.
worker_system_disk_size	int	No	The system disk size of worker nodes. Unit : GiB.

Return information

Response line

```
HTTP/1.1 202 Accepted
```

Custom response headers

None. For more information, see [Public response headers](#).

Response body

```
{
  "cluster_id": "string",
  "request_id": "string",
  "task_id": "string"
}
```

```
}
```

Example

Sample request

```
POST /api/v2/clusters/Cccfd68c474454665ace07efce924**** HTTP/1.1
<Common request headers>
{
  "login_password":"xxxxxxx",
  "worker_instance_charge_type":"PostPaid",
  "vswitch_ids":["vsw-2zes3rfz7bmk0nxxxxxxx"],
  "worker_instance_types":["ecs.ebmg5s.24xlarge"],
  "worker_system_disk_category":"cloud_efficiency",
  "worker_system_disk_size":120,
  "worker_data_disk":true,
  "worker_data_disk_category":"cloud_efficiency",
  "worker_data_disk_size":200,
  "tags":[],
  "count":1,
  "disable_rollback":false
}
```

Sample response

```
HTTP/1.1 202 Accepted
<Common response headers>
{
  "cluster_id": "Cccfd68c474454665ace07efce924****",
  "request_id": "687C5BAA-D103-4993-884B-C35E4314A1E1",
  "task_id": "T-5a54309c80282e39ea00002f"
}
```

3.9 Obtain the cluster kubeconfig file

You can call **DescribeClusterUserKubeconfig** to describe how to obtain the kubeconfig file that is used to configure access to the Kubernetes cluster. The file contains the identity information of the current user.

Request information

Request line

```
GET /k8s/:clusterid/user_config
```

Special request header

None. See [Public request headers](#).

Response information

Response line

```
HTTP/1.1 200 OK
```

Special response header

None. See [Public response headers](#).

Response body

```
{  
  "config": "string"  
}
```

Response body explanation

Name	Type	Description
config	String	Kubeconfig file used by the current user to access the cluster.

Example

Request example

```
GET /k8s/c5b5e80b0b64a4bf6939d2d8fbbc5ded7/user_config HTTP/1.1  
<Public request header>
```

Response example

```
HTTP/1.1 200 Ok  
<Public response header>  
{  
  "config": "xxxxxxx"
```

```
}
```

3.10 Change the cluster tag

You can call **ModifyClusterTags** to change the cluster tag.

Request information

Request line

```
POST /clusters/{cluster_id}/tags HTTP/1.1
```

Table 3-15: Request parameters

Parameter	Type	Required	Description
cluster_id	string	Yes	The ID of the cluster.

Custom request headers

None. For more information, see [Public request headers](#).

Request body

```
[ {  "key": "xxxxxxx",
    "value": "xxxxxxx"
}]
```

Table 3-16: Request body parameters

Parameter	Type	Required	Description
key	string	No	The key of the new tag.
value	string	No	The value of the new tag.

Response information

Response line

```
HTTP/1.1 200 OK
```

Custom response headers

None. For more information, see [Public response headers](#).

Examples

Sample requests

```
POST /clusters/c5b5e80b0b64a4bf6939d2d8fbbc5****/tags HTTP/1.1
<Common request headers>
[ {  "key": "xxxxxxx",
    "value": "xxxxxxx" } ]
```

Sample responses

```
HTTP/1.1 200 Ok
<Common response headers>
```

3.11 Query resource tags

You can call **ListTagResources** to query tags that are attached to one or more resources.

Request line

```
GET /tags HTTPS|HTTP
```

Request parameters

Parameter	Type	Required	Example	Description
resource_type	String	Yes	CLUSTER	The type of the resource.
next_token	String	No	1d2db86sca4384811e0b5e8707e*****	The start token for the next query.
resource_ids	String	No	["xxxxx","xxxxx"]	The list of cluster IDs that you want to query.
tags	String	No	[{"key":"env","value":"dev"}, {"key":"dev", "value":"IT"}]	The list of tags to be queried , which is a JSON string and contains up to 20 subitems.



Note:

- Specify either parameter resource_ids or tags. Otherwise, an exception occurs.
- The query results only include visible tags.

Response parameters

Parameter	Type	Example	Description
request_id	String	F99407AB-2FA9-489E-A259-40CF6D*****	The ID of the request.
tag_resources	Struct	-	The group of tags and resources.
tag_resource	Array	-	The tag returned for the specified resource.
resource_id	String	xxxxxx	The ID of the cluster.
resource_type	String	ALIYUN::CS::CLUSTER	The type of the resource.
tag_key	String	env	The key of the returned tag.
tag_value	String	dev	The value of the returned tag.

Examples

Sample requests

```
GET /tags?
resource_ids=["xxx","yyy"]
&resource_type=CLUSTER
&tags=[{"key":"env","value":"dev"}, {"key":"dep","value":"IT"}]
```

Sample success responses

XML format

```
<request_id>3D8795D9-8FF5-46B2-86E6-E3B4073169BE</request_id>
<tag_resources>
  <tag_resource>
    <resource_type>ALIYUN::CS::CLUSTER</resource_type>
    <resource_id>xxxxx****</resource_id>
    <tag_key>env</tag_key>
    <tag_value>dev</tag_value>
  </tag_resource>
</tag_resources>
```

JSON format

```
{
  "request_id": "3D8795D9-8FF5-46B2-86E6-E3B4073169BE",
  "tag_resources": {
    "tag_resource": [
      {
```

```
    "resource_type": "ALIYUN::CS::CLUSTER",
    "resource_id": "xxxx****",
    "tag_key": "env",
    "tag_value": "dev"
  }
}
```

3.12 Add an ECS instance to a cluster

You can call the `DescribeEdgeClusterAttachScripts` operation to add an existing ECS instance to a cluster.



Note:

When you add an ECS instance to a cluster, the system disk of the instance is replaced. We recommend that you back up data in advance.

Request information

Request line

```
POST /clusters/{cluster_id}/attach HTTP/1.1
```

Table 3-17: Request parameters

Parameter	Type	Required	Description
cluster_id	string	Yes	The ID of the cluster that you want to add nodes to.

Custom request headers

None. For more information, see [Public request headers](#).

Request body

```
{ "password": "root password of the ECS instance", "instances": "instances that you want to add to the cluster", "format_disk": "whether to format data disks", "key_pair": "key pair" "tags": "node tags in an array" }
```

Table 3-18: Request body parameters

Parameter	Type	Required	Description
password	String	Yes	The password for logging on to the ECS instance. The password must be 8 to 30 characters in length and contain three of the following four types of characters : uppercase letters , lowercase letters, digits, and special characters.
instances	Array	Yes	The instances that you want to add to the cluster.
format_disk	bool	No	Whether to format data disks.

Parameter	Type	Required	Description
key_pair	string	No	The name of the key pair. - For Windows instances, this parameter is ignored. Default is empty. The password parameter takes effect even if this parameter is set. - For Linux instances, if this parameter is set, password authentication is disabled.
tags	list	No	The tags of the nodes. - key: The name of the tag. - value: The value of the tag.

Response information

Response line

```
HTTP/1.1 202 OK
```

Custom response headers

None. For more information, see [Public response headers](#).

Response body

```
{
  "list": [
    {
      "code": "200",
```

```

        "instanceId": "i-2zee3oiwcyoz7kwd****",
        "message": "successful"
      },
      {
        "code": "200",
        "instanceId": "i-2ze0lgm3y6iylcbt****",
        "message": "successful"
      }
    ],
    "task_id": "T-5a544aff80282e39ea000039"
  }

```

Examples

Sample requests

```

POST /clusters/Cccfd68c474454665ace07efce924****/attach HTTP/1.1
<Common request headers> {  "password": "Hello1234",
  "tags": [],
  "instances": [
    "i-xxxx",
    "i-yyyy"
  ]
}

```

Sample responses

```

HTTP/1.1 202 Accepted
<Common response headers> {
  "list": [
    {
      "code": "200",
      "instanceId": "i-xxxx",
      "message": "successful"
    },
    {
      "code": "200",
      "instanceId": "i-yyyy",
      "message": "successful"
    }
  ],
  "task_id": "T-5a544aff80282e39ea000039"
}

```

```
}
```

3.13 Add a node to a managed edge cluster

You can call the DescribeEdgeClusterAttachScripts operation to obtain a script. You can then execute the script on an existing node to add the node to a managed edge cluster.

Request information

Request line

```
POST /clusters/{cluster_id}/attachscript HTTP/1.1
```

Table 3-19: Request parameters

Parameter	Type	Required	Description
cluster_id	string	Yes	The ID of the cluster that you want to add nodes to.

Custom request headers

None. For more information, see [Public request headers](#).

Request body

```
{ "options": "{
  \"flanneliface\": \"network interface used by flannel\", \"enableIptables\": \"whether
to enable the iptables proxy mode\", \"manageRuntime\": \"whether to use edgeadm to
manage the runtime\", \"nodeNameStrategy\": \"node naming strategy\",
  \"nodeName\": \"node name\",
  \"nodeNamePrefix\": \"node name prefix\",
  \"enabledAddons\": [
    \"kube-proxy\",
    \"flannel\",
    \"coredns\"
  ]
}"
```

Table 3-20: Request body parameters

Parameter	Type	Required	Description
options	json	Yes	The node configurations.

Parameter	Type	Required	Description
enabledAddons	array	No	The add-ons to be installed. Default is empty, indicating that no add-on is installed. For worker nodes, set this parameter to ["kube-proxy","flannel","coredns"].
enableIptables	bool	No	Whether to enable iptables. Default is true.
flannelIface	string	No	The network interface used by flannel. Default is eth0.
gpuVersion	string	No	The GPU version of the node. Default is empty. Currently, Nvidia_Tesla_T4 is supported.
manageRuntime	bool	No	Whether to use edgeadm to install and check the runtime.
nodeName	string	No	The name of the node. This parameter takes effect only when parameter <code>nodename-strategy</code> is set to customized.

Parameter	Type	Required	Description
nodeNamePrefix	string	No	The prefix of the node name. - By default, the hostname is used as the prefix. - Otherwise, the parameter is set to a prefixed UUID.
nodeNameStrategy	bool	No	The node naming strategy. Valid values: - hostname: The hostname of the node. - random: A random string. - randomWithPrefix : A random string with a specified prefix. - customized: A custom name. Default is hostname.

Response information

Response line

```
HTTP/1.1 202 OK
```

Custom response headers

None. For more information, see [Public response headers](#).

Response body

```
"wget http://aliacs-k8s-cn-hangzhou.oss-cn-hangzhou.aliyuncs.com/public/pkg/run/attach/1.12.6-aliyunedge.1/edgeadm -O edgeadm; chmod u+x edgeadm; ./edgeadm join --openapi-token=XXX --node-spec="{\"flanneliface\":\"eth0\", \"enableiptables\":
```

```
true,\"assumeYes\":true,\"manageRuntime\":true,\"nodeNameStrategy\": \"hostname\", \"enabledAddons\": [\"kube-proxy\", \"flannel\", \"coredns\"]}"
```

Examples

Sample requests

```
POST /clusters/Cccfd68c47445466XXXX/attachscript HTTP/1.1
<Common request headers> {  "options": "{
  \"flannelface\": \"eth0\",
  \"enableIptables\": true,
  \"manageRuntime\": true,
  \"nodeNameStrategy\": \"hostname\",
  \"enabledAddons\": [
    \"kube-proxy\",
    \"flannel\",
    \"coredns\"
  ]
}"
}
```

Sample responses

```
HTTP/1.1 202 Accepted
<Common response headers> "wget http://aliacs-k8s-cn-hangzhou.oss-cn-hangzhou
.aliyuncs.com/public/pkg/run/attach/1.12.6-aliyunedge.1/edgeadm -O edgeadm;
chmod u+x edgeadm; ./edgeadm join --openapi-token=XXX --node-spec="{ \"flannelface\": \"eth0\", \"enableIptables\": true, \"assumeYes\": true, \"manageRuntime\": true, \"nodeNameStrategy\": \"hostname\", \"enabledAddons\": [\"kube-proxy\", \"flannel\", \"coredns\"]}"
```

3.14 Add existing ENS instances to a managed edge cluster

You can call **AttachInstances** to add existing Edge Node Service (ENS) instances to a managed edge cluster.

Request information

```
POST /clusters/{cluster_id}/attach HTTP/1.1
```

Table 3-21: Request line parameters

Parameter	Type	Required	Description
cluster_id	string	Yes	The ID of the cluster.

Custom request headers

No custom request headers are used. For more information, see [Public request headers](#).

Request body

```
{
```

```
"instances": "A list of the ENS instances that you want to add.",
"is_edge_worker": "Specifies whether to configure the worker node as an edge node.",
}
```

Table 3-22: Request body parameters

Parameter	Type	Required	Description
instances	Array	Yes	A list of the ENS instances that you want to add.
is_edge_worker	bool	Yes	Specifies whether to configure the worker node as an edge node. Set the value to true.

Response information

Response line

```
HTTP/1.1 202 OK
```

Custom response headers

No custom response headers are used. For more information, see [Public response headers](#).

Response body

```
{
  "list": [
    {
      "code": "200",
      "instanceId": "i-5j2zjis****",
      "message": "successful"
    },
    {
      "code": "200",
      "instanceId": "i-5j443uf****",
      "message": "successful"
    }
  ],
  "task_id": "T-5d6f733d9e408ec74f000002"
}
```

Examples

Sample requests

```
POST /clusters/Cccfd68c474454665ace07efce924****/attach HTTP/1.1
<Common request header>
{
  "is_edge_worker": true
}
```

```
"instances": [
  "i-xxxx",
  "i-yyyy"
]
```

Sample success responses

```
HTTP/1.1 202 Accepted
<Common response header>
{
  "list": [
    {
      "code": "200",
      "instanceId": "i-xxxx",
      "message": "successful"
    },
    {
      "code": "200",
      "instanceId": "i-yyyy",
      "message": "successful"
    }
  ],
  "task_id": "T-5d6f733d9e408ec74f000002"
}
```

3.15 View a cluster

You can call **DescribeClusterDetail** to view the cluster details according to the cluster ID.

Request Information

Request line (RequestLine)

```
GET /clusters/{cluster_id} HTTP/1.1
```

Request line parameter (URI Param)

Name	Type	Required	Description
cluster_id	string	Yes	Custer ID

Special request header (RequestHead)

None. See [Public request headers](#).

Request body requestbody

None.

Return information

Response line (ResponseLine)

```
HTTP/1.1 200 OK
```

Special response header (ResponseHead)

None. See [Public response headers](#).

Response body (ResponseBody)

```
{
  "agent_version": "string",
  "cluster_id": "string",
  "created": "datetime",
  "external_loadbalancer_id": "string",
  "master_url": "string",
  "name": "string",
  "network_mode": "string",
  "region_id": "string",
  "security_group_id": "string",
  "size": "numbers",
  "state": "string",
  "updated": "datetime",
  "vpc_id": "string",
  "vswitch_id": "string"
}
```

Response body explanation

Cluster format

Name	Type	Description
agent_version	string	The Agent version.
cluster_id	string	The cluster ID, which is the unique identifier of the cluster.
created	string	The creation time of the cluster.
external_loadbalancer_id	string	The Server Load Balancer instance ID of the cluster.
master_url	string	The master address of the cluster, which is used to connect to the cluster to perform operations. For more information, see #unique_29 .

Name	Type	Description
name	string	The cluster name, which is specified when you create the cluster and is unique for each account.
network_mode	string	The network mode of the cluster (Virtual Private Cloud (VPC)).
region_id	string	The ID of the region where the cluster is located.
security_group_id	string	The security group ID.
size	string	The number of nodes.
state	string	The cluster status.
updated	string	The last update time.
vpc_id	string	The VPC ID.
vswitch_id	string	The VSwitch ID.

Examples

Request example

```
GET /clusters/C5b5e80b0b64a4bf6939d2d8fbbc5ded7 HTTP/1.1
<Public request header>
```

Response example

```
HTTP/1.1 200 Ok
<Public response header>
{
  "agent_version": "0.5-e56dab3",
  "cluster_id": "c978ca3eaacd3409a9437db07598f****",
  "created": "2015-12-11T03:52:40Z",
  "external_loadbalancer_id": "1518f2b7e4c-cn-beijing-bt****",
  "master_url": "https://182.92.245.56:17589",
  "name": "my-python-cluster-039de960",
  "network_mode": "vpc",
  "region_id": "cn-beijing",
  "security_group_id": "sg-25yqjuxhz",
  "size": 5,
  "state": "running",
  "updated": "2015-12-15T15:01:58Z",
  "vpc_id": "",
  "vswitch_id": ""
}
```

```
}
```

3.16 View all clusters

You can call **DescribeClusters** to view all the clusters you have created in Container Service, including swarm clusters and Kubernetes clusters.

Request information

Request line (RequestLine)

```
GET /clusters HTTP/1.1
```

Special request header (RequestHead)

None. See [Public request headers](#).

Request body (RequestBody)

None.

Response information

Response line (ResponseLine)

```
HTTP/1.1 200 OK
```

Special response header (ResponseHead)

None. See [Public response headers](#).

Response body (ResponseBody)

```
[
  {
    "agent_version": "string",
    "cluster_id": "string",
    "created": "datetime",
    "external_loadbalancer_id": "string",
    "master_url": "string",
    "name": "string",
    "network_mode": "string",
    "region_id": "string",
    "security_group_id": "string",
    "size": "numbers",
    "state": "string",
    "updated": "datetime",
    "vpc_id": "string",
    "vswitch_id": "string"
  }
]
```

Response body explanation

Cluster format

Name	Type	Description
agent_version	string	The Agent version.
cluster_id	string	The cluster ID, which is the unique identifier of the cluster.
created	string	The creation time of the cluster.
external_loadbalancer_id	string	The Server Load Balancer instance ID of the cluster.
master_url	string	The master address of the cluster, which is used to connect to the cluster to perform operations. For more information, see #unique_31 .
name	string	The cluster name, which is specified when you create the cluster and is unique for each account.
network_mode	string	The network mode of the cluster (Virtual Private Cloud (VPC)).
region_id	string	The ID of the region where the cluster is located.
security_group_id	string	The security group ID.
size	string	The number of nodes.
state	string	The cluster status.
updated	string	The last update time.
vpc_id	string	The VPC ID.
vswitch_id	string	The VSwitch ID.

Example

Request example

```
GET /clusters HTTP/1.1
```

<Public request header>

Response example

```
HTTP/1.1 200 OK
<Public response header>
[
  {
    "agent_version": "0.5-e56dab3",
    "cluster_id": "c978ca3eaacd3409a9437db07598f1f69",
    "created": "2015-12-11T03:52:40Z",
    "external_loadbalancer_id": "1518f2b7e4c-cn-beijing-btc-a01",
    "master_url": "https://182.92.245.56:17589",
    "name": "my-python-cluster-039de960",
    "network_mode": "vpc",
    "region_id": "cn-beijing",
    "security_group_id": "sg-25yqjuxhz",
    "size": 5,
    "state": "running",
    "updated": "2015-12-15T15:01:58Z",
    "vpc_id": "",
    "vswitch_id": ""
  },
  {
    "agent_version": "0.5-e56dab3",
    "cluster_id": "c1eb19e0093204cbb86c3a80334d2129e",
    "created": "2015-12-15T14:26:58Z",
    "external_loadbalancer_id": "151a6099de1-cn-beijing-btc-a01",
    "master_url": "https://182.92.245.56:11905",
    "name": "my-test-cluster-002b3f3d",
    "network_mode": "vpc",
    "region_id": "cn-beijing",
    "security_group_id": "sg-25rg2ws9f",
    "size": 1,
    "state": "running",
    "updated": "2015-12-15T14:43:55Z",
    "vpc_id": "",
    "vswitch_id": ""
  }
]
```

3.17 Remove a node

You can call the DeleteClusterNodes operation to remove worker nodes from a Kubernetes cluster.



Note:

- When you remove a node, pods running on the node will be migrated to other nodes, which may cause service interruptions. We recommend that you remove nodes during off-peak hours.
- Unexpected errors may occur when you remove a node. We recommend that you back up your data in advance.

- The node that you choose to remove will be set to the unschedulable state.
- You can remove worker nodes only.

Request information

Request line

```
POST /clusters/{cluster_id}/nodes HTTP/1.1
```

Request parameters

Parameter	Type	Required	Description
cluster_id	string	Yes	The ID of the cluster from which you want to remove nodes.

Custom request headers

None. For more information, see [Public request headers](#).

Request body

```
{ "nodes": "names of nodes that you want to remove in an array", "release_node": "whether to release corresponding ECS instances when nodes are removed" }
```

Request body parameters

Parameter	Type	Required	Description
nodes	Array	Yes	The nodes that you want to remove.
release_node	bool	No	Whether to release corresponding ECS instances when nodes are removed . By default, ECS instances are not released.

Response information

Response line

```
HTTP/1.1 202 OK
```

Custom response headers

None. For more information, see [Public response headers](#).

Examples

Sample requests

```
POST /clusters/Cccfd68c474454665ace07efce924****/attach HTTP/1.1
<Common request headers> {  "password": "Hello1234",
    "nodes": [
      "i-2ze74qyw9s3ge4k0****",
      "i-2zeipuf9l4tj659q****"
    ]
  }
```

Sample responses

```
HTTP/1.1 202 Accepted
<Common response headers>
```

3.18 Delete a cluster

You can call **DeleteCluster** to describe how to delete a cluster according to the cluster ID and release all node resources.

Request information

Request line

```
DELETE /clusters/{cluster_id} HTTP/1.1
```

Request line parameter (URI Param)

Name	Type	Required	Description
cluster_id	String	Yes	Custer ID

Special request header

None. See [Public request headers](#).

Request body

None.

Response information

Response line

```
HTTP/1.1 202 Accepted
```

Special response header

None. See [Public response headers](#).

Response body

None

Example**Request example**

```
DELETE /clusters/Cccfd68c474454665ace07efce924f75f HTTP/1.1  
<Public request header>
```

Response sample

```
HTTP/1.1 202 Accepted  
<Public response header>
```