Alibaba Cloud
Alibaba Cloud Container Service for Kubernetes

FAQ

Issue: 20190802

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 informatio n, 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 `cd / d C :/ windows` command to enter the Windows system folder.
`Italics`	It is used for parameters and variables.	`bae log list --instanceid` *`Instance_ID`*
[] or [a|b]	It indicates that it is a optional value, and only one item can be selected.	`ipconfig` *`[-all|-t]`*

Style	Description	Example
{} or {a|b}	It indicates that it is a required value, and only one item can be selected.	`swich` *{stand | slave}*

Contents

1 General FAQ

Q: What is Alibaba Cloud Container Service for Kubernetes?

Alibaba Cloud Container Service for Kubernetes (ACK) is a fully-managed service compatible with Kubernetes. ACK helps you focus on your applications, eliminating the need to manage the container infrastructure.

It provides enterprise-level high-performance and flexible management of Kubernetes containerized applications throughout the application lifecycle. This service simplifies cluster creation and expansion and integrates Alibaba Cloud capabilities in virtualization, storage, network, and security, providing an improved running environment for Kubernetes containerized applications.

Q: When should I use ACK?

ACK fits perfectly for scenarios such as DevOps continuous delivery, Microservice architecture, hybrid cloud architecture and auto-scaling. Here we take auto scaling as an example: ACK can automatically scale out or in the application according to the incoming traffic, without manual intervention. In this way, the system is elastic and stable when facing traffic surge, and resource utility rate gets improved as well.

Q: How does ACK work?

Container Service for Kubernetes is adapted and enhanced on the basis of native Kubernetes. This service simplifies cluster creation, scaling and upgrade, and integrates Alibaba Cloud capabilities in computing, storage, network, and security, providing an improved running environment for Kubernetes containerized applications.

Q: What are the advantages of ACK over a local host in creating a Kubernetes cluster?

- Ease of use
  - Supports creating Kubernetes clusters with one click in the Container Service console.
  - Supports upgrading Kubernetes clusters with one click in the Container Service console.

  You may have to deal with self-built Kubernetes clusters of different versions at the same time, including version 1.11.x, 1.12.x, and later. Container Service upgrade solution performs rolling update by using images and uses the backup policy of complete metadata, which allows you to conveniently roll back to the previous version.

- Supports scaling Kubernetes clusters conveniently in the Container Service console.

  Container Service Kubernetes clusters allow you to expand or contract the capacity vertically with one click to respond to the peak of the data analysis business quickly.

- Powerful functions

Function	Description
Network	- High-performance Virtual Private Cloud (VPC) network plug-in. - Supports network policy and flow control. - Container Service provides you with continuous network integration and the best network optimization.
Server Load Balancer	- Supports creating Internet or intranet Server Load Balancer instances. - If your self-built Kubernetes clusters are implemented by using the self-built Ingress, releasing the business frequently may cause pressure on Ingress configuration and higher error probabilities. The Server Load Balancer solution of Container Service supports Alibaba Cloud native high-availability Server Load Balancer, and can automatically modify and update the network configurations. This solution has been used by a large number of users for a long time, which is more stable and reliable than self-built Kubernetes.
Storage	Container Service integrates with Alibaba Cloud cloud disk , Network Attached Storage (NAS), and block storage, and provides the standard FlexVolume drive. Self-built Kubernetes clusters cannot use the storage resources on the cloud. Alibaba Cloud Container Service provides the best seamless integration.

Function	Description
O&M	- Integrates with Alibaba Cloud Log Service and CloudMonitor. - Supports auto scaling.
Image repository	- Provides high availability. - Supports image scan. - Supports P2P distribution and replication across regions. The self-built image repository may crash if you pull images from millions of clients at the same time. Enhance the reliability of the image repository by using the image repository of Alibaba Cloud Container Service, which reduces the O&M burden and upgrade pressure.
Stability	- The dedicated team guarantees the stability of the container. - Each Linux version and Kubernetes version are provided to you after strict tests. - Container Service provides the Docker CE to reveal all the details and promotes the repair capabilities of Docker. If you have issues such as Docker Engine hangs, network problems, and kernel compatibility, Container Service provides you with the best practices.
High availability	- Supports multiple zones. - Supports backup and disaster recovery.
Technical support	- Provides the Kubernetes upgrade capabilities. Supports upgrading a Kubernetes cluster to the latest version with one click. - Alibaba Cloud container team is responsible for solving problems about containers in your environment.

Q: How can I get started with ACK?

Please see Workflow for more details on how to get started with ACK.

2 Failed to create a Kubernetes cluster

Check the cause of failure

You can check the cause of cluster creation failure by viewing the cluster creation events.

Log on to the Resource Orchestration Service (ROS) console.

Select the region where the cluster resides. Then, click Manage at the right of the cluster. In the left-side navigation pane, click Event. Rest the pointer on the failed event to view the specific error message of the failure.

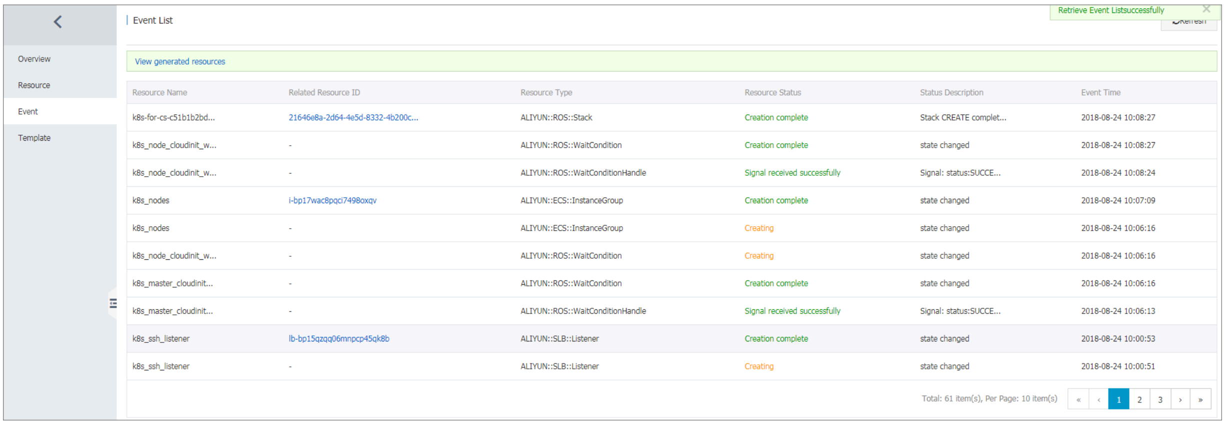

Resource Name	Related Resource ID	Resource Type	Resource Status	Status Description	Event Time
k8s-for-cs-c51b1b2bd...	21646e8a-2d64-4e5d-8332-4b200c...	ALIYUN::ROS::Stack	Creation complete	Stack CREATE complet...	2018-08-24 10:08:27
k8s_node_cloudinit_w...	-	ALIYUN::ROS::WaitCondition	Creation complete	state changed	2018-08-24 10:08:27
k8s_node_cloudinit_w...	-	ALIYUN::ROS::WaitConditionHandle	Signal received successfully	Signal: status:SUCCE...	2018-08-24 10:08:24
k8s_nodes	i-bp17wac8pqci7498oxqv	ALIYUN::ECS::InstanceGroup	Creation complete	state changed	2018-08-24 10:07:09
k8s_nodes	-	ALIYUN::ECS::InstanceGroup	Creating	state changed	2018-08-24 10:06:16
k8s_node_cloudinit_w...	-	ALIYUN::ROS::WaitCondition	Creating	state changed	2018-08-24 10:06:16
k8s_master_cloudinit...	-	ALIYUN::ROS::WaitCondition	Creation complete	state changed	2018-08-24 10:06:16
k8s_master_cloudinit...	-	ALIYUN::ROS::WaitConditionHandle	Signal received successfully	Signal: status:SUCCE...	2018-08-24 10:06:13
k8s_ssh_listener	lb-bp15qzqq06mnpcp45qk8b	ALIYUN::SLB::Listener	Creation complete	state changed	2018-08-24 10:00:53
k8s_ssh_listener	-	ALIYUN::SLB::Listener	Creating	state changed	2018-08-24 10:00:51

Total: 61 item(s), Per Page: 10 item(s)

Failure codes and solutions

- Resource CREATE failed: ResponseException: resources.k8s_SNat_Eip: Elastic IP address quota exceeded Code: QuotaExceeded.Eip

  Solution: Release unused EIPs, or open a ticket to raise the EIP quota.
- Resource CREATE failed: ResponseException: resources.k8s_master_slb_internet: The maximum number of SLB instances is exceeded. Code: ORDER.QUANTITY_INVALID

  Solution: Release unused SLB instances, or open a ticket to raise the SLB quota.
- Resource CREATE failed: ResponseException: resources.k8s_vpc: VPC quota exceeded. Code: QuotaExceeded.Vpc

  Solution: Release unused VPCs, or open a ticket to raise the VPC quota.

- Status Code: 403 Code: InvalidResourceType.NotSupported Message: This resource type is not supported;

  Solution: No ECS in stock or the type is not supported. Select other ECS specifications and try again.
- Resource CREATE failed: ResponseException: resources.k8s_master_1: The specified image does not support cloud-init. Code: ImageNotSupportCloudInit

  Solution: When a custom image is used to create a cluster, the custom image must be based on the latest CentOS image.
- Resource CREATE failed: ResponseException: resources.k8s_nodes: The resource is out of stock in the specified zone. Please try other types, or choose other regions and zones. Code: OperationDenied.NoStock

  Solution: The instances of your selected specifications are sold out. Select other availability zones or specifications, and try again.
- Resource CREATE failed: ResponseException: resources.k8s_NAT_Gateway: A route entry already exists, which CIDR is '0.0.0.0/0' Code: RouterEntryConflict.Duplicated

  Solution: Current route table of the VPC includes system route entries. Delete the system route entries, or clear the Configure SNAT for VPC check box, and try again.
- Resource CREATE failed: ResponseException: resources.KubernetesWorkerRole: The number of role is limited to 200. Code: LimitExceeded.Role

  Solution: The number of RAM roles has reached the quota. Delete some RAM roles, or open a ticket to raise the quota.
- Resource CREATE failed: ResponseException: resources.k8s_NAT_Gateway: The Account failed to create order. Code: OrderFailed

  Solution: Failed to create an order. Open a ticket for consultation.
- Resource CREATE failed: ResponseException: resources.k8s_master_1: This operation is forbidden by Aliyun RiskControl system. Code: Forbidden.RiskControl

  Solution: An exception occurs to your account. For more information, see What do I do if I get a security notification?.
- Resource CREATE failed: WaitConditionFailure: resources.k8s_node_cloudinit_wait_cond: See output value for more details.

  Solution: Failed to create a cluster. Try again later, or open a ticket for consultation.

- Resource CREATE failed: WaitConditionTimeout: resources.k8s_master1_cloudinit_wait_cond: 0 of 2 received:

  Solution: Failed to create a cluster. Try again later, or open a ticket for consultation.
- Resource CREATE failed: ResponseException: resources.k8s_master_1: The request processing has failed due to some unknown error. Code: UnknownError

  Solution: Unknown error. Try again later, or open a ticket for consultation.
- Resource CREATE failed: ResponseException: resources.k8s_nodes: The request processing has failed due to some unknown error. Code: UnknownError

  Solution: Unknown error. Try again later, or open a ticket for consultation.

3 Failed to delete Kubernetes clusters: ROS stack cannot be deleted

Root cause

Some resources are manually added (for example, manually add a VSwitch under the Virtual Private Cloud (VPC) created by Resource Orchestration Service (ROS)) under the resources created by ROS. ROS does not have permissions to delete those resources. This causes ROS to fail to process the VPC when deleting the Kubernetes resources and then the cluster fails to be deleted.

> **Note:**
> For more information about the resources automatically created by ROS when the Kubernetes cluster is created, see Create a Kubernetes cluster.

Solutions

1. If the cluster fails to be deleted (the cluster status is Failed to delete), go to the ROS console.

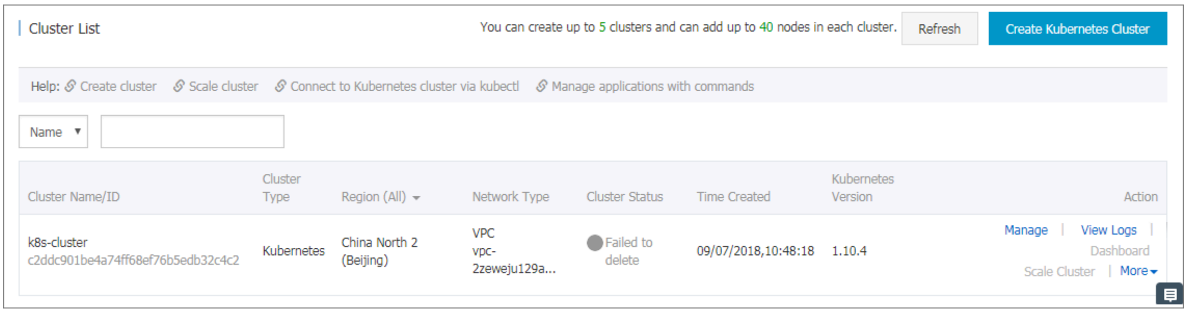
Cluster List
Refresh
Create Kubernetes Cluster
Help: Create cluster
Scale cluster
Connect to Kubernetes cluster via kubectl
Manage applications with commands
Name ▾

You can create up to 5 clusters and can add up to 40 nodes in each cluster.

Cluster Name/ID	Cluster Type	Region (All)	Network Type	Cluster Status	Time Created	Kubernetes Version	Action
k8s-cluster c2ddc901be4a74ff68ef76b5edb32c4c2	Kubernetes	China North 2 (Beijing)	VPC vpc-2zeweju129a...	● Failed to delete	09/07/2018,10:48:18	1.10.4	Manage | View Logs | Dashboard Scale Cluster | More ▾

2. Select the region in which the cluster resides and find the stack `k8s - for - cs -{ cluster - id }` corresponding to the cluster. You can see the status is Failed to delete.

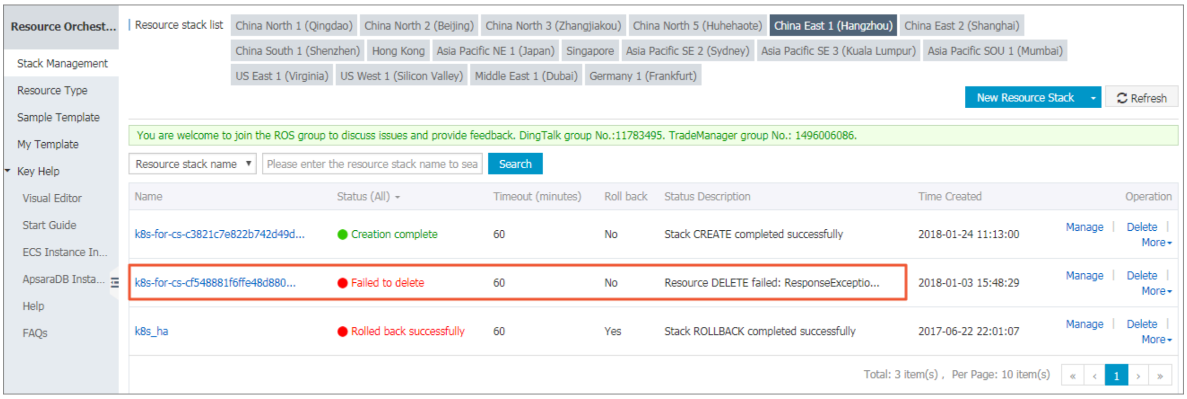

You are welcome to join the ROS group to discuss issues and provide feedback. DingTalk group No.:11783495. TradeManager group No.: 1496006086.

Name	Status (All)	Timeout (minutes)	Roll back	Status Description	Time Created	Operation
k8s-for-cs-c3821c7e822b742d49d...	● Creation complete	60	No	Stack CREATE completed successfully	2018-01-24 11:13:00	Manage | Delete | More ▾
k8s-for-cs-cf548881f6ffe48d880...	● Failed to delete	60	No	Resource DELETE failed: ResponseExceptio...	2018-01-03 15:48:29	Manage | Delete | More ▾
k8s_ha	● Rolled back successfully	60	Yes	Stack ROLLBACK completed successfully	2017-06-22 22:01:07	Manage | Delete | More ▾

Total: 3 item(s) , Per Page: 10 item(s) « ‹ 1 › »

3. Click the stack name to go to the stack details page. Click Resource in the left-side navigation pane.

   You can see what resources failed to be deleted. In this example, the VSwitch under Server Load Balancer failed to be deleted.

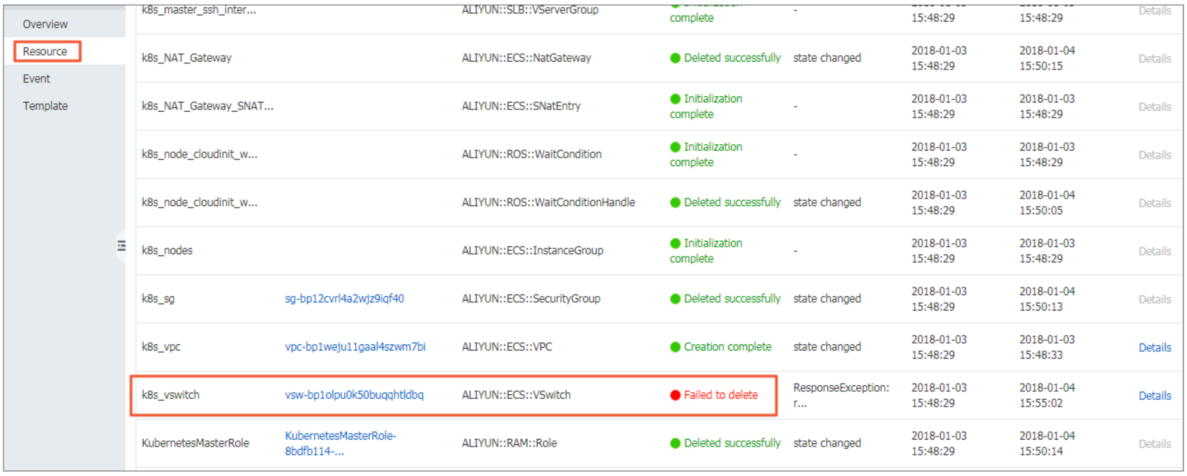

k8s_master_ssh_inter...		ALIYUN::SLB::VServerGroup	● complete	-	15:48:29	15:48:29	Details
k8s_NAT_Gateway		ALIYUN::ECS::NatGateway	● Deleted successfully	state changed	2018-01-03 15:48:29	2018-01-04 15:50:15	Details
k8s_NAT_Gateway_SNAT...		ALIYUN::ECS::SNatEntry	● Initialization complete	-	2018-01-03 15:48:29	2018-01-03 15:48:29	Details
k8s_node_cloudinit_w...		ALIYUN::ROS::WaitCondition	● Initialization complete	-	2018-01-03 15:48:29	2018-01-03 15:48:29	Details
k8s_node_cloudinit_w...		ALIYUN::ROS::WaitConditionHandle	● Deleted successfully	state changed	2018-01-03 15:48:29	2018-01-04 15:50:05	Details
k8s_nodes		ALIYUN::ECS::InstanceGroup	● Initialization complete	-	2018-01-03 15:48:29	2018-01-03 15:48:29	Details
k8s_sg	sg-bp12cvrl4a2wjz9iqf40	ALIYUN::ECS::SecurityGroup	● Deleted successfully	state changed	2018-01-03 15:48:29	2018-01-04 15:50:13	Details
k8s_vpc	vpc-bp1weju11gaal4szwm7bi	ALIYUN::ECS::VPC	● Creation complete	state changed	2018-01-03 15:48:29	2018-01-03 15:48:33	Details
k8s_vswitch	vsw-bp1olpu0k50buqqhtldbq	ALIYUN::ECS::VSwitch	● Failed to delete	ResponseException: r...	2018-01-03 15:48:29	2018-01-04 15:55:02	Details
KubernetesMasterRole	KubernetesMasterRole-8bdfb114-...	ALIYUN::RAM::Role	● Deleted successfully	state changed	2018-01-03 15:48:29	2018-01-04 15:50:14	Details

4. Go to the console in which the resource that failed to be deleted resides and find that resource.

   In this example, log on to the VPC console and find the VPC in which the cluster resides. Find the VSwitch that failed to be deleted under that VPC.

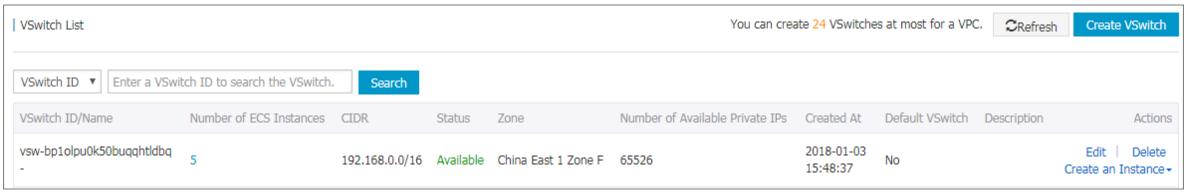
VSwitch List
You can create 24 VSwitches at most for a VPC.
Refresh
Create VSwitch
VSwitch ID ▾
Enter a VSwitch ID to search the VSwitch.
Search

VSwitch ID/Name	Number of ECS Instances	CIDR	Status	Zone	Number of Available Private IPs	Created At	Default VSwitch	Description	Actions
vsw-bp1olpu0k50buqqhtldbq -	5	192.168.0.0/16	Available	China East 1 Zone F	65526	2018-01-03 15:48:37	No		Edit | Delete Create an Instance ▾

5. Click Delete at the right of the VSwtich to manually delete it.

   In this example, the VSwitch has resources to release and cannot be deleted.

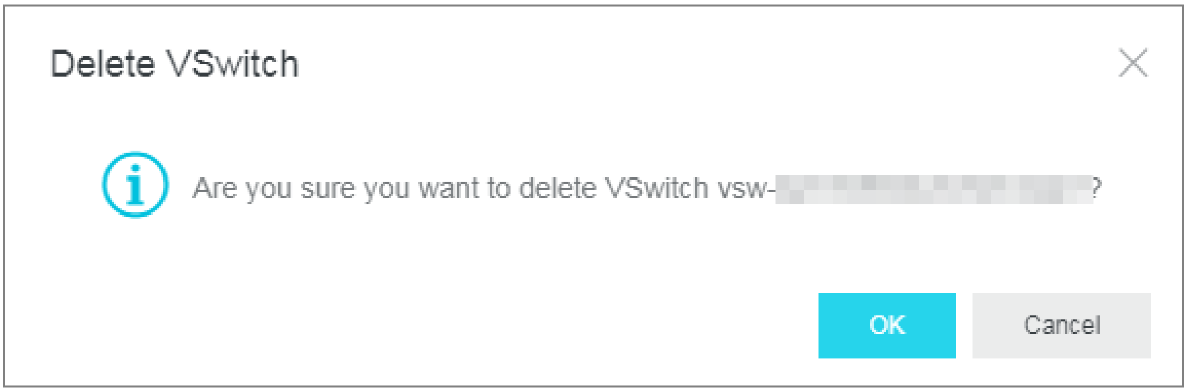

Are you sure you want to delete VSwitch vsw-[redacted]?

   Manually release the resources under this VSwitch and try to delete this VSwitch again.

6. **Manually delete all the resources that failed to be deleted under the Kubernetes cluster in this way and try to delete the Kubernetes cluster again.**

4 How do I collect Kubernetes diagnosis information when a Kubernetes cluster exception or a cluster node exception occurs?

Context

If exceptions occur to the Kubernetes cluster, you need to collect diagnosis information on the Master node.

If exceptions occur to Worker nodes, you need to collect diagnosis information on the Master node and the abnormal Worker nodes.

Procedure

1. Download the diagnosis script on the Master and Worker nodes, and add the permission to run the script.

```
curl  - o  / usr / local / bin / diagnose_k  8s . sh   http ://
aliacs - k8s - cn - hangzhou . oss - cn - hangzhou . aliyuncs . com
/ public / diagnose / diagnose_k  8s . sh
chmod   u + x  / usr / local / bin / diagnose_k  8s . sh
```

2. Run the diagnosis script.

```
diagnose_k  8s . sh
......
+  echo  ' please   get   diagnose_1  514939155 . tar . gz   for
diagnostic  s '   ## The   name   of   a   generated   log   file
varies   each   time   when   you   run   the   diagnosis   script
.
 please   get   diagnose_1  514939155 . tar . gz   for
diagnostic  s
+  echo  ' upload   diagnose_1  514939155 . tar . gz '
 Upload   the   diagnose_1  514939155 . tar . gz   file .
```

3. List and upload the generated log file.

```
cd  / usr / local / bin
ls  - ltr | grep   diagnose_1  514939155 . tar . gz          ##
Replace   this   example   file   name   with   the   name   of
your   generated   log   file .
```

5 Upgrade Helm manually

Log on to the master node of the Kubernetes cluster, see Connect to a Kubernetes cluster by using kubectl.

Execute the following command:

```
helm  init  -- tiller - image  registry . cn - hangzhou . aliyuncs
. com / acs / tiller : v2 . 11 . 0  -- upgrade
```

The image address can use the VPC domain name of the region corresponding to the image. For example, the image address of a machine in the Hangzhou region can be replaced by registry-vpc.cn-hangzhou.aliyuncs.com/acs/tiller:v2.11.0.

Wait for `tiller` passing through health check. Then you can execute `helm version` to view the upgraded version.

Note:

Only the Helm server version is upgraded here. To use the Helm client, download the corresponding client binary.

Helm 2.11.0 client download address: https://github.com/helm/helm/releases/tag/v2.11.0. Currently, the latest version of Helm supported by Alibaba Cloud is 2.11.0.

After the Helm client and server are both upgraded, you can see the following information by executing the `helm  version` :

```
$  helm  version
 Client : & version . Version { SemVer :" v2 . 11 . 0 ",  GitCommit :"
 2e55dbe1fd  b5fdb96b75  ff144a3394  89417b146b ",  GitTreeSta  te :"
 clean "}
 Server : & version . Version { SemVer :" v2 . 11 . 0 ",  GitCommit :"
 2e55dbe1fd  b5fdb96b75  ff144a3394  89417b146b ",  GitTreeSta  te :"
 clean "}
```

6 How to use private images in Kubernetes clusters

```
kubectl   create   secret   docker - registry   regsecret  -- docker
- server = registry - internal . cn - hangzhou . aliyuncs . com  --
docker - username = abc @ aliyun . com  -- docker - password = xxxxxx
 -- docker - email = abc @ aliyun . com
```

where:

- regsecret: Specifies the secret key name and the name is customizable.
- —docker-server: Specifies the Docker repository address.
- —docker-username: Specifies the user name of the Docker repository.
- —docker-password: Specifies the logon password of the Docker repository.
- —docker-email: Specifies the email address (optional).

Add secret key parameters in the YML file.

```
containers :
   -  name :  foo
      image :  registry - internal . cn - hangzhou . aliyuncs . com /
abc / test : 1 . 0
imagePullS  ecrets :
   -  name :  regsecret
```

where:

- `imagePullS  ecrets` declares that a secret key must be specified when you pull the image.
- `regsecret` must be the same as the preceding secret key name.
- The Docker repository name in `image` must be the same as that in `-- docker - server`.

For more information, see the official documentation Use private repository.

Implement keyless orchestration

To avoid referencing keys each time when using private images to deploy, you can add secret to the default service account of namespace. For more information, see Add ImagePullSecrets to a service account.

First find the secret created to pull the private image.

```
#  kubectl   get   secret   regsecret
 NAME           TYPE                               DATA        AGE
```

```
regsecret     kubernetes . io / dockerconf  igjson     1
13m
```

In this example, manually configure the default service account default of the namespace to use this secret as the imagePullSecret.

Create a *`sa . yaml`* configuration file to export the configurations of the service account default to this file.

```
kubectl   get   serviceacc  ounts   default  - o   yaml  > ./ sa .
yaml

cat    sa . yaml

apiVersion :  v1
kind :  ServiceAcc  ount
metadata :
  creationTi  mestamp :  2015 - 08 - 07T22 : 02 : 39Z
  name :  default
  namespace :  default
  resourceVe  rsion : " 243024 "            ## Pay   attention   to
  this   item
  selfLink : / api / v1 / namespaces / default / serviceacc  ounts /
default
  uid :  052fb0f4 - 3d50 - 11e5 - b066 - 42010af0d7  b6
secrets :
-  name :  default - token - uudgeoken - uudge
```

Execute the `vim   sa . yaml` command to delete resourceVersion and add the secret configuration item, imagePullSecrets which is used to pull images. The modified configuration is as follows:

```

apiVersion :  v1
kind :  ServiceAcc  ount
metadata :
  creationTi  mestamp :  2015 - 08 - 07T22 : 02 : 39Z
  name :  default
  namespace :  default
  selfLink : / api / v1 / namespaces / default / serviceacc  ounts /
default
  uid :  052fb0f4 - 3d50 - 11e5 - b066 - 42010af0d7  b6
secrets :
-  name :  default - token - uudge
imagePullS  ecrets :                  ## Add   this   item
-  name :  regsecret

```

Use the *`sa . yaml`* configuration file to replace the service account configurations of default.

```
kubectl   replace   serviceacc  ount   default  - f  ./ sa . yaml
```

```
serviceacc  ount  " default "  replaced
```

Execute the `kubectl  create  - f` command to create a tomcat orchestration as an example.

```
apiVersion :  apps / v1beta2  #  for   versions   before   1 . 8 . 0
  use   apps / v1beta1
kind :  Deployment
metadata :
  name :  tomcat - deployment
  labels :
    app :  tomcat
spec :
  replicas :  1
  selector :
    matchLabel  s :
      app :  tomcat
  template :
    metadata :
      labels :
        app :  tomcat
    spec :
      containers :
     -  name :  tomcat
        image :  registry - internal . cn - hangzhou . aliyuncs . com
/ abc / test : 1 . 0               # Replace   this   with   your
own   private   image   address
        ports :
       -  containerP  ort :  8080contai  nerPort :  8080
```

If you have configured properly, the pod starts successfully. Execute the `kubectl  get   pod   tomcat - xxx  - o   yaml` command. The following configuration items are displayed:

```
spec :
  imagePullS  ecrets :
 -  nameregsec  retey
```

7 Volume FAQ

What do I do if a volume cannot be mounted to the Kubernetes cluster?

Check whether the flexvolume plugin is installed

Run the following command on the Master node:

```
#  kubectl   get   pod   - n   kube - system  |  grep   flexvolume

 flexvolume - 4wh8s              1 / 1        Running    0
   8d
 flexvolume - 65z49              1 / 1        Running    0
   8d
 flexvolume - bpc6s              1 / 1        Running    0
   8d
 flexvolume - l8pml              1 / 1        Running    0
   8d
 flexvolume - mzkpv              1 / 1        Running    0
   8d
 flexvolume - wbfhv              1 / 1        Running    0
   8d
 flexvolume - xf5cs              1 / 1        Running    0
   8d
```

If the flexvolume plugin is installed, check whether flexvolume pods is running and that the number of running pods is the same as the number of nodes.

If no flexvolume plugin is installed, see Install the plug-in.

If the flexvolume pod is not running, see the running logs of the plugin to further analyze the cause.

Check whether the dynamic storage plugin is installed

To use the dynamic storage function of a cloud disk, you must install the dynamic storage plugin. Run the following command to check whether the dynamic storage plugin is installed:

```
#  kubectl   get   pod   - n   kube - system  |  grep   alicloud -
 disk

 alicloud - disk - controller - 8679c9fc76 - lq6zb      1 / 1
 Running     0     7d
```

If no dynamic storage plugin is installed, see Install the plug-in.

If the pod is not running, see the running logs of the plugin to further analyze the cause.

How do I view storage logs?

View flexvolume logs by running commands on Master1 node

Run the get command to view the error pod as follows:

```
#  kubectl   get   pod  - n   kube - system  |  grep   flexvolume
```

Run the log command to view the logs of the error pod as follows:

```
#  kubectl   logs   flexvolume - 4wh8s  - n   kube - system
#  kubectl   describe   pod   flexvolume - 4wh8s  - n   kube - system

# The   last   several   lines   of   the   pod   descriptio  n
 describe   the   status   of   the   pod .  You   can   analyze
 errors   according   to   the   status   of   the   pod .
```

View the driver logs of a cloud disk, NAS, and OSS:

```
#  View   the   persistent   logs   on   the   host   node .
#  If   a   volume   cannot   be   mounted   to   a   pod ,  run   the
   following   command   to   view   the   address   of   the   node
 where   the   pod   resides :

#  kubectl   describe   pod   nginx - 97dc96f7b - xbx8t  |  grep
 Node
 Node :  cn - hangzhou . i - bp19myla3u  vnt6zihejb / 192 . 168 . 247
 . 85
 Node - Selectors :  < none >

#  Log   on   to   the   node   to   view   logs .

#  ssh   192 . 168 . 247 . 85
#  ls  / var / log / alicloud / flexvolume *
 flexvolume  _disk . log    flexvolume  _nas . log    flexvolume  _o #
 ss . log

 The   logs   mounted   to   the   cloud   disk ,  NAS ,  and   OSS
 are   displayed .
```

Run the following commands on Master1 node to view provsioner plugin logs

Run the get command to view the error pod as follows:

```
#  kubectl   get   pod  - n   kube - system  |  grep   alicloud - disk
```

Run the log command to view the logs of the error pod as follows:

```
#  kubectl   logs   alicloud - disk - controller - 8679c9fc76 - lq6zb
 - n   kube - system
#  kubectl   describe   pod   alicloud - disk - controller -
 8679c9fc76 - lq6zb  - n   kube - system
```

```
# The last several lines of the pod descriptio n
 describe the status of the pod . You can analyze
 errors according to the status of the pod .
```

View kubelet logs

```
# If a volume cannot be mounted to a pod , run the
   following command to view the address of the node
 where the pod resides :

# kubectl describe pod nginx - 97dc96f7b - xbx8t | grep
 Node
 Node : cn - hangzhou . i - bp19myla3u vnt6zihejb / 192 . 168 . 247
 . 85
 Node - Selectors : < none >

# Log on to the node to view kubelet logs :

# ssh 192 . 168 . 247 . 85
# journalctl - u kubelet - r - n 1000 &> kubelet . log

# The value of - n indicates the number of log lines
   that you expect to display .
```

The preceding content describes the methods to obtain the logs of errors that occurred to flexvolume, provsioner, and kubelet. If you cannot fix the errors according to the logs, we recommend that you submit a ticket with the log informatio n for further consultation.

Cloud disk volume FAQ

What do I do if a cloud disk fails to be mounted to the cluster and a timeout error is displayed?

If the cluster node is added manually, this problem may be caused by insufficient STS permissions. We recommend that you manually configure RAM permissions. For more information, see Use the instance RAM role in the console.

What do I do if a cloud disk fails to be mounted to the cluster and a Size error is displayed?

Depending on the type of cloud disk you create, you must confirm that the following requirements are met:

Note:

- The minimum capacity of a basic cloud disk is 5 GiB.
- The minimum capacity of an Ultra disk is 20 GiB.
- The minimum capacity of an SSD disk is 20 GiB.

What do I do if a cloud disk fails to be mounted to the cluster and a zone error is displayed?

To mount a cloud disk to your Kubernetes cluster, you must select the cloud disk that is in the same region and zone as the ECS instances used by the Kubernetes cluster.

What do I do if the input/output error is displayed by the cloud disk after the system is upgraded?

1. Upgrade the flexvolume to v1.9.7-42e8198 or later.
2. Recreate the faulty pods.

Run the following command to upgrade the flexvolume:

```
#  kubectl   set   image   daemonset / flexvolume   acs - flexvolume =
 registry . cn - hangzhou . aliyuncs . com / acs / flexvolume : v1 . 9
 . 7 - 42e8198  - n   kube - system
```

To obtain the latest flexvolume version, log on to the Container Registry console, click Search in the left-side navigation pane, and search for acs/flexvolume.

NAS volume FAQ

What do I do if it takes a long time to mount a NAS file system to a Kubernetes cluster?

This problem may occur if the NAS file system contains a large amount of data and you set the chmod parameter in the mount template. We recommend that you remove the chmod parameter setting.

What do I do if a NAS file system fails to be mounted to a Kubernetes cluster and a timeout error is displayed?

Check whether the NAS mount point and the cluster are in the same VPC. Otherwise, the NAS file system cannot be mounted to the cluster.

OSS volume FAQ

What do I do if an OSS bucket fails to be mounted to a Kubernetes cluster?

- Check whether the submitted Access Key is valid.
- Check whether the URL used to mount the OSS bucket is accessible through network.

What do I do if the OSS mount directory within the container becomes unavailable after the cluster is upgraded?

If you upgrade your Kubernetes cluster or restart a kubelet, the ossfs process restarts because the container network restarts.

After the ossfs process restarts, the mapping to the container directory from your host becomes invalid. In this case, you need to restart the container or recreate the pod.

We recommend that you configure the liveness probe for heath checks so that the container or pod can automatically restart when such a problem occurs. For more information, see Use an Alibaba Cloud OSS bucket as a persistent volume.

8 Do I select the Terway or Flannel plugin for my Kubernetes cluster network?

This topic describes Terway and Flannel, two network plugins provided by Container Service for creating a Kubernetes cluster. This information helps you select an appropriate network plugin when creating a cluster.

When you create a Kubernetes cluster, Alibaba Cloud Container Service provides two network plugins: Terway and Flannel.

- Flannel: a simple and stable community Flannel CNI plugin. Flannel works with the high-speed network of Alibaba Cloud VPC, providing a high-performance and stable container network for clusters. However, it provides only a few simple features. For example, it does not support the Kubernetes Network Policy.
- Terway: a network plugin developed by Alibaba Cloud Container service. It is fully compatible with Flannel, and can allocate Alibaba Cloud Elastic Network Interfaces (ENIs) to containers. It can also define the access policies between containers according to the Kubernetes Network Policy. In addition, it supports bandwidth limiting for individual containers. If you do not need to use a Network Policy, we recommend that you select Flannel. In other cases, we recommend that you select Terway. For more information, see Terway network plugin.

9 How to manually install alicloud-application-controller

By default, alicloud-application-controller is installed in Alibaba Cloud Container Service in version 1.10.4 and later to provide the release based on custom resource definition (CRD).

> **Note:**
> In the Kubernetes cluster of the latest version, alicloud-application-controller is installed by default. In Kubernetes clusters of old versions, manually install alicloud-application-controller and the oldest version of Kubernetes cluster must be 1.9.3.

Use the `kubectl  create  - f  alicloud - applicatio  n - controller . yml` command to deploy alicloud-application-controller. In *`alicloud - applicatio  n - controller . yml`* , enter the following orchestration template:

```
apiVersion :  extensions / v1beta1
kind :  Deployment
metadata :
  name :  alicloud - applicatio  n - controller
  labels :
    owner :  aliyun
    app :  alicloud - applicatio  n - controller
  namespace :  kube - system
spec :
  replicas :  1
  selector :
    matchLabel  s :
      owner :  aliyun
      app :  alicloud - applicatio  n - controller
  template :
    metadata :
      labels :
        owner :  aliyun
        app :  alicloud - applicatio  n - controller
      annotation  s :
        scheduler . alpha . kubernetes . io / critical - pod : ''
    spec :
      toleration  s :
     -  effect :  NoSchedule
        operator :  Exists
        key :  node - role . kubernetes . io / master
     -  effect :  NoSchedule
        operator :  Exists
        key :  node . cloudprovi  der . kubernetes . io / uninitiali
zed
      containers :
       -  name :  alicloud - applicatio  n - controller
          image :  registry . cn - hangzhou . aliyuncs . com / acs /
aliyun - app - lifecycle - manager : 0 . 1 - c8d5da8
```

```
          imagePullP  olicy :  IfNotPrese  nt
      serviceAcc  ount :  admin
```

10 Kubernetes cluster network failures caused by security group settings

This topic describes the Kubernetes cluster network failures caused by cluster security group settings, and provides corresponding resolutions.

Symptom

Containers cannot communicate with each other over the Kubernetes cluster network .

Causes

- A relevant ingress security group rule is removed. The following are the details of the rule: the ingress Authorization Objects is Pod Network CIDR and the Protocol Type is All.
- Newly added ECS instances and the Kubernetes cluster are located in different security groups.

Resolution

Cause 1: A relevant ingress security group rule is removed. The following are details of the rule: the ingress Authorization Objects is Pod Network CIDR and the Protocol Type is All.

1. Log on to the Container Service console.
2. In the left-side navigation pane under Kubernetes, click Clusters.
3. Click the target cluster name to view cluster details.
4. In the Cluster Resource area, click VPC.

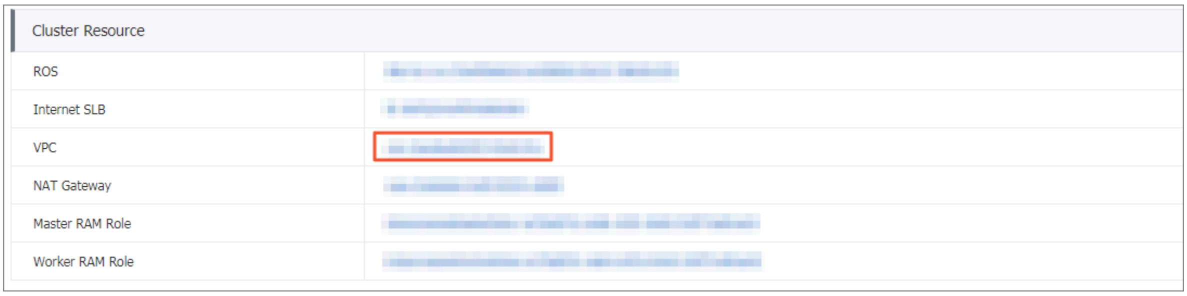

Cluster Resource

ROS

Internet SLB

VPC

NAT Gateway

Master RAM Role

Worker RAM Role

5. In the Network Resources area, click the number on the right of Security Group.

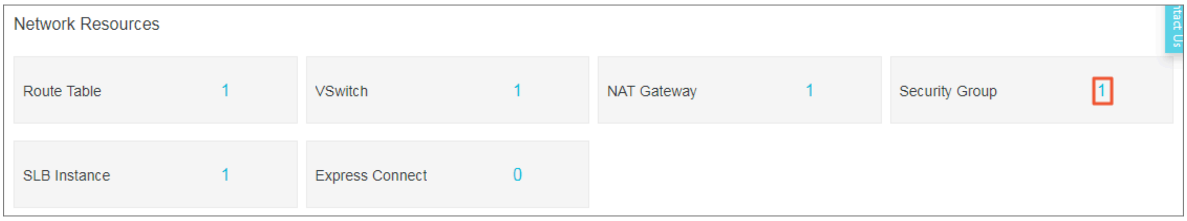

Network Resources

Route Table	1	VSwitch	1	NAT Gateway	1	Security Group	1
SLB Instance	1	Express Connect	0				

6. In the Actions column of the target security group, click Add Rules.
7. On the Ingress page, click Add Security Group Rule in the upper-right corner.

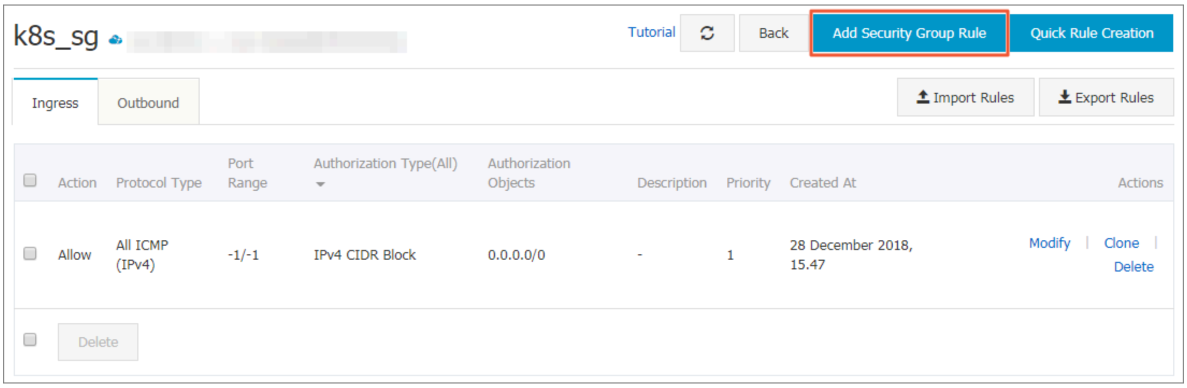

□	Action	Protocol Type	Port Range	Authorization Type(All)	Authorization Objects	Description	Priority	Created At	Actions
□	Allow	All ICMP (IPv4)	-1/-1	IPv4 CIDR Block	0.0.0.0/0	-	1	28 December 2018, 15.47	Modify | Clone | Delete

8. Set Protocol Type and Authorization Objects.

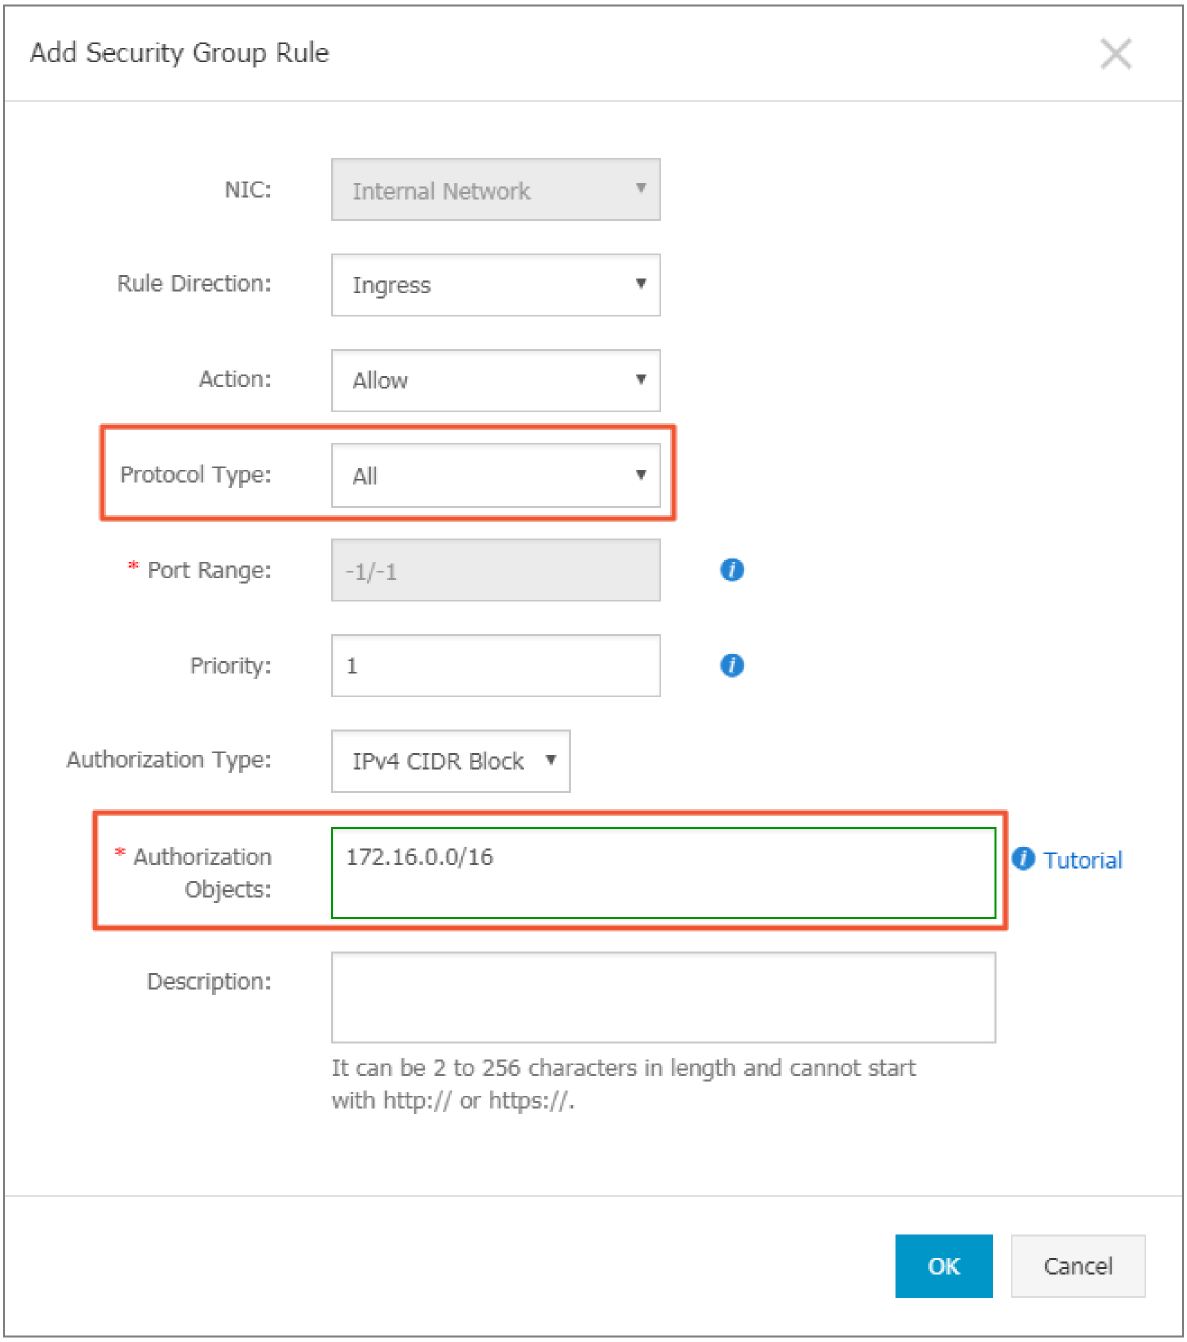

Add Security Group Rule

NIC: Internal Network

Rule Direction: Ingress

Action: Allow

Protocol Type: All

* Port Range: -1/-1

Priority: 1

Authorization Type: IPv4 CIDR Block

* Authorization Objects: 172.16.0.0/16

Description:

It can be 2 to 256 characters in length and cannot start with http:// or https://.

Select All from the Protocol Type drop-down list.

Enter the Pod Network CIDR of the Kubernetes cluster for Authorization Objects.

Note:

- You can view the Pod Network CIDR of the Kubernetes cluster in the Cluster Information area on the cluster basic information page.

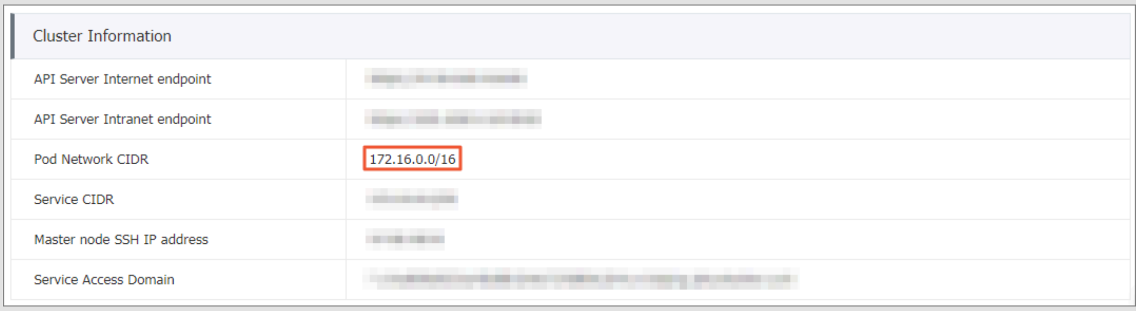

Cluster Information

API Server Internet endpoint	[redacted]
API Server Intranet endpoint	[redacted]
Pod Network CIDR	172.16.0.0/16
Service CIDR	[redacted]
Master node SSH IP address	[redacted]
Service Access Domain	[redacted]

- For more information about theAuthorization Objects settings, see Scenarios.

Verify the results

The required ingress security group rule is added. The following are details of the rule: the ingress Authorization Objects is Pod Network CIDR and the Protocol Type is All.

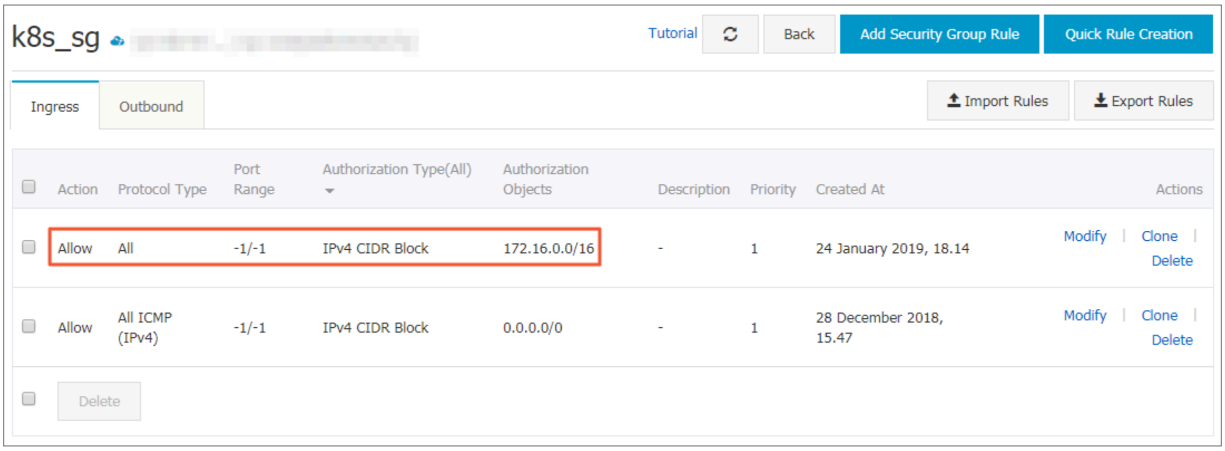

k8s_sg [redacted]

□	Action	Protocol Type	Port Range	Authorization Type(All)	Authorization Objects	Description	Priority	Created At	Actions
□	Allow	All	-1/-1	IPv4 CIDR Block	172.16.0.0/16	-	1	24 January 2019, 18.14	Modify | Clone | Delete
□	Allow	All ICMP (IPv4)	-1/-1	IPv4 CIDR Block	0.0.0.0/0	-	1	28 December 2018, 15.47	Modify | Clone | Delete
□	Delete								

Cause 2: Newly added ECS instances and the Kubernetes cluster are located in different security groups.

1. Log on to the Container Service console.
2. In the left-side navigation pane, click Clusters.
3. Click the target Cluster Name.
4. In the Cluster Resource area, click VPC.

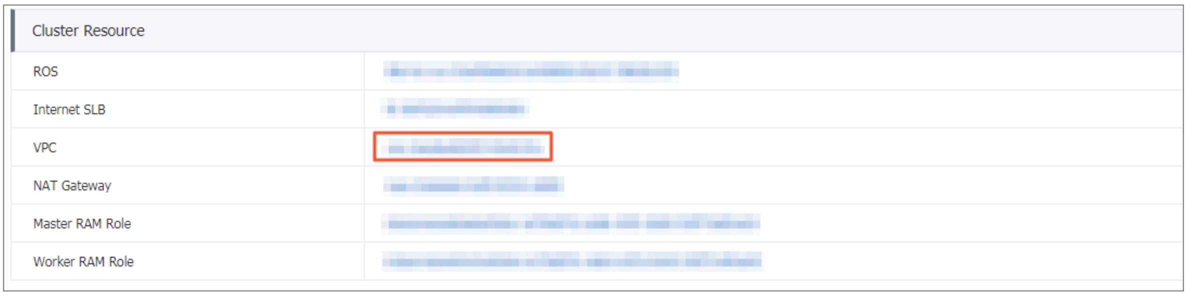

Cluster Resource

ROS	[redacted]
Internet SLB	[redacted]
VPC	[redacted]
NAT Gateway	[redacted]
Master RAM Role	[redacted]
Worker RAM Role	[redacted]

5. On the VPC Details page, click the number on the right of Security Group in the Network Resources area.

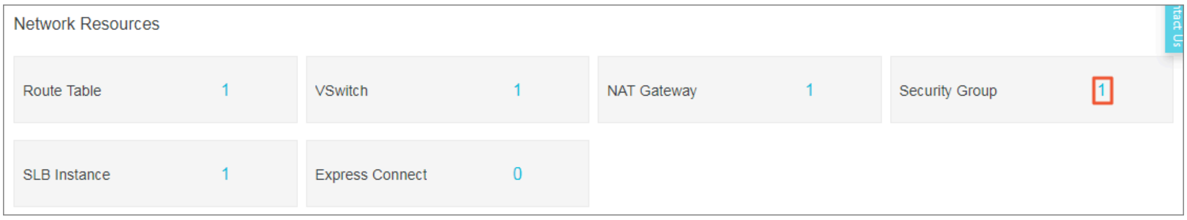

Network Resources

Route Table	1	VSwitch	1	NAT Gateway	1	Security Group	1
SLB Instance	1	Express Connect	0				

6. On the Security Groups page, view the target security group name.

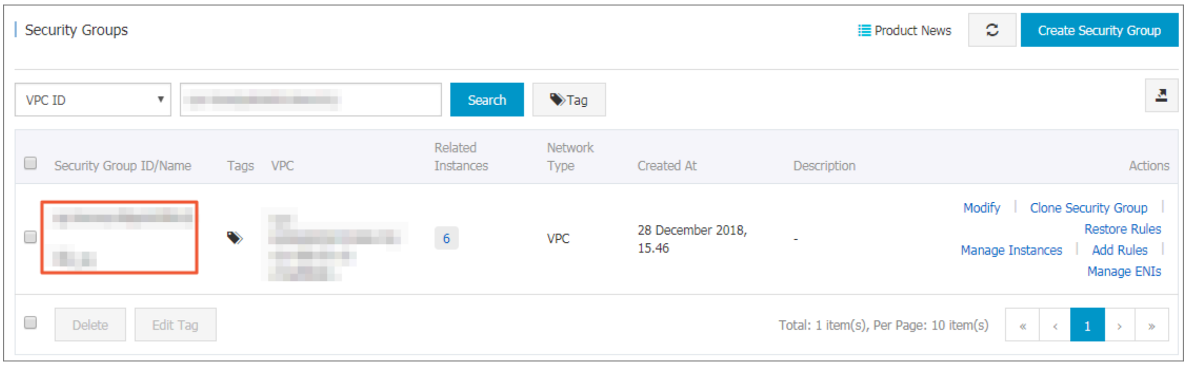

Security Groups

□	Security Group ID/Name	Tags	VPC	Related Instances	Network Type	Created At	Description	Actions
□	[redacted] [redacted]		[redacted]	6	VPC	28 December 2018, 15.46	-	Modify | Clone Security Group | Restore Rules Manage Instances | Add Rules | Manage ENIs

□ Delete Edit Tag

Total: 1 item(s), Per Page: 10 item(s) « ‹ 1 › »

7. In the Elastic Compute Service console, click Instances in the left-side navigation pane.

8. On the Instances page, choose More > Network and Security Group > Add to Security Group in the Actions column of the target instance.

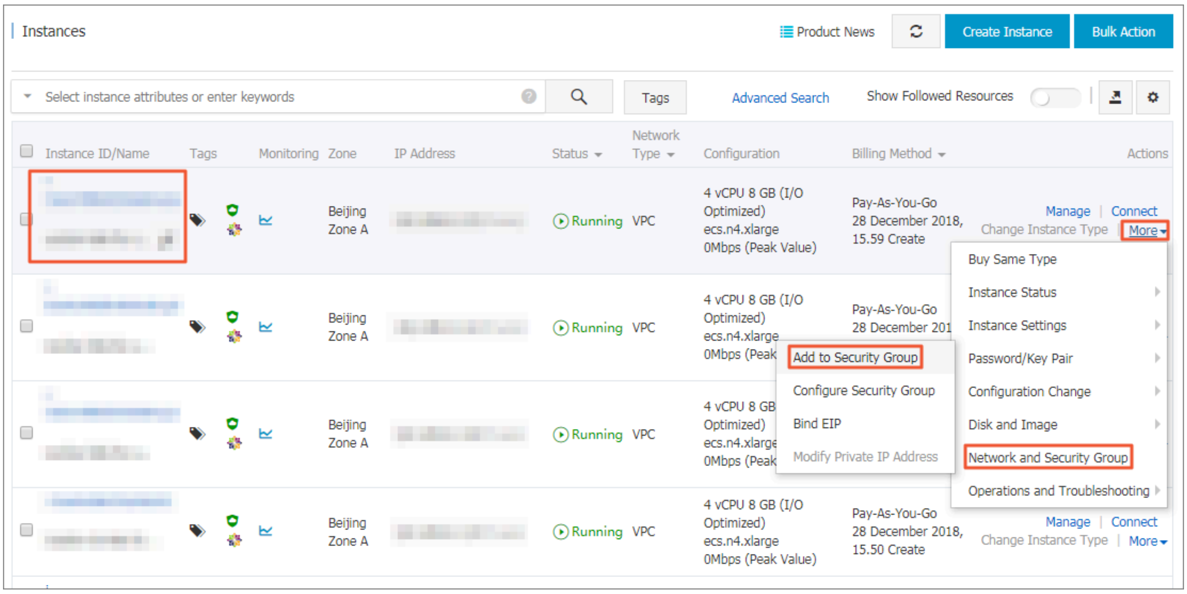

Instances

□	Instance ID/Name	Tags	Monitoring	Zone	IP Address	Status ▾	Network Type ▾	Configuration	Billing Method ▾	Actions
□	[redacted] [redacted]			Beijing Zone A	[redacted]	Running	VPC	4 vCPU 8 GB (I/O Optimized) ecs.n4.xlarge 0Mbps (Peak Value)	Pay-As-You-Go 28 December 2018, 15.59 Create	Manage | Connect Change Instance Type | More ▾
□	[redacted] [redacted]			Beijing Zone A	[redacted]	Running	VPC	4 vCPU 8 GB (I/O Optimized) ecs.n4.xlarge 0Mbps (Peak	Pay-As-You-Go 28 December 201	
□	[redacted] [redacted]			Beijing Zone A	[redacted]	Running	VPC	4 vCPU 8 GB Optimized) ecs.n4.xlarge 0Mbps (Peak		
□	[redacted] [redacted]			Beijing Zone A	[redacted]	Running	VPC	4 vCPU 8 GB (I/O Optimized) ecs.n4.xlarge 0Mbps (Peak Value)	Pay-As-You-Go 28 December 2018, 15.50 Create	Manage | Connect Change Instance Type | More ▾

9. Click the drop-down arrow on the right of the Security Group box, and enter the cluster security group name obtained in step 6.

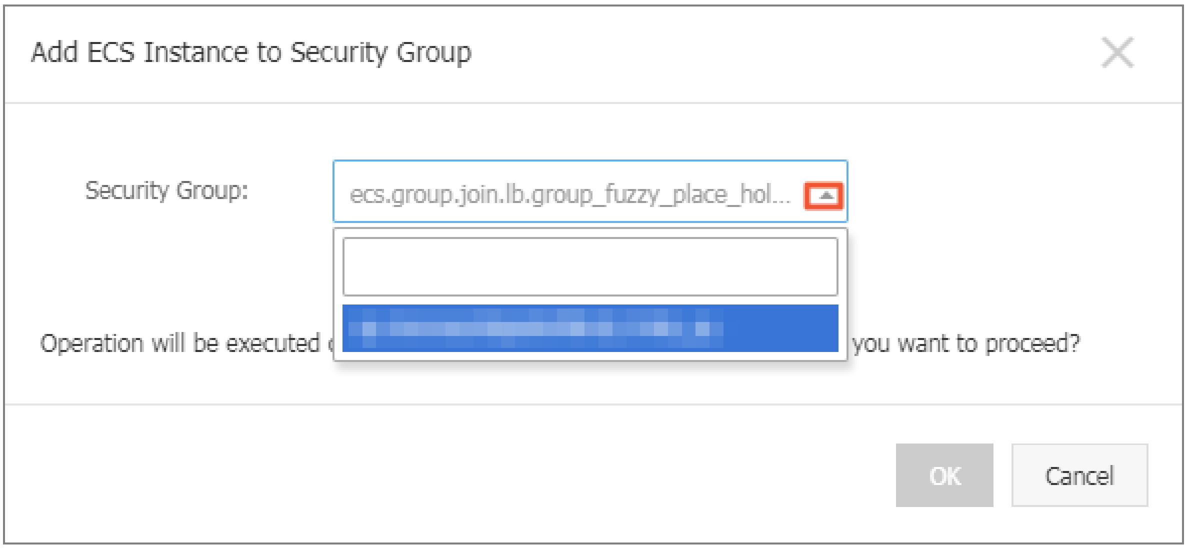

Add ECS Instance to Security Group

Security Group: ecs.group.join.lb.group_fuzzy_place_hol...

[redacted]

Operation will be executed ... you want to proceed?

10.Click OK.

Verify the results

1. In the left-side navigation pane of the Elastic Computer Service console, click Instances and then click the target instance.
2. In the left-side navigation pane, click Security Groups.
3. The Security Groups area shows that the target ECS instance has been added to the security group to which the Kubernetes cluster belongs.

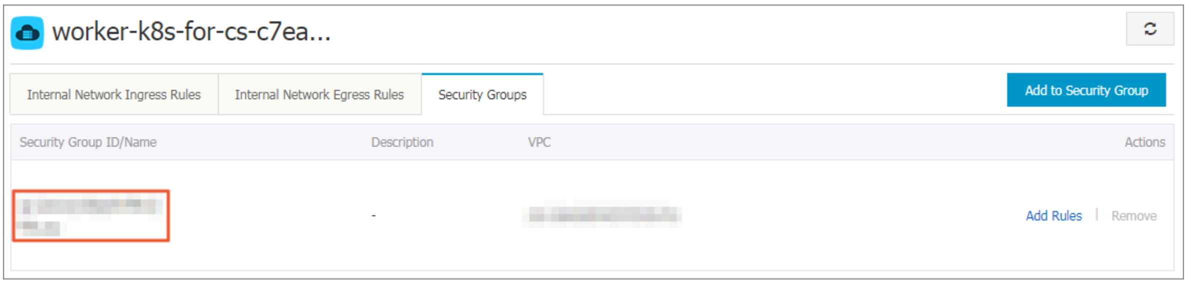

worker-k8s-for-cs-c7ea...

Security Group ID/Name	Description	VPC	Actions
[redacted]	-	[redacted]	Add Rules | Remove

11 Istio FAQ

This topic lists common Istio FAQ and their corresponding solutions.

What do I do if the services in the cluster cannot access external URLs?

Symptom

The services in the cluster cannot access external URLs.

Cause

By default, this is because the pod in the Istio service mesh uses iptables to transparently forward all outbound traffic to the sidecar. The sidecar can only handle the traffic destined for addresses within the cluster.

Solutions

- Define `ServiceEntry` to call external services.
- Configure Istio to allow access to a specific range of IP addresses.

For more information, see Control Egress Traffic.

What do I do if Tiller is in an earlier version?

Symptom

The following message is displayed during the installation process:

```
Can't install release with errors: rpc error: code = Unknown desc = Chart incompatible with Tiller v2.7.0
```

Cause

Your current version of Tiller needs to be upgraded.

Solution

Run one of the following two commands to upgrade the Tiller version.

Note:

You must upgrade Tiller to V2.10.0 or later.

Run the following command to upgrade Tiller to V2.11.0:

```
helm  init  -- tiller - image  registry . cn - hangzhou . aliyuncs
. com / acs / tiller : v2 . 11 . 0  -- upgrade
```

Run the following command to upgrade Tiller to V2.10.0:

```
helm  init  -- tiller - image  registry . cn - hangzhou . aliyuncs
. com / acs / tiller : v2 . 10 . 0  -- upgrade
```

Note:

After you upgrade Tiller to a new version, we recommend that you upgrade the client to the corresponding version. For more information, see Install a client.

What do I do if Custom Resource Definitions (CRDs) are in an invalid version?

Symptom

The following message is displayed when you create Istio for the first time or upgrade from Istio 1.0:

```
Can ' t  install  release  with  errors : rpc  error : code
 =  Unknown  desc  = apiVersion  " networking . istio . io /
v1alpha3 "  in  ack - istio / charts / pilot / templates / gateway .
yaml  is  not  available
```

Cause

CRDs do not exist or are of an earlier version.

Note:

This problem occurs only in Helm 2.10.0 and earlier versions. For Helm later than V2.10.0, the system automatically upgrades CRDs.

Solution

1. Download the latest Istio. For more information, see Download the release.

2. Run the following command to install the latest CRDs:

```
$ kubectl apply - f install / kubernetes / helm / istio /
 templates / crds . yaml - n istio - system
```

3. If you have enabled `certmanage r`, you must run the following command to install the relevant CRDs:

```
$ kubectl apply - f install / kubernetes / helm / istio /
 charts / certmanage r / templates / crds . yaml
```

What do I do if Istio cannot be installed when I log on as a RAM user?

Symptom

The following message or a similar one is displayed during the installation process:

```
Error from server ( Forbidden ): error when retrieving
current configurat ion of :
Resource : " apiextensi ons . k8s . io / v1beta1 , Resource =
customreso urcedefini tions ", GroupVersi onKind : " apiextensi
ons . k8s . io / v1beta1 , Kind = CustomReso urceDefini tion "
```

Cause

The RAM user does not have permission to install Istio.

Solutions

- Log on to Alibaba Cloud by using the primary account.
- Grant the RAM user the required permissions. For example, you can grant the RAM user the `cluster - admin` custom role. For more information, see Grant RBAC permissions to a RAM user.

What do I do if CRDs are not removed after Istio is uninstalled?

Symptom

CRDs are not removed after Istio is uninstalled.

Cause

The system does not remove CRDs when you uninstall Istio.

Solution

1. If you use Helm later than V2.9.0, perform step 2 directly. If you use Helm 2.9.0 or earlier, you must first run the following command to delete Job resources:

```
$ kubectl -n istio-system delete job --all
```

2. Run the following command to delete CRDs:

```
$ kubectl delete -f install/kubernetes/helm/istio/
 templates/crds.yaml -n istio-system
```

What do I do if a custom resource is retained after I delete Istio?

Symptom

After Istio is deleted from a Kubernetes cluster, a custom resource is retained.

Cause

You only deleted the CRD.

Solution

1. Run the `kubectl edit istio -n istio-system istio-config` command.

2. Delete `- istio-operator.finializer.alibabacloud.com` in the `finalizers` parameter.

12 Can I manually set a security group for a Kubernetes cluster?

No, an existing security group cannot be manually set for a Kubernetes cluster. A security group can only be set automatically for a Kubernetes cluster when the cluster is created.

13 How do I customize a RAM role for a Kubernetes cluster?

You cannot create a RAM role manually. However, when cluster Worker nodes are created, a Worker RAM role is automatically created for the Kubernetes cluster. You can then add policies to the Worker RAM role to grant the role the required permissions.

1. Log on to the Container Service console.
2. In the left-side navigation pane under Container Service-Kubernetes, choose Clusters > Clusters.
3. Click the target cluster name.

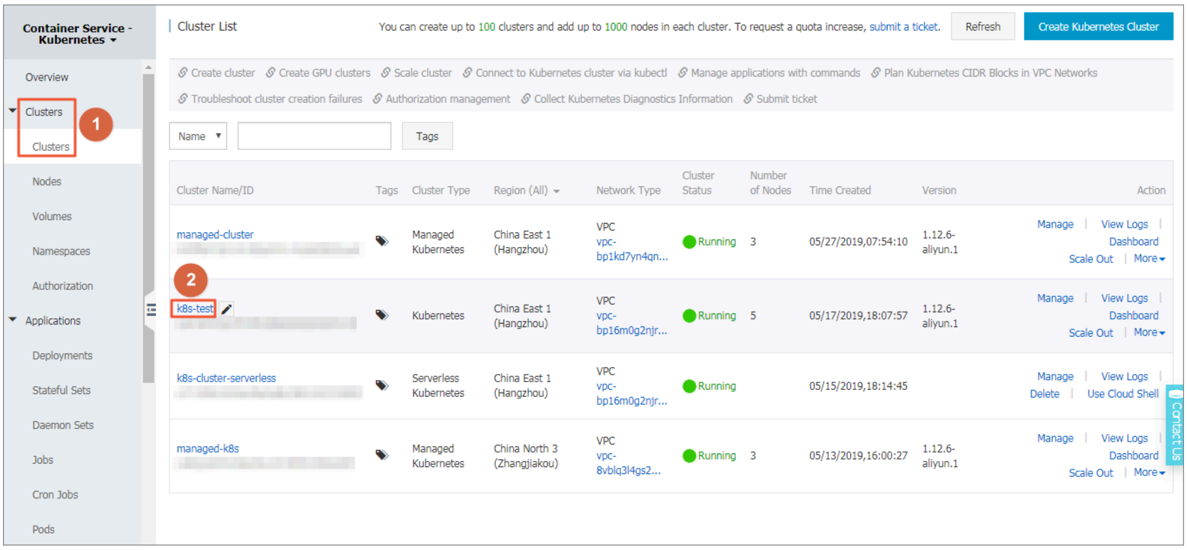
Container Service - Kubernetes
Overview
Clusters
Clusters
1
Nodes
Volumes
Namespaces
Authorization
Applications
Deployments
Stateful Sets
Daemon Sets
Jobs
Cron Jobs
Pods
Cluster List
Refresh
Create Kubernetes Cluster
Create cluster
Create GPU clusters
Scale cluster
Connect to Kubernetes cluster via kubectl
Manage applications with commands
Plan Kubernetes CIDR Blocks in VPC Networks
Troubleshoot cluster creation failures
Authorization management
Collect Kubernetes Diagnostics Information
Submit ticket
Name
Tags
2
Contact Us

You can create up to 100 clusters and add up to 1000 nodes in each cluster. To request a quota increase, submit a ticket.

Cluster Name/ID	Tags	Cluster Type	Region (All)	Network Type	Cluster Status	Number of Nodes	Time Created	Version	Action
managed-cluster [redacted]		Managed Kubernetes	China East 1 (Hangzhou)	VPC vpc-bp1kd7yn4qn...	Running	3	05/27/2019,07:54:10	1.12.6-aliyun.1	Manage | View Logs | Dashboard Scale Out | More
k8s-test [redacted]		Kubernetes	China East 1 (Hangzhou)	VPC vpc-bp16m0g2njr...	Running	5	05/17/2019,18:07:57	1.12.6-aliyun.1	Manage | View Logs | Dashboard Scale Out | More
k8s-cluster-serverless [redacted]		Serverless Kubernetes	China East 1 (Hangzhou)	VPC vpc-bp16m0g2njr...	Running		05/15/2019,18:14:45		Manage | View Logs | Delete | Use Cloud Shell
managed-k8s [redacted]		Managed Kubernetes	China North 3 (Zhangjiakou)	VPC vpc-8vblq3l4gs2...	Running	3	05/13/2019,16:00:27	1.12.6-aliyun.1	Manage | View Logs | Dashboard Scale Out | More

4. In the Cluster Resources area, click Worker RAM Role.

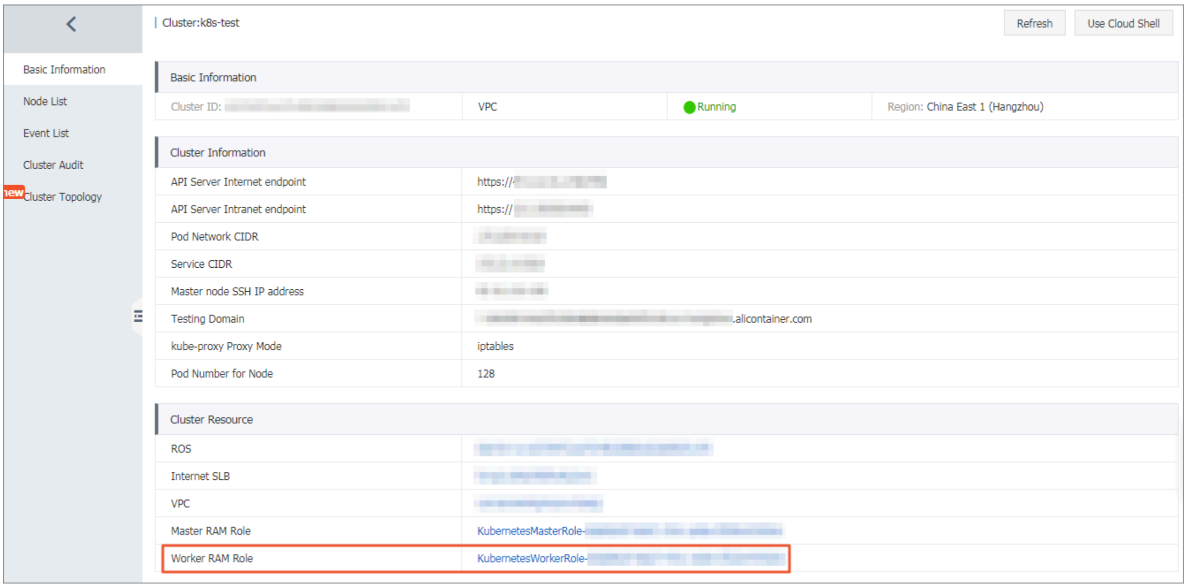

Basic Information

Cluster ID: [redacted]	VPC	Running	Region: China East 1 (Hangzhou)

Cluster Information

API Server Internet endpoint	https://[redacted]
API Server Intranet endpoint	https://[redacted]
Pod Network CIDR	[redacted]
Service CIDR	[redacted]
Master node SSH IP address	[redacted]
Testing Domain	[redacted].alicontainer.com
kube-proxy Proxy Mode	iptables
Pod Number for Node	128

Cluster Resource

ROS	[redacted]
Internet SLB	[redacted]
VPC	[redacted]
Master RAM Role	KubernetesMasterRole-[redacted]
Worker RAM Role	KubernetesWorkerRole-[redacted]

5. On the RAM Roles page, click the policy name on the Permission tab page to view the policy details.

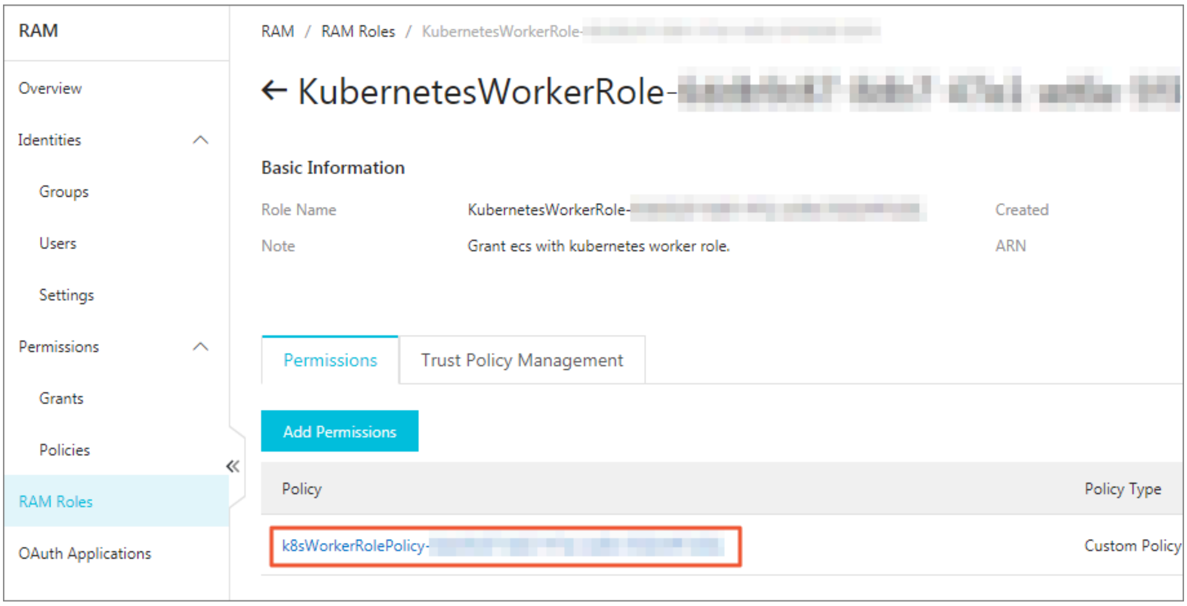

← KubernetesWorkerRole-[redacted]

Basic Information

Role Name	KubernetesWorkerRole-[redacted]	Created
Note	Grant ecs with kubernetes worker role.	ARN

Policy	Policy Type
k8sWorkerRolePolicy-[redacted]	Custom Policy

6. On the Policy Document tab page, click Modify Policy Document.

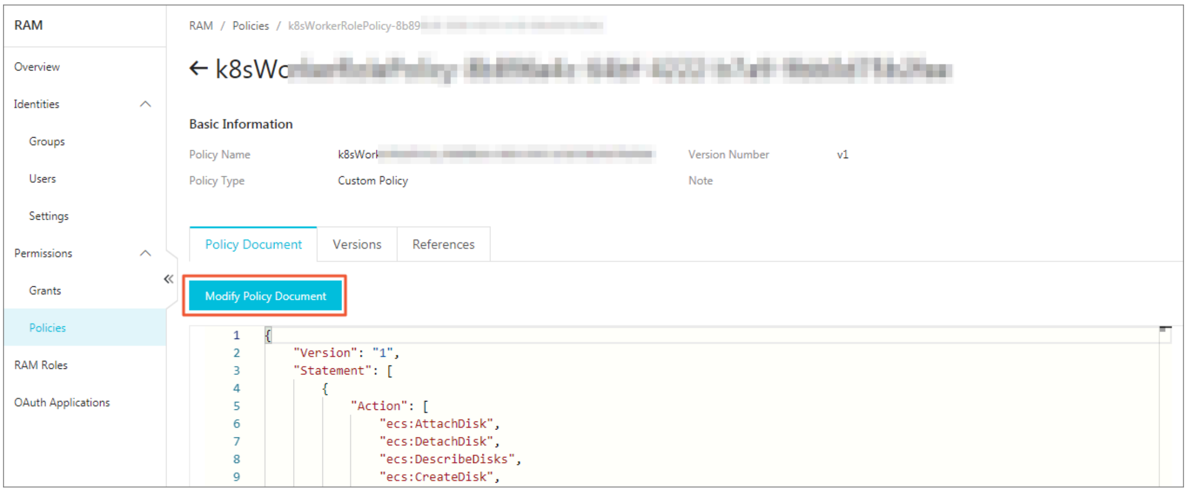

RAM / Policies / k8sWorkerRolePolicy-8b89[redacted]

← k8sWo[redacted]

Basic Information

Policy Name k8sWork[redacted]

Version Number v1

Policy Type Custom Policy

Note

```
{
    "Version": "1",
    "Statement": [
        {
            "Action": [
                "ecs:AttachDisk",
                "ecs:DetachDisk",
                "ecs:DescribeDisks",
                "ecs:CreateDisk",
```

7. On the displayed page, add the target policies to the Policy Document area, and then click OK.

   In this example, the policies containing the permissions of scaling and deleting clusters are added to the policy document. For more information about permissions supported by a Kubernetes cluster, see Table 1.

```
{
            " Action ": [
              " cs : ScaleClust  er ",
              " cs : DeleteClus  ter "
            ],
            " Resource ": "*",
            " Effect ": " Allow "
        }
```

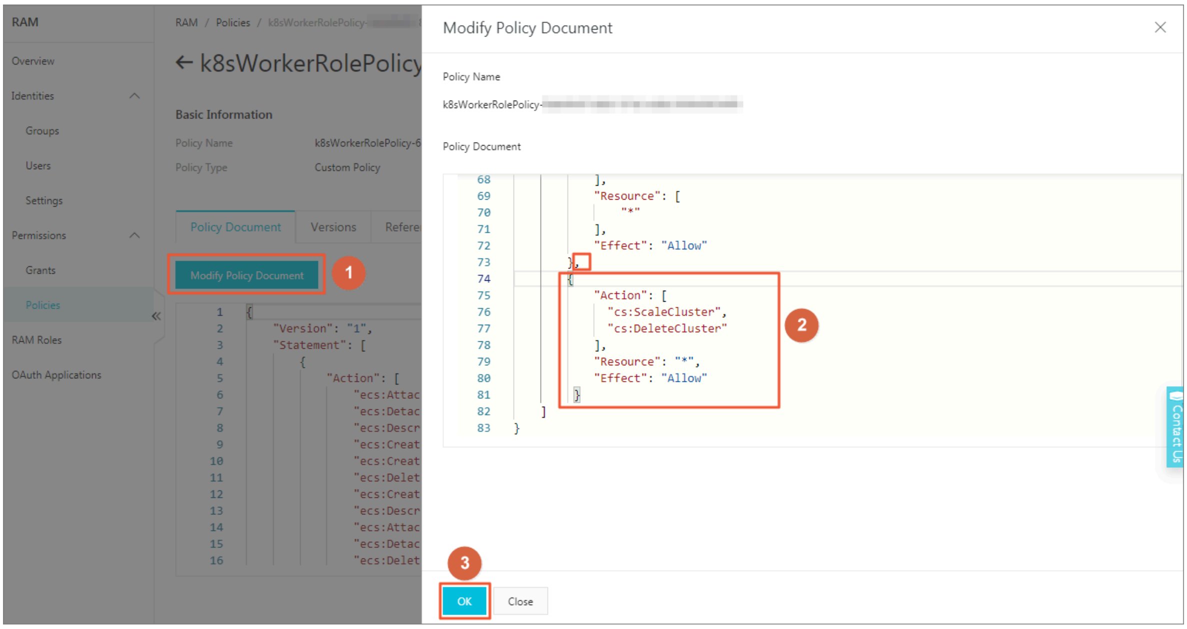

Modify Policy Document

Policy Name

k8sWorkerRolePolicy-[redacted]

Policy Document

```
                ],
                "Resource": [
                    "*"
                ],
                "Effect": "Allow"
            },
            {
                "Action": [
                  "cs:ScaleCluster",
                  "cs:DeleteCluster"
                ],
                "Resource": "*",
                "Effect": "Allow"
            }
        ]
    }
```

14 How do I rename an SLB instance that I created in Cloud Controller Manager (CCM) of V1.9.3.10 or an earlier version?

To rename the SLB instance, you must manually add a specific tag to the SLB instance before you can rename it. To do so complete the following steps:

Note:

You can directly rename an SLB instance that you created in CCM of a version later than V1.9.3.10. This is because the system automatically adds a tag to the SLB instance created in these versions of CCM.

1. Log on to the Master node of the target Kubernetes cluster. For more information, see Connect to a Kubernetes cluster by using kubectl.
2. Run the `# kubectl get svc -n ${ namespace } ${ service }` command to view the IP address of the service provided by the target SLB instance.

   Note:

   - You must replace ${namespace} and ${service} with the target namespace and service names.
   - The service type is displayed as `LoadBalancer`.

   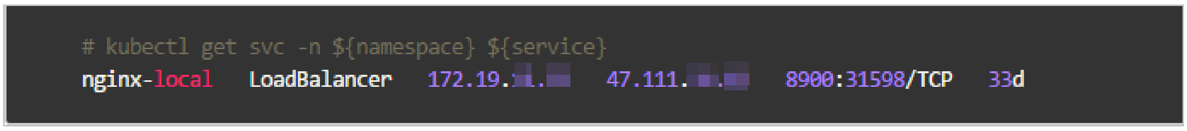

   ```
   # kubectl get svc -n ${namespace} ${service}
   nginx-local   LoadBalancer   172.19.[redacted].[redacted]   47.111.[redacted].[redacted]   8900:31598/TCP   33d
   ```

3. Run the following command to generate the tag required by the target SLB instance:

   ```
   # kubectl get svc -n ${ namespace } ${ service } -o jsonpath ="{. metadata . uid }"| awk -F "-" '{ print " kubernetes . do . not . delete : " substr (" a "$ 1 $ 2 $ 3 $ 4 $ 5 , 1 , 32 )}'
   ```

   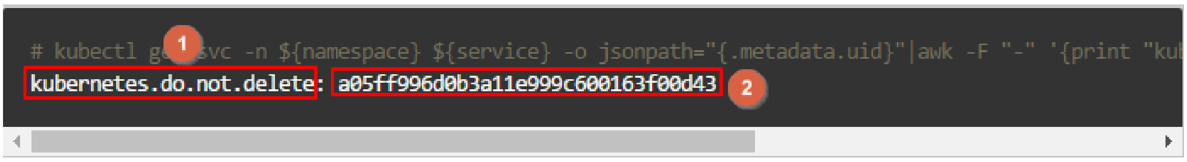

   ```
   # kubectl get svc -n ${namespace} ${service} -o jsonpath="{.metadata.uid}"|awk -F "-" '{print "ku
   kubernetes.do.not.delete: a05ff996d0b3a11e999c600163f00d43
   ```

4. Log on to the SLB consoleSLB console. Then, select the corresponding region to which the target SLB instance belongs, and use the IP addressed obtained in step 2 to search for the SLB instance.
5. Add a tag for the target SLB instance according to the obtained key and its value shown in the preceding figure (in step 3). For more information, see Manage tags.

15 Cluster certificate update FAQ

When do I have to update my cluster certificates?

You can update your cluster certificates two months before your cluster certificates expire.

How can I update cluster certificates?

You can click the red button in the console to update the certificates automatically, or update them by following these steps:

1. Update data according to the following table.

Node type	Backup data
Master	· /etc/kubernetes/ · /var/lib/kubelet/pki · /etc/systemd/system/kubelet.service.d/10-kubeadm.conf · /etc/kubeadm/ · *Necessary service data* Note: If the `/ var / lib / kubelet / pki` file contains no data, or if no *necessary service data* is available, do not back up them.
Worker	· /etc/kubernetes/ · /etc/systemd/system/kubelet.service.d/10-kubeadm.conf · /var/lib/kubelet/pki/* · *Necessary service data* Note: If the `/ var / lib / kubelet / pki /*` file contains no data, or if no *necessary service data* is available, do not back up them.

2. Update certificates according to the following tables.

Table 15-1: Master node

Certificate or conf	Path
· apiserver.crt · apiserver.key	/etc/kubernetes/pki
· apiserver-kubelet-client.crt · apiserver-kubelet-client.key	/etc/kubernetes/pki
· front-proxy-client.crt · front-proxy-client.key	/etc/kubernetes/pki
· dashboard.crt · dashboard.key	/etc/kubernetes/pki/dashboard
· kubelet.crt · kubelet.key Note: If the `kubelet . key` file contains no data, do not update it.	/var/lib/kubelet/pki Note: If the file contains no data, do not update it.
admin.conf	/etc/kubernetes
kube.conf	/etc/kubernetes
controller-manager.conf	/etc/kubernetes
scheduler.conf	/etc/kubernetes
kubelet.conf	/etc/kubernetes
config	~/.kube/

Certificate or conf	Path
· kubelet-client-current.pem or kubelet-client.crt\\ · kubelet-client.key Note: If the `kubelet-client.key` file contains no data, do not update it.	/var/lib/kubelet/pki Note: If the file contains no data, do not update it.

Table 15-2: Worker node

Certificate or conf	Path
· kubelet.crt · kubelet.key Note: If the `kubelet.key` file contains no data, do not update it.	/var/lib/kubelet/pki Note: If the file contains no data, do not update the file.
· kubelet-client-current.pem or kubelet-client.crt · kubelet-client.key Note: If the `kubelet-client.key` file contains no data, do not update it.	/var/lib/kubelet/pki Note: If the file contains no data, do not update the file.
kubelet.conf	/etc/kubernetes

How long does it take to update a cluster certificate?

Each update takes 5 to 10 minutes.

Note:

- When the certificates are being updated, the native components of Kubernetes will be restarted for a while. Therefore, we recommend that you perform the update during off-peak hours.
- The update does not affect existing online services.