

Alibaba Cloud Aliyun Container for Kubernetes Developer Guide

Issue: 20190801

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 content of the Alibaba Cloud website, including but not limited to works, products, images, archives, information, materials, website architecture, website graphic layout, and webpage design, 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 the Alibaba Cloud website, product programs, or content 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.

Generic conventions

Table -1: Style conventions

Style	Description	Example
	This warning information 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.
	This warning information 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 restore business.
	This indicates warning information, supplementary instructions, and other content that the user must understand.	 Notice: Take the necessary precautions to save exported data containing sensitive information.
	This indicates supplemental instructions, best practices, tips, and other content that is good to know for the user.	 Note: You can use Ctrl + A to select all files.
>	Multi-level menu cascade.	Settings > Network > Set network type
Bold	It is used for buttons, menus, page names, and other UI elements.	Click OK.
Courier font	It is used for commands.	Run the <code>cd / d C :/ windows</code> command to enter the Windows system folder.
<i>Italics</i>	It is used for parameters and variables.	<code>bae log list --instanceid Instance_ID</code>
[] or [a b]	It indicates that it is an optional value, and only one item can be selected.	<code>ipconfig [-all -t]</code>

Style	Description	Example
<code>{}</code> or <code>{a b}</code>	It indicates that it is a required value, and only one item can be selected.	<code>swich {stand slave}</code>

Contents

Legal disclaimer.....	I
Generic conventions.....	I
1 Cluster API call method.....	1
1.1 Overview.....	1
1.2 Common parameters.....	1
1.3 Request structure.....	4
1.4 Return results.....	4
1.5 Signature.....	4
2 Cluster API list.....	11
2.1 Create a Kubernetes cluster.....	11
2.2 Create a managed Kubernetes cluster.....	23
2.3 Obtain the cluster kubeconfig file.....	30
2.4 Add existing ECS instances to a Kubernetes cluster.....	31
2.5 Download the cluster certificate.....	35
2.6 View a cluster.....	36
2.7 View all clusters.....	38
2.8 Delete a cluster.....	41
3 Use Container Service through CLI.....	43
3.1 View all clusters.....	43
3.2 View cluster information.....	43
3.3 Create a Kubernetes cluster.....	44
3.4 Expand a cluster.....	47
3.5 Add existing ECS instances to a cluster.....	48
3.6 Delete a cluster.....	49

1 Cluster API call method

1.1 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](#)

1.2 Common parameters

Public request headers

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

Parameter	Description	Options
Authorizat ion	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. applicatio n / json and applicatio n / xml are supported.	Required
x - acs - version	The API version. The current version is 2015 - 12 - 15 .	Required
x - acs - region - id	A region indicates the physical location of an Elastic Compute Service (ECS) instance.	Required
x - acs - signature - nonce	The unique random number used to prevent network replay attacks. Different random numbers must be used for different requests.	Required

Parameter	Description	Options
x - acs - 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 : applicatio n / 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 -
ea8fde61ea e3
Authorizat ion : acs < yourAccess KeyId >:< yourSignat ure >
x - acs - signature - version : 1 . 0
Date : Wed , 16 Dec 2015 11 : 18 : 47 GMT
x - acs - signature - method : HMAC - SHA1
Content - Type : applicatio n / 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 - AC049166F2 16 </
RequestId >
| <!-- Return result data -->
</ Interface name + Response >
```

JSON example:

```
{
  " RequestId ": " 4C467B38 - 3910 - 447D - 87BC - AC049166F2 16 "
  /* Response data */
}
```

```
}
```

1.3 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.

1.4 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.

1.5 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 "
+ Canonicali zedHeaders + "\ n "
+ Canonicali zedResourc e ))
```

- `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`.
- `Canonicali zedHeaders` indicates a combination of fields started with `x -` `acs -` 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: 6U4ALMkKSj0PYbeQSHqg mA ==
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-253l****l"}
```

Request construction process

Calculate `Content-Length` and `Content-MD5`

`Content-Length`: The length of the body content.



Note:

No space or line break is at the beginning of the example body.

```
body : {" password ": " Just $****"," instance_t ype ": " ecs . m2
. medium "," name ": " my - test - cluster - 9708 ****"," size ":
1 ," network_mo de ": " vpc "," data_disk_ category ": " cloud ","
data_disk_ size ": 10 ," ecs_image_ id ": " m - 253ll ****"}
Content - Length : 210
```

Content - MD5 : The MD5 calculation process.

```
body : {" password ": " Just $****"," instance_t ype ": " ecs . m2
. medium "," name ": " my - test - cluster - 9708 ****"," size ":
1 ," network_mo de ": " vpc "," data_disk_ category ": " cloud ","
data_disk_ size ": 10 ," ecs_image_ id ": " m - 253ll ****"}
# Calculate the MD5 value of the body .
md5 ( body ): e94e002cc9 0a4a3d0f61 b790487aa0 98
# Convert the MD5 value to a byte array . Convert
every two hexadecimal symbols of the MD5 value to
a byte .
# For example , e9 -> 1111111111 1111111111 111111010 01 ->
- 23
bytes ( md5 ( body )): {[ - 23 ], [ 78 ], [ 0 ], [ 44 ], [ - 55 ], [ 10
], [ 74 ], [ 61 ], [ 15 ], [ 97 ], [ - 73 ], [ - 112 ], [ 72 ], [ 122
], [ - 96 ], [ - 104 ]}
# Convert the obtained byte array to a Base64 string
base64 ( bytes ( md5 ( body ))) : 6U4ALMkKSj 0PYbeQSHqg mA ==
Content - MD5 : 6U4ALMkKSj 0PYbeQSHqg mA ==
```

Process Canonicali zedHeaders

```
# List all headers started with ' x - acs - '.
x - acs - version : 2015 - 12 - 15
x - acs - signature - nonce : ca480402 - 7689 - 43ba - acc4 -
4d2013d9d8 d4
x - acs - signature - version : 1 . 0
x - acs - signature - method : HMAC - SHA1
X - Acs - Region - Id : cn - beijing
# Convert the request name to lowercase letters , delete
the spaces at the beginning and end of each line
, and sort the headers in the lexicograp hically
ascending order . Delete any space at either side of
a separator between request header and content .
# Note : No line break is in the last line .
x - acs - region - id : cn - beijing
x - acs - signature - method : HMAC - SHA1
x - acs - signature - nonce : fbf6909a - 93a5 - 45d3 - 8b1c -
3e03a79167 99
x - acs - signature - version : 1 . 0
x - acs - version : 2015 - 12 - 15
```

Calculate Canonicali zedResourc e

In the example, the length of Canonicali zedResourc e is 27.



Note:

An \ n line break is at the end of the first line.

```
/ clusters ? param1 = value1 & param2 = value2
```

Calculate Signature

Assemble SignatureS tring . 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
applicatio n / json
6U4ALMkKSj 0PYbeQSHqg mA ==
applicatio n / json ; charset = utf - 8
Wed , 16 Dec 2015 12 : 20 : 18 GMT
x - acs - region - id : cn - beijing
x - acs - signature - method : HMAC - SHA1
x - acs - signature - nonce : fbf6909a - 93a5 - 45d3 - 8b1c -
3e03a79167 99
x - acs - signature - version : 1 . 0
x - acs - version : 2015 - 12 - 15
/ clusters ? param1 = value1 & param2 = value2
```

Calculate Signature

```
# Use Access Key Secret to encrypt the signature
string . In the example , the accessKeyS cret is
access_key _secret .
hmac - sha1 ( SignatureS tring ) : fee03d405e 421ebaf514
adec881038 c4b313584d
# Convert the encrypted string to a byte array ,
similar to the Content - MD5 calculatio n method .
# Convert the byte array into a Base64 string to
get the final signature string .
base64 ( bytes ( hmac - sha1 ( SignatureS tring ))) : ZmVlMDNkND
A1ZTQyMWVi YWY1MTRhZG VjODgxMDM4 YzRiMzEzNT g0ZA ==
Signature : ZmVlMDNkND A1ZTQyMWVi YWY1MTRhZG VjODgxMDM4
YzRiMzEzNT g0ZA ==
```

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 : 6U4ALMkKSj 0PYbeQSHqg mA ==
x - acs - version : 2015 - 12 - 15
Accept : applicatio n / 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 -
3e03a79167 99
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 - 253ll ****" }
```

2 Cluster API list

2.1 Create a Kubernetes cluster

This topic describes how to create a Kubernetes cluster and a specified number of nodes.

Request information

Request line

```
POST / clusters HTTP / 1 . 1
```

Special request header

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

Request body

```
{
  "disable_rollback": "whether to roll back the
cluster if the cluster creation fails "
  "name": "cluster name "
  "timeout_mins": "cluster creation timeout"
  "cluster_type": "cluster type, namely, Kubernetes "
  "region_id": "region "
  "vpcid": "Virtual Private Cloud (VPC) ID "
  "master_vswitch_ids": "the IDs of VSwitches on
Master nodes. Set three VSwitches that are located in
different regions to ensure the cluster is highly
available ",
  "master_instance_types": "the type of ECS instances
that used by the Master nodes. Set three types of
ECS instances "
  "master_count": "the number of Master nodes. Valid
value: 3 | 5 ."
  "container_cidr": "pod Classless Inter - Domain Routing
(CIDR) block "
  "service_cidr": "service CIDR block "
  "ssh_flags": "whether to enable SSH access over the
Internet "
  "cloud_monitor_flags": "whether to install the cloud
monitoring plugin "
  "login_password": "password used to log on to the
node by using SSH. Use either this parameter or
the key_pair ."
  "key_pair": "key pair name. Use either this
parameter or login_password ."
  "master_instance_charge_type": "Master instance payment
type that includes PostPaid and PrePaid "
  "master_period_unit": "Subscription unit that which
includes month and year, and takes effect only for
the PrePaid payment type "
```

```

    "master_period ":" Subscription period that takes
effect only for the PrePaid payment type "
    "master_auto_renew ":" whether to enable master node
automatic renew "
    "master_auto_renew_period ":" Master node renew period
"
    "master_system_disk_category ":" Master node system
disk type "
    "master_system_disk_size ":" Master system disk size "
    "master_data_disk ":" whether to mount data disk to
the Master node "
    "master_data_disk_category ":" Master node data disk
type "
    "master_data_disk_size ":" Master node data disk
size "
    "worker_instance_charge_type ":" Worker node payment
type that includes PrePaid and PostPaid "
    "worker_period_unit ":" Subscription unit that includes
Month and Year, and takes effect only for the
PrePaid payment type "
    "worker_period ":" Subscription period that takes
effect only for the PrePaid payment type "
    "worker_auto_renew ":" whether to enable worker node
automatic renew . Available values are true and false ."
    "worker_auto_renew_period ":" Worker node renew period
"
    "worker_instance_types ":" Worker instance types "
    "worker_vswitch_ids ":" IDs of VSwitches . You must
set one VSwitch at least , and can set five
VSwitches at most ",
    "worker_system_disk_category ":" Worker node system
disk type "
    "worker_system_disk_size ":" Worker node system disk
size "
    "worker_data_disk ":" whether to mount data disks to
the Worker node "
    "worker_data_disk_category ":" Worker node data disk
type "
    "worker_data_disk_size ":" Worker node data disk size
"
    "num_of_nodes ":" number of Worker nodes "
    "snat_entry ":" whether to set an SNAT entry "
    "public_slb ":" whether to associate an EIP with the
intranet SLB instance ",
    "cpu_policy ":" static | none ",
    "node_port_range ":" the range of the node port .
The default is 30 , 000 to 65 , 535 ",
    "proxy_mode ":" the proxy mode . Valid value : iptables
| ipvs ",
    "addons ":" the addon . This is an object in the
array format .",
    "tags ":" the tag to be add to the cluster . This
is an object in the array format .",
}

```

Request body description

Name	Type	Required	Description
cluster_type	String	Yes	The cluster type.

Name	Type	Required	Description
key_pair	String	Yes	key pair name. Use either this parameter or login_password .
login_password	String	Yes	SSH logon password. The password must be a string of 8 to 30 characters and contain uppercase letters, lowercase letters, numbers, and symbols. Choose either this parameter or key_pair .
master_instance_charge_type	String	Yes	Master node payment type <ul style="list-style-type: none"> PrePaid: Subscription. PostPaid: Pay-As-You-Go. This is the default payment type.
master_instance_types	list	Yes	The types of ECS instances that used by the Master nodes. For more information, see Instance type families .
master_system_disk_category	String	Yes	The type of system disks attached to Master nodes. Available values are: <ul style="list-style-type: none"> cloud_efficiency: indicates an Ultra cloud disk. cloud_ssd: indicates an SSD cloud disk.
master_system_disk_size	Int	Yes	The size of a system disk attached to Master nodes, in GiB.
master_vswitch_ids	list	Yes	The IDs of VSwitches on Master nodes. You must set one VSwitch at least, and can set three VSwitches at most. Set three VSwitches that are located in different regions to ensure the cluster is highly available.
num_of_nodes	Int	Yes	The number of Worker nodes. The value is in the range of 0 to 300.


Name	Type	Required	Description
name	String	Yes	The cluster name. A cluster name can contain uppercase letters, lowercase letters, Chinese characters, numbers, and hyphens (-).
region_id	String	Yes	The ID of the region in which the cluster is located.
snat_entry	bool	Yes	Whether to set an SNAT entry for the container network. <ul style="list-style-type: none"> · If the existing VPC has the capability to access the Internet, you need to set this parameter to false. · If the existing VPC does not have the capability to access the Internet, setting this parameter to true indicates to set an SNAT entry; setting this parameter to false indicates not to set any SNAT entry and the Internet cannot be accessed.
worker_instance_types	list	Yes	The types of ECS instances used by Worker nodes. For more information, see Instance type families .
worker_system_disk_category	String	Yes	The type of system disks attached to Worker nodes.
worker_system_disk_size	Int	Yes	The size of a system disk attached to Worker nodes, in GiB. <ul style="list-style-type: none"> ·
worker_vswitch_ids	list	Yes	The IDs of VSwitches on Worker nodes.

Name	Type	Required	Description
addons	list	No	<p>The addons installed on the Kubernetes cluster.</p> <ul style="list-style-type: none"> · An addon requires the following parameters: <ul style="list-style-type: none"> - name: Required. - version: Optional. By default, the latest version is used. - config: Optional. · Network plugin: Select Flannel or Terway. · Log Service plugin: Optional . If you do not install this plugin, the audit logs of the cluster cannot be collected.
container_cidr	String	No	<p>The pod CIDR block, which cannot overlap with the VPC CIDR block. If you choose to enable the system to automatically create a VPC, the 172.16.0.0/16 pod CIDR block is used by default.</p>
cloud_monitor_flags	bool	No	<p>Whether to install the cloud monitoring plugin.</p> <ul style="list-style-type: none"> · true: indicates to install the plugin · false: indicates not to install the plugin
cpu_policy	String	No	<p>The CPU policy. If the Kubernetes cluster is V1.12.6 or later, you can set static or none. The default is none.</p>

Name	Type	Required	Description
disable_rollback	bool	No	<p>Whether to roll back the cluster if the cluster creation fails.</p> <ul style="list-style-type: none"> · true: indicates not to roll back if the cluster creation fails. · false: indicates to roll back if the cluster creation fails. <p>The default is true. If you select to roll back, resources generated in the creation process will be released. We recommend that you do not select false.</p>
master_auto_renew	bool	No	<p>Whether to enable Master node automatic renew. This parameter setting takes effect only if the value of master_instance_charge_type is set to PrePaid . Available values are:</p> <ul style="list-style-type: none"> · true: indicates to enable automatic renew · false: indicates not to enable automatic renew
master_auto_renew_period	Int	No	<p>The automatic renew period. This parameter setting takes effect and becomes required only if you select the PrePaid payment type and enable automatic renew. Available values for different Subscription period units are as follows:</p> <ul style="list-style-type: none"> · { “1” , “2” , “3” } for PeriodUnit = Week · { “1” , “2” , “3” , “6” , “12” } for PeriodUnit = Month
master_count	Int	No	<p>The number of Master nodes. Valid value: 3 5. The default is 3.</p>

Name	Type	Required	Description
master_data_disk	bool	No	Whether to mount data disks to the Master node. Available values are: <ul style="list-style-type: none"> · true: indicates to mount data disks · false: indicates not to mount data disks. This is the default .
master_data_disk_category	String	No	Master node data disk type. This parameter setting takes effect only if data disks are mounted to the Master node. Available values are: <ul style="list-style-type: none"> · cloud: indicates a basic cloud disk · cloud_efficiency: indicates an Ultra cloud disk · cloud_ssd: indicates an SSD cloud disk
master_data_disk_size	Int	No	Master node data disk size in GiB. This parameter setting takes effect only if data disks are mounted to the Master node.
master_period	Int	No	Subscription period. This parameter setting takes effect and becomes required only if the value of <code>marker_instance_charge_type</code> is set to <code>PrePaid</code> . The parameter values for different Subscription period units are as follows: <ul style="list-style-type: none"> · { "1" , "2" , "3" , "4" } for <code>PeriodUnit = Week</code> · { "1" , "2" , "3" , "4" , "5" , "6" , "7" , "8" , "9" , "12" , "24" , "36" , "48" , "60" } for <code>PeriodUnit = Month</code>

Name	Type	Required	Description
master_period_unit	String	No	<p>If you select the PrePaid payment type, you need to specify a Subscription period that uses either of the following units:</p> <ul style="list-style-type: none"> • Week: one week is used as the timing unit. • Month: one month is used as the timing unit.
node_port_range	String	No	The range of the node port. The default is 30,000 to 65,535
proxy_mode	String	No	The kube-proxy mode. The modes iptables and ipvs are supported. The default is iptables.
public_slb	bool	No	<p>Whether or not to enable the Internet API server:</p> <ul style="list-style-type: none"> • true: indicates to enable the Internet API server. This is the default value. • false: indicates not to create the Internet API server. If you select this value, only the intranet API server will be created.
service_cidr	String	No	The service CIDR block that cannot overlap with the VPC CIDR block or the pod CIDR block. If you choose to enable the system to automatically create a VPC, the 172.19.0.0/20 service CIDR block is used by default.
ssh_flags	bool	No	<p>Whether to enable SSH access over the Internet.</p> <ul style="list-style-type: none"> • true: yes • false: no

Name	Type	Required	Description
tags	list	No	<p>The tag to be added to the cluster.</p> <ul style="list-style-type: none"> · key: tag name. · value: tag value.
timeout_mins	Int	No	The timeout (in minutes) for creating the cluster resource stack. The default value is 60.
vpcid	String	No	<p>VPC ID, which can be empty. If you do not set any VPC, the system automatically creates a VPC that belongs to the CIDR block 192.168.0.0/16.</p> <div>  Note: The vpcid parameter and the vswitchid parameter must be both set to empty or relevant values. </div>
worker_auto_renew	bool	No	<p>Whether to enable Worker node automatic renew. Available values are:</p> <ul style="list-style-type: none"> · true: indicates to enable automatic renew · false: indicates not to enable automatic renew
worker_auto_renew_period	Int	No	<p>The automatic renew period. This parameter setting takes effect and becomes required only if you select the PrePaid payment type and enable automatic renew. Available values for different Subscription period units are as follows:</p> <ul style="list-style-type: none"> · { "1", "2", "3" } for PeriodUnit = Week · { "1", "2", "3", "6", "12" } for PeriodUnit = Month

Name	Type	Required	Description
worker_data_disk	String	No	Whether to mount data disks. Available values are: <ul style="list-style-type: none"> · true: indicates to mount data disks to Worker nodes · false: indicates not to mount data disks to Worker nodes
worker_data_disk_category	Int	No	Data disk type. This parameter setting takes effect only if you mount data disks to the Worker nodes. Available values are: <ul style="list-style-type: none"> · cloud: indicate a basic cloud disk · cloud_efficiency: indicates an Ultra cloud disk · cloud_ssd: indicates an SSD cloud disk
worker_data_disk_size	String	No	Data disk size in GiB. This parameter setting takes effect only if data disks are mounted to the Worker node.
worker_instance_charge_type	String	No	Worker node payment type. The default value is PostPaid. Available values are: <ul style="list-style-type: none"> · PrePaid: indicates Subscription · PostPaid: indicates Pay-As-You-Go

Name	Type	Required	Description
worker_period	Int	No	Subscription period. This parameter setting takes effect and becomes required only if the value of <code>worker_instance_charge_type</code> is set to <code>PrePaid</code> . The parameter values for different Subscription period units are as follows: <ul style="list-style-type: none"> · { "1", "2", "3", "4" } for <code>PeriodUnit = Week</code> · { "1", "2", "3", "4", "5", "6", "7", "8", "9", "12", "24", "36", "48", "60" } for <code>PeriodUnit = Month</code>
worker_period_unit	String	No	Whether to enable Worker node automatic renewal. Available values are: <ul style="list-style-type: none"> · <code>true</code>: indicates to enable automatic renewal · <code>false</code>: indicates not to enable automatic renewal

Response information

Response line

```
HTTP / 1 . 1 202 Accepted
```

Special response header

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

Response body

Name	Type	Description
cluster_id	String	The cluster ID.
request_id	String	The request ID.

Name	Type	Description
task_id	String	The task ID. It is automatically assigned by the system. You can use it to check the status of a task.

Examples

Request example

```
POST / clusters HTTP / 1 . 1
< Public response header >
{
  " cluster_ty pe ":" Kubernetes ",
  " name ":" my - test - Kubernetes - cluster ",
  " region_id ":" cn - beijing ",
  " disable_ro llback ": true ,
  " timeout_mi ns ": 60 ,
  " kubernetes _version ":" 1 . 12 . 6 - aliyun . 1 ",
  " snat_entry ": true ,
  " public_slb ": false ,
  " cloud_moni tor_flags ": false ,
  " node_cidr_ mask ":" 25 ",
  " proxy_mode ":" iptables ",
  " tags ":[],
  " addons ":[{" name ":" flannel "}, {" name ":" nginx - ingress -
controller "}],
  " node_port_ range ":" 30000 - 32767 ",
  " login_pass word ":" test ****",
  " cpu_policy ":" none ",
  " master_cou nt ": 3 ,
  " master_vsw itch_ids ":[" vsw - 2ze48rkq46 4rsdts ****"," vsw -
2ze48rkq46 4rsdts1 ****"," vsw - 2ze48rkq46 4rsdts1 ****"],
  " master_ins tance_type s ":[" ecs . sn1 . medium "," ecs . sn1 .
medium "," ecs . sn1 . medium "],
  " master_sys tem_disk_c ategory ":" cloud_effi ciency ",
  " master_sys tem_disk_s ize ": 40 ,
  " worker_ins tance_type s ":[" ecs . sn2 . 3xlarge "],
  " num_of_nod es ": 3 ,
  " worker_sys tem_disk_c ategory ":" cloud_effi ciency ",
  " worker_sys tem_disk_s ize ": 120 ,
  " vpcid ":" vpc - 2zegvl5eta h5requ0 ****",
  " worker_vsw itch_ids ":[" vsw - 2ze48rkq46 4rsdts ****"],
  " container_ cidr ":" 172 . 20 . XX . XX / 16 ",
  " service_ci dr ":" 172 . 21 . XX . XX / 20 "
}
```

Response example

```
HTTP / 1 . 1 202 Accepted
< Public response header >
{
  " cluster_id ":" cb95aa626a 47740afbf6 aa099b650 ****",
  " request_id ":" 687C5BAA - D103 - 4993 - 884B - C35E4314A1 E1
",
  " task_id ":" T - 5a54309c80 282e39ea00 002f "
```

```
}
```

2.2 Create a managed Kubernetes cluster

This topic describes how to create a managed Kubernetes cluster. You must specify the number of Worker nodes for the cluster.

Request information

Request line

```
POST / clusters HTTP / 1 . 1
```

Special request header

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

Request body

```
{
  "disable_rollback": "whether or not to roll back if the cluster fails to be created",
  "name": "cluster name",
  "timeout_mins": "timeout for creating the cluster",
  "cluster_type": "cluster type, Managed Kubernetes",
  "region_id": "region",
  "vpcid": "Virtual Private Cloud (VPC) ID",
  "vswitch_ids": "IDs of VSwitches. You must set one VSwitch at least, and can set five VSwitches at most.",
  "container_cidr": "pod Classless Inter-Domain Routing (CIDR)",
  "service_cidr": "service CIDR",
  "cloud_monitor_flags": "whether or not to install the cloud monitoring plug-in",
  "login_password": "password used to log on to the node by using SSH. Use either this parameter or the key_pair.",
  "key_pair": "key pair name. use either this parameter or login_password.",
  "worker_instance_charge_type": "worker node payment type PrePaid | PostPaid",
  "worker_periode_unit": "subscription unit, which includes month and year, and takes effect only for the prepaid type.",
  "worker_periode": "subscription period, which takes effect only for the prepaid type",
  "worker_auto_renew": "worker node auto renew true | false",
  "worker_auto_renew_period": "worker node renew period",
  "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 or not to mount data
    disks true | false ",
    " data_disk_category ":" data disk category ",
    " worker_data_disk_size ":" data disk size ",
    " num_of_nodes ":" number of worker nodes ",
    " snat_entry ":" whether or not to configure the
    SNATEntry ",
    " public_slb ":" whether to associate an EIP with the
    intranet SLB instance ",
    " proxy_mode ":" the proxy mode . Valid value : iptables
    | ipvs ",
    " addons ":" the addon . This is an object in the
    array format .",
    " tags ":" the tag to be add to the cluster . This
    is an object in the array format .",
  }


```

Request body explanation

Name	Type	Required	Description
cluster_type	string	Yes	The cluster type.
key_pair	string	Yes	key pair name. Use either this parameter or login_password .
login_password	string	Yes	SSH logon password. The password must be a string of 8 to 30 characters and contain uppercase letters, lowercase letters, numbers, and symbols. Choose either this parameter or key_pair .
name	string	Yes	The cluster name. A cluster name can contain uppercase letters, lowercase letters, Chinese characters, numbers, and hyphens (-).
num_of_nodes	int	Yes	The number of Worker nodes. The value is in the range of 0 to 300.
region_id	string	Yes	The ID of the region in which the cluster is located.

Name	Type	Required	Description
snat_entry	bool	Yes	Whether or not to configure the SNAT for VPC. This parameter must be set to true if a VPC is created automatically. Configure the parameter value according to the outbound capability if you select to use an existing VPC to create the cluster.
vswitch_ids	list	Yes	The IDs of VSwitches. You must set one VSwitch at least, and can set three VSwitches at most.
worker_system_disk_category	string	Yes	The type of system disks attached to Worker nodes.
worker_system_disk_size	int	Yes	The size of a system disk attached to Worker nodes, in GiB.
addons	list	No	<p>The addons installed on the Kubernetes cluster.</p> <ul style="list-style-type: none"> • An addon requires the following parameters: <ul style="list-style-type: none"> - name: Required. - version: Optional. By default, the latest version is used. - config: Optional. • Network plugin: Select Flannel or Terway. • Log Service plugin: Optional. If you do not install this plugin, the audit logs of the cluster cannot be collected.
container_cidr	string	No	The pod CIDR block, which cannot overlap with the VPC CIDR block. If you choose to enable the system to automatically create a VPC, the 172.16.0.0/16 pod CIDR block is used by default.

Name	Type	Required	Description
cloud_monitor_flags	bool	No	Whether to install the cloud monitoring plugin.
disable_rollback	bool	No	Whether to install the cloud monitoring plugin. <ul style="list-style-type: none"> · true: indicates to install the plugin · false: indicates not to install the plugin
proxy_mode	string	No	The kube-proxy mode. The modes iptables and ipvs are supported. The default is iptables.
public_slb	bool	No	Whether or not to enable the Internet API server: <ul style="list-style-type: none"> · true: indicates to enable the Internet API server. This is the default value. · false: indicates not to create the Internet API server. If you select this value, only the intranet API server will be created.
service_cidr	string	No	The service CIDR block that cannot overlap with the VPC CIDR block or the pod CIDR block. If you choose to enable the system to automatically create a VPC, the 172.19.0.0/20 service CIDR block is used by default.
tags	list	No	The tag to be added to the cluster. <ul style="list-style-type: none"> · key: tag name. · value: tag value.
timeout_minutes	int	No	The timeout (in minutes) for creating the cluster resource stack. The default value is 60.

Name	Type	Required	Description
vpcid	string	No	<p>VPC ID, which can be empty. If you do not set any VPC, the system automatically creates a VPC that belongs to the CIDR block 192.168.0.0/16.</p> <div>  Note: The vpcid parameter and the vswitchid parameter must be both set to empty or relevant values. </div>
worker_auto_renew	bool	No	<p>Whether to enable Worker node automatic renewal. Available values are:</p> <ul style="list-style-type: none"> · true: indicates to enable automatic renewal · false: indicates not to enable automatic renewal
worker_auto_renew_period	int	No	<p>The automatic renewal period. This parameter setting takes effect and becomes required only if you select the PrePaid payment type and enable automatic renewal. Available values for different Subscription period units are as follows:</p> <ul style="list-style-type: none"> · { "1" , "2" , "3" } for PeriodUnit = Week · { "1" , "2" , "3" , "6" , "12" } for PeriodUnit = Month
worker_data_disk	string	No	<p>Whether to mount data disks. Available values are:</p> <ul style="list-style-type: none"> · true: indicates to mount data disks to Worker nodes · false: indicates not to mount data disks to Worker nodes

Name	Type	Required	Description
worker_data_disk_category	int	No	Data disk type.
worker_data_disk_size	string	No	Data disk size
worker_instance_charge_type	string	No	Worker node payment type. <ul style="list-style-type: none"> PrePaid: indicates Subscription PostPaid: indicates Pay-As-You-Go
worker_period	int	No	Subscription period. This parameter setting takes effect and becomes required only if the value of worker_instance_charge_type is set to PrePaid . The parameter values for different Subscription period units are as follows: <ul style="list-style-type: none"> { "1" , "2" , "3" , "4" } for PeriodUnit = Week { "1" , "2" , "3" , "4" , "5" , "6" , "7" , "8" , "9" , "12" , "24" , "36" , "48" , "60" } for PeriodUnit = Month
worker_period_unit	string	No	Whether to enable Worker node automatic renew. Available values are: <ul style="list-style-type: none"> true: indicates to enable automatic renew false: indicates not to enable automatic renew

Response information

Response line

```
HTTP / 1 . 1 202 Accepted
```

Special response header

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

Response body

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

Examples

Request example

```
POST / clusters HTTP / 1 . 1
< Public response header >
{
  " 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 ,
  " public_slb ": false ,
  " node_cidr_mask ":" 255 ",
  " proxy_mode ":" ipvs ",
  " tags ":[],
  " addons ":[{" name ":" flannel "}, {" name ":" nginx - ingress - controller "}],
  " worker_instance_type_s ":[" 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 - 2zegvl5eta h5requ09ne c ",
  " container_cidr ":" 172 . 20 . 0 . 0 / 16 ",
  " service_cidr ":" 172 . 21 . 0 . 0 / 20 ",
  " vswitch_ids ":[" vsw - 2ze48rkq46 4rsdts1 ****"],
  " login_password ":" test @ 19 ****"
}
```

Response example

```
HTTP / 1 . 1 202 Accepted
< Public response header >
{
  " cluster_id ":" cb95aa626a 47740afbf6 aa099b65 ****",
```

```
{
  " request_id ": " 687C5BAA - D103 - 4993 - 884B - C35E4314A1  E1
",
  " task_id ": " T - 5a54309c80  282e39ea00  002f "
}
```

2.3 Obtain the cluster kubeconfig file

This topic describes 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_conf i g
```

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 / c5b5e80b0b 64a4bf6939 d2d8fbbc5d ed7 / user_conf i g
HTTP / 1 . 1
```

```
< Public request header >
```

Response example

```
HTTP / 1 . 1 200 Ok
< Public response header >
{
  " config ": " xxxxxxxx "
}
```

2.4 Add existing ECS instances to a Kubernetes cluster

This topic describes how to add existing Elastic Compute Service (ECS) instances to a cluster.



Note:

The system disk is replaced in the process of adding ECS instances. Therefore, you need to back up the data in advance.

Request information

Request line

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

```
{
  " password ": " the password used by the root account
to log on to an ECS instance ",
  " instances ": " the IDs of the ECS instances that
you want to add to the target cluster .",
  " format_dis k ": " whether to format disks . Value :
true | false ",
  " key_pair ": " the key pair ",
  " tags ": " the tag added to the node . This is an
object of the array format ."
}
```

Request body parsing

Name	Type	Required	Description
password	String	Yes	Password of the ECS instance. The password must be a string of 8 to 30 characters and contain uppercase letters, lowercase letters, numbers, and symbols.
instances	Array	Yes	The IDs of the ECS instances that you want to add to the target cluster.
format_disk	Bool	No	Whether to format a disk.

Name	Type	Required	Description
key_pair	String	No	<p>Set the name of the a key pair.</p> <ul style="list-style-type: none">· If you add an ECS instance that runs the Windows OS , ignore this parameter. By default, this parameter is set as null. Any settings of this parameter do not take effect or affect the password.· If you add an ECS instance that runs the Linux OS, you can set this parameter . In addition, the password parameter is initialized as forbidden.
tags	List	No	<p>Add a tag to the node.</p> <ul style="list-style-type: none">· key: indicates the tag name.· value: indicates the tag value.

Response information

Response line

```
HTTP / 1 . 1    202    OK
```

Special response header

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

Response body

```
{
  " list ": [
    {
      " code ": " 200 ",
      " instanceId ": " i - 2zee3oiwcy  oz7kwdo8bt ",
      " message ": " successful "
    },
    {
      " code ": " 200 ",
      " instanceId ": " i - 2ze0lgm3y6  iylcbtcypf ",
      " message ": " successful "
    }
  ],
  " task_id ": " T - 5a544aff80  282e39ea00  0039 "
}
```

Example

Request example

```
POST / clusters / Cccfd68c47  4454665ace  07efce924f  75f / attach
HTTP / 1 . 1
< Public  request  header >
{
  " password ": " Hello1234 ",
  " tags ": [],
  " instances ": [
    " i - xxxx ",
    " i - yyyy "
  ]
}
```

Response sample

```
HTTP / 1 . 1    202    Accepted
< Public  response  header >
{
  " list ": [
    {
      " code ": " 200 ",
      " instanceId ": " i - xxxx ",
      " message ": " successful "
    },
    {
      " code ": " 200 ",
      " instanceId ": " i - yyyy ",

```

```
      " message ": " successful "
    },
    " task_id ": " T - 5a544aff80 282e39ea00 0039 "
  }
}
```

2.5 Download the cluster certificate

This topic describes how to download the certificate used by the currently logged on user to access the cluster.

Request information

Request line

```
GET / clusters /: clusterid / certs
```

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 description

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

Examples

Request example

```
GET / clusters / C5b5e80b0b 64a4bf6939 d2d8fbbc5d ed7 / certs
HTTP / 1 . 1
```

```
< Public request header >
```

Response example

```
HTTP / 1 . 1 200 Ok
< Public response header >
{
  " config ": " xxxxxxxx "
}
```

2.6 View a cluster

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_vers ion ": " string ",
  " cluster_id ": " string ",
  " created ": " datetime ",
  " external_l oadbalance _r_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 Access Kubernetes clusters by using SSH .
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.

Name	Type	Description
updated	string	The last update time.
vpc_id	string	The VPC ID.
vswitch_id	string	The VSwitch ID.

Examples

Request example

```
GET / clusters / C5b5e80b0b 64a4bf6939 d2d8fbbc5d ed7 HTTP / 1
. 1
< Public request header >
```

Response example

```
HTTP / 1 . 1 200 Ok
< Public response header >
{
  " agent_vers ion ": " 0 . 5 - e56dab3 ",
  " cluster_id ": " c978ca3eaa cd3409a943 7db07598f1 f69 ",
  " created ": " 2015 - 12 - 11T03 : 52 : 40Z ",
  " external_l oadbalance r_id ": " 1518f2b7e4 c - cn - beijing -
btc - a01 ",
  " master_url ": " https :// 182 . 92 . 245 . 56 : 17589 ",
  " name ": " my - python - cluster - 039de960 ",
  " network_mo de ": " vpc ",
  " region_id ": " cn - beijing ",
  " security_g roup_id ": " sg - 25yqjuxhz ",
  " size ": 5 ,
  " state ": " running ",
  " updated ": " 2015 - 12 - 15T15 : 01 : 58Z ",
  " vpc_id ": "",
  " vswitch_id ": ""
}
```

2.7 View all clusters

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.

Name	Type	Description
master_url	string	The master address of the cluster, which is used to connect to the cluster to perform operations. For more information, see Connect to a Kubernetes cluster by using kubectl .
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. For more information, see Cluster lifecycle .
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 ": " c978ca3eaa cd3409a943 7db07598f1 f69 ",
    " created ": " 2015 - 12 - 11T03 : 52 : 40Z ",
```



```

    "external_loadbalance_r_id": "1518f2b7e4 c - 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": "cleb19e009 3204cbb86c 3a80334d21 29e ",
    "created": "2015 - 12 - 15T14 : 26 : 58Z ",
    "external_loadbalance_r_id": "151a6099de 1 - 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": ""
  }
]

```

2.8 Delete a cluster

This topic describes 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 / Cccfd68c47 4454665ace 07efce924f 75f HTTP /  
1 . 1  
< Public request header >
```

Response sample

```
HTTP / 1 . 1 202 Accepted  
< Public response header >
```

3 Use Container Service through CLI

3.1 View all clusters

This topic describes how to view all the clusters that you have created through Alibaba Cloud Container Service for Kubernetes.

API request and response

Request format

```
aliyun  cs  GET  / clusters
```

Response result

```
[
  {
    "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"
  }
]
```

3.2 View cluster information

This topic describes how to view the cluster details according to the cluster ID. For API descriptions, see [View a cluster](#).

API request and response

Request format

```
aliyun  cs  GET  / clusters /< cluster_id >
```

Response result

```
{
  "agent_version": "string",
  "cluster_id": "string",
}
```

```

    "created": "datetime",
    "external_loadbalance_r_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"
  }

```

3.3 Create a Kubernetes cluster

This topic describes how to create a Kubernetes cluster. For more information about the related APIs, see [Create a Kubernetes cluster](#), [Create a multi-zone Kubernetes cluster](#), and [Create a multi-zone Kubernetes cluster](#).

API request and response

Request format

```

aliyun cs POST / clusters --header "Content-Type =
application/json" --body "$(cat create.json)"

```

Parameter descriptions:

- `--header`: Specify Content-Type as application/json.
- `--body`: This is the body content to be sent to the server. The content can be read from a local file and must be in JSON format. The content of `create.json` is as follows:

Dedicated Kubernetes cluster

```

{
  "disable_rollback": "disable_rollback": "indicates
whether to roll back the cluster if the cluster
creation fails",
  "name": "cluster name",
  "timeout_mins": "cluster creation timeout",
  "cluster_type": "Kubernetes",
  "region_id": "region",
  "vpcid": "VPC ID",
  "master_vswitch_ids": "the IDs of VSwitches on
Master nodes. You must set three VSwitches that are
located in different zones to ensure high availability",
  "master_instance_types": "the types of ECS
instances used by Master nodes. You must set three
instance types.",
  "master_count": "the number of Master nodes. Valid
values: 3 | 5."
}

```

```

" container_ cidr ": " pod Classless Inter - Domain Routing
( CIDR ) block ",
" service_ci dr ": " service CIDR block ",
" ssh_flags ": " enable SSH access over the Internet .
Valid values : true | false .",
" cloud_moni tor_flags ": " install the CloudMonit or
plugin . Valid values : true | false .",
" login_pass word ": " password used to log on to
cluster nodes by using SSH . This parameter and the
key_pair parameter are mutually exclusive .",
" key_pair ": " the key pair used to log on to
cluster nodes by using SSH . This parameter and the
login_pass word parameter are mutually exclusive .",
" master_ins tance_char ge_type ": " the billing method of
Master node instances . Valid values : PostPaid | PrePaid
",
" master_per iod_unit ": " Subscripti on unit that includes
Month and Year . This parameter takes effect only
for the PrePaid billing method ",
" master_per iod ": " Subscripti on period , which takes
effect only for the PrePaid billing method ",
" master_aut o_renew ": " enable Master node automatic
renew . Valid values : true | false .",
" master_aut o_renew_pe riod ": " Master node renew period
",
" master_sys tem_disk_c ategory ": " Master node system
disk type ",
" master_sys tem_disk_s ize ": " Master node system disk
size ",
" master_dat a_disk ": " mount data disks to the Master
node . Valid values : true | false .",
" master_dat a_disk_cat egory ": " Master node data disk
type ",
" master_dat a_disk_siz e ": " Master node data disk size
",
" worker_ins tance_char ge_type ": " the billing method of
Worker node instances . Valid values : PostPaid | PrePaid
",
" worker_per iod_unit ": " Subscripti on unit that includes
Month and Year . This parameter takes effect only
for the PrePaid billing method ",
" worker_per iod ": " Subscripti on period , which takes
effect only for the PrePaid billing method ",
" worker_aut o_renew ": " enable Worker node automatic
renew . Valid values : true | false .",
" worker_aut o_renew_pe riod ": " Worker node renew period
",
" worker_ins tance_type s ": " Worker node instance types
. ",
" worker_vsw itch_ids ": " the IDs of VSwitches . The
value range of the number of VSwitches is 1 to 5
. ",
" worker_sys tem_disk_c ategory ": " Worker node system
disk type ",
" worker_sys tem_disk_s ize ": " Worker node system disk
size ",
" worker_dat a_disk ": " mount data disks to the Worker
node . Valid values : true | false .",
" worker_dat a_disk_cat egory ": " Worker node data disk
type ",
" worker_dat a_disk_siz e ": " Worker node data disk size
",
" num_of_nod es ": " the number of Worker nodes ",

```

```

    "snat_entry": "Configure an SNAT entry. Valid values
: true | false.",
    "public_slb": "associate an EIP with the intranet
SLB instance. Valid values: true | false.",
    "cpu_policy": "static | none",
    "node_port_range": "value range of node ports. The
default value range is 30000 to 32767",
    "proxy_mode": "iptables | ipvs",
    "addons": "set addons. This is an array object.",
    "tags": "set tags for the cluster. This is an
array object.",
}

```

Managed Kubernetes cluster

```

{
  "disable_rollback": "indicates whether to roll back the
cluster if the cluster creation fails"
  "name": "cluster name"
  "timeout_mins": "cluster creation timeout"
  "cluster_type": "Managed Kubernetes"
  "region_id": "region."
  "vpcid": "VPC ID"
  "vswitch_ids": "the IDs of VSwitches. The value range
of the number of VSwitches is 1 to 5.",
  "container_cidr": "pod CIDR block"
  "service_cidr": "service CIDR block"
  "cloud_monitor_flags": "whether to install the cloud
monitoring plugin"
  "login_password": "password used to log on to the
node by using SSH. Use either this parameter or the
key_pair."
  "key_pair": "key pair name. use either this parameter
or login_password."
  "worker_instance_charge_type": "Worker node payment type
that includes PrePaid and PostPaid"
  "worker_period_unit": "Subscription unit that includes
Month and Year. This parameter takes effect only for
the PrePaid payment type"
  "worker_period": "Subscription period, which takes
effect only for the PrePaid payment type."
  "worker_auto_renew": "whether to enable Worker node
automatic renew. Available values are true and false."
  "worker_auto_renew_period": "Worker node renew period"
  "worker_instance_types": "Worker instance types"
  "worker_system_disk_category": "Worker node system disk
type"
  "worker_system_disk_size": "Worker node system disk
size"
  "worker_data_disk": "whether to mount data disks to
the worker node. Available values are true and false."
  "worker_data_disk_category": "data disk type"
  "worker_data_disk_size": "data disk size"
  "num_of_nodes": "number of worker nodes"
  "snat_entry": "whether to set an SNAT entry"
  "public_slb": "associate an EIP with the intranet SLB
instance. Valid values: true | false.",
  "proxy_mode": "iptables | ipvs",
  "addons": "set addons. This is an array object.",
  "tags": "set tags for the cluster. This is an array
object.",
}

```

```
}
```

Serverless Kubernetes cluster

```
{
  "cluster_type": "Ask",
  "name": "indicates the name of a serverless Kubernetes cluster",
  "region_id": "indicates a region",
  "nat_gateway": "indicates whether to create a NAT gateway. Valid value: true | false.",
  "private_zone": "indicates whether to enable PrivateZone for service discovery in a VPC. Valid value: true | false.",
  "vpc_id": "indicates the ID of a VPC. If you do not set this parameter, then the system automatically creates a VPC.",
  "vswitch_id": "indicates a Vswitch ID"
}
```

Response result

```
{
  "cluster_id": "c61cf53052 4474386a7a b5a1c192 ****",
  "request_id": "348D4C9C - 9105 - 4A1B - A86E - B58F0F8755 75",
  "task_id": "T - 5ad724ab94 a2b109e800 0004"
}
```

3.4 Expand a cluster

This topic describes how to increase the number of nodes in the cluster. For API descriptions, see [Scale out or in a cluster](#).

API request and response

Request format

```
aliyun cs PUT / clusters /< cluster_id > -- header "Content - Type = application / json" -- body "$( cat scale . json )"
```

Parameter descriptions:

- **-- header** : Specify Content - Type as application / json .
- **-- body** : This is the body content to be sent to the server. The content can be read from a local file and must be in JSON format. The content of `scale . json` is as follows:

Kubernetes cluster

```
{ "disable_rollback ": " indicates whether to roll back
the cluster if cluster expansion fails ",
" timeout_mins ": " timeout period for creating a cluster
",
" worker_instance_type ": " Worker instance type ",
" login_password ": " password used to log on to the
node by using SSH ",
" num_of_nodes ": " number of Worker nodes "
}
```

Response result

```
{
  " cluster_id ": " c61cf53052 4474386a7a b5a1c192a0 d57 ",
  " request_id ": " 348D4C9C - 9105 - 4A1B - A86E - B58F0F8755 75
",
  " task_id ": " T - 5ad724ab94 a2b109e800 0004 "
}
```

3.5 Add existing ECS instances to a cluster

This topic describes how to add existing ECS instances to a cluster. For API descriptions, see [Add existing ECS instances to a Kubernetes cluster](#).

API request and response

Request format

```
aliyun cs POST / clusters /< cluster_id >/ attach -- header "
Content - Type = applicatio n / json " -- body "$( cat attach .
json )"
```

Parameter description:

- **-- header** : Specify Content - Type as applicatio n / json .
- **-- body** : This is the body content to be sent to the server. The content can be read from a local file and must be in JSON format. The content of `attach . json` is as follows:

```
{
  " password ": " password used to log on to the
ECS instance by using SSH ",
  " instances ": " ECS instance array ",
  " ecs_image_id ": " image ID ",
  " release_eip_flag ": " whether to release Elastic
IP ( EIP ) after you configure the cluster "
}
```



```
}
```

Response result

```
{
  "list": [
    {
      "code": "200",
      "instanceId": "i-2zee3oiwcy oz7kwdo8bt",
      "message": "successful"
    },
    {
      "code": "200",
      "instanceId": "i-2ze0lgm3y6 iylcbtcypf",
      "message": "successful"
    }
  ],
  "task_id": "T-5a544aff80 282e39ea00 0039"
}
```

3.6 Delete a cluster

This topic describes how to delete a cluster according to the cluster ID, and release all node resources of the cluster. For API descriptions, see [Delete a cluster](#).

API request and response

Request format

```
aliyun cs DELETE / clusters /< cluster_id >
```

Response result

None.