

# Alibaba Cloud Alibaba Cloud Container Service for Kubernetes

FAQ

Issue: 20190920

## 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.	 <b>Danger:</b> 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.	 <b>Warning:</b> 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.	 <b>Notice:</b> 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.	 <b>Note:</b> You can use Ctrl + A to select all files.
>	Multi-level menu cascade.	Settings > Network > Set network type
<b>Bold</b>	It is used for buttons, menus, page names, and other UI elements.	Click <b>OK</b> .
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 <i>Instance_ID</i></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 General FAQ.....	1
2 Failed to create a Kubernetes cluster.....	4
3 Failed to delete Kubernetes clusters: ROS stack cannot be deleted.....	7
4 How do I collect Kubernetes diagnosis information when a Kubernetes cluster exception or a cluster node exception occurs?.....	10
5 Upgrade Helm manually.....	11
6 How to use private images in Kubernetes clusters.....	12
7 Do I select the Terway or Flannel plugin for my Kubernetes cluster network?.....	15
8 How to manually install alicloud-application-controller.....	16
9 Can I manually set a security group for a Kubernetes cluster?.....	18
10 How do I customize a RAM role for a Kubernetes cluster?...	19
11 How do I use a RAM user account to grant RBAC permissions to any other RAM user?.....	22
12 How do I rename an SLB instance that I created in Cloud Controller Manager (CCM) of V1.9.3.10 or an earlier version?.....	24
13 Failed to start a container on the GPU node.....	26
14 How do I combine kubeconfig files of multiple clusters?....	28

# 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	<ul style="list-style-type: none"> <li>- High-performance Virtual Private Cloud (VPC) network plug-in.</li> <li>- Supports network policy and flow control.</li> <li>- Container Service provides you with continuous network integration and the best network optimization.</li> </ul>
Server Load Balancer	<ul style="list-style-type: none"> <li>- Supports creating Internet or intranet Server Load Balancer instances.</li> <li>- 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.</li> </ul>
Storage	<p>Container Service integrates with Alibaba Cloud cloud disk , Network Attached Storage (NAS), and block storage, and provides the standard FlexVolume drive.</p> <p>Self-built Kubernetes clusters cannot use the storage resources on the cloud. Alibaba Cloud Container Service provides the best seamless integration.</p>

Function	Description
O&M	<ul style="list-style-type: none"> <li>- Integrates with Alibaba Cloud Log Service and CloudMonitor.</li> <li>- Supports auto scaling.</li> </ul>
Image repository	<ul style="list-style-type: none"> <li>- Provides high availability.</li> <li>- Supports image scan.</li> <li>- Supports P2P distribution and replication across regions.</li> </ul> <p>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&amp;M burden and upgrade pressure.</p>
Stability	<ul style="list-style-type: none"> <li>- The dedicated team guarantees the stability of the container.</li> <li>- Each Linux version and Kubernetes version are provided to you after strict tests.</li> <li>- 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.</li> </ul>
High availability	<ul style="list-style-type: none"> <li>- Supports multiple zones.</li> <li>- Supports backup and disaster recovery.</li> </ul>
Technical support	<ul style="list-style-type: none"> <li>- Provides the Kubernetes upgrade capabilities. Supports upgrading a Kubernetes cluster to the latest version with one click.</li> <li>- Alibaba Cloud container team is responsible for solving problems about containers in your environment.</li> </ul>

Q: How can I get started with ACK?

Please see [#unique\\_4](#) for more details on how to get started with ACK.

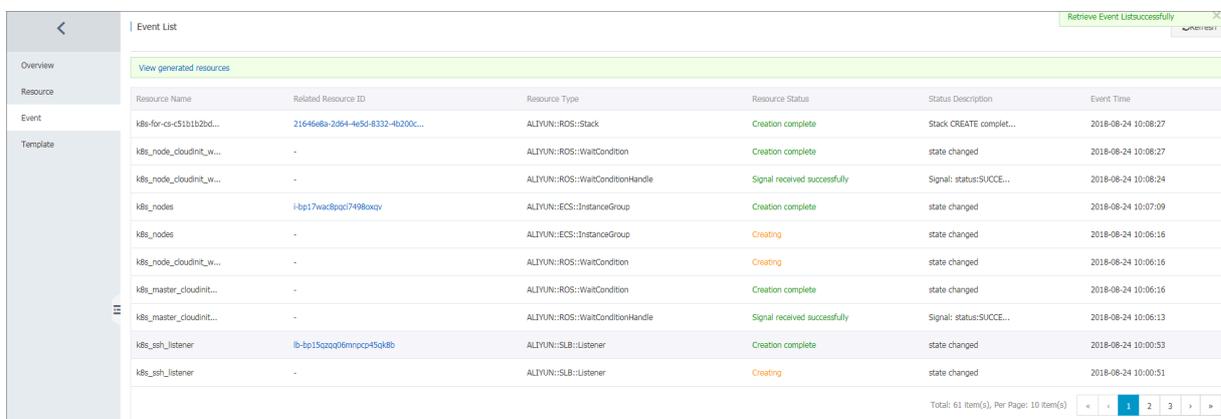
## 2 Failed to create a Kubernetes cluster

### 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-2d54-4e5d-8332-4b200c...	ALYUN::ROS::Stack	Creation complete	Stack CREATE complet...	2018-08-24 10:08:27
k8s_node_cloudinit_w...	-	ALYUN::ROS::WaitCondition	Creation complete	state changed	2018-08-24 10:08:27
k8s_node_cloudinit_w...	-	ALYUN::ROS::WaitConditionHandle	Signal received successfully	Signal: status:SUCCE...	2018-08-24 10:08:24
k8s_nodes	i-bp17wac9pcc7498wqv	ALYUN::ECS::InstanceGroup	Creation complete	state changed	2018-08-24 10:07:09
k8s_nodes	-	ALYUN::ECS::InstanceGroup	Creating	state changed	2018-08-24 10:06:16
k8s_node_cloudinit_w...	-	ALYUN::ROS::WaitCondition	Creating	state changed	2018-08-24 10:06:16
k8s_master_cloudinit...	-	ALYUN::ROS::WaitCondition	Creation complete	state changed	2018-08-24 10:06:16
k8s_master_cloudinit...	-	ALYUN::ROS::WaitConditionHandle	Signal received successfully	Signal: status:SUCCE...	2018-08-24 10:06:13
k8s_slb_listener	lb-bp15zqzq96mrpcc45ol8b	ALYUN::SLB::Listener	Creation complete	state changed	2018-08-24 10:00:53
k8s_slb_listener	-	ALYUN::SLB::Listener	Creating	state changed	2018-08-24 10:00:51

### 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.

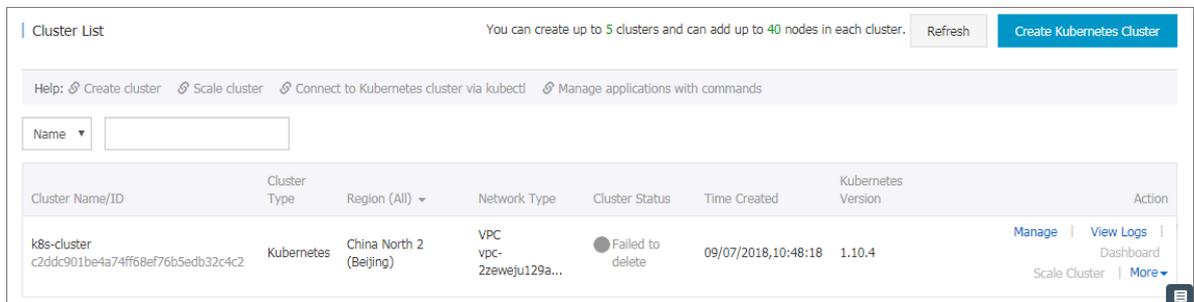


### Note:

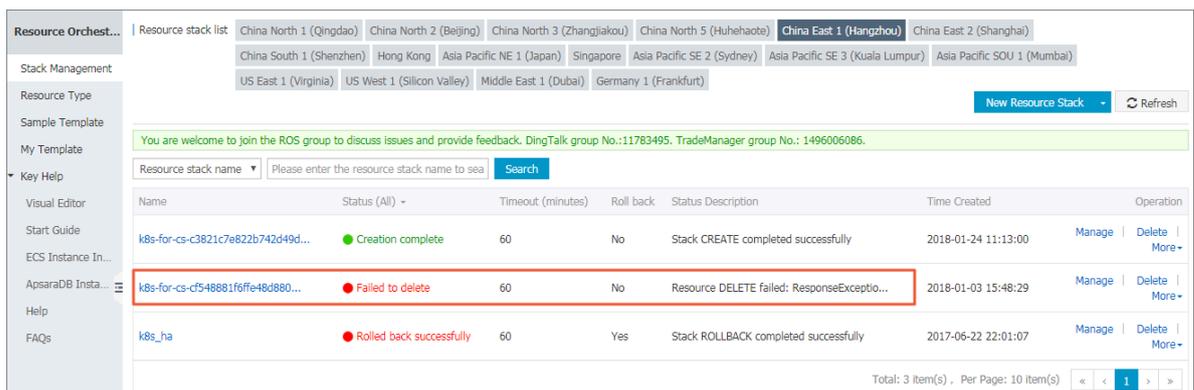
For more information about the resources automatically created by ROS when the Kubernetes cluster is created, see [#unique\\_7](#).

## Solutions

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



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.



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.

Resource	Resource Name	Resource Type	Status	Reason	Created At	Deleted At	Details
k8s_master_ssh_inter...		ALYUN::SLB::VServerGroup	complete		15:48:29	15:48:29	Details
k8s_NAT_Gateway		ALYUN::ECS::NatGateway	Deleted successfully	state changed	2018-01-03 15:48:29	2018-01-04 15:50:15	Details
k8s_NAT_Gateway_SNAT...		ALYUN::ECS::SNatEntry	Initialization complete		2018-01-03 15:48:29	2018-01-03 15:48:29	Details
k8s_node_cloudinit_w...		ALYUN::ROS::WaitCondition	Initialization complete		2018-01-03 15:48:29	2018-01-03 15:48:29	Details
k8s_node_cloudinit_w...		ALYUN::ROS::WaitConditionHandle	Deleted successfully	state changed	2018-01-03 15:48:29	2018-01-04 15:50:05	Details
k8s_nodes		ALYUN::ECS::InstanceGroup	Initialization complete		2018-01-03 15:48:29	2018-01-03 15:48:29	Details
k8s_sg	sg-bp12cvr14a2wjz9iqf40	ALYUN::ECS::SecurityGroup	Deleted successfully	state changed	2018-01-03 15:48:29	2018-01-04 15:50:13	Details
k8s_vpc	vpc-bp1weju11gaal4szwm7bi	ALYUN::ECS::VPC	Creation complete	state changed	2018-01-03 15:48:29	2018-01-03 15:48:33	Details
k8s_vswitch	vsw-bp1olpu0k50buqhtdbq	ALYUN::ECS::VSwitch	Failed to delete	ResponseException: r...	2018-01-03 15:48:29	2018-01-04 15:55:02	Details
KubernetesMasterRole	KubernetesMasterRole-8bdfb114...	ALYUN::RAM::Role	Deleted successfully	state changed	2018-01-03 15:48:29	2018-01-04 15:50:14	Details

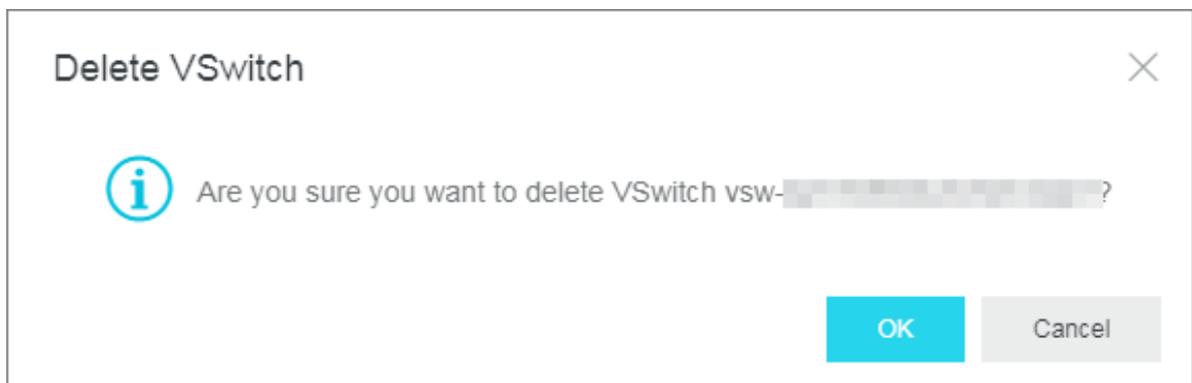
4. Go to the console in which the resource that failed to be deleted resides and find that resource.

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 ID/Name	Number of ECS Instances	CIDR	Status	Zone	Number of Available Private IPs	Created At	Default VSwitch	Description	Actions
vsw-bp1olpu0k50buqhtdbq	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 VSwitch to manually delete it.

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



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 [#unique\\_10](#).

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 : "
2e55dbef1fd b5fdb96b75 ff144a3394 89417b146b ", GitTreeState : "
clean "}
Server: & version . Version { SemVer : " v2 . 11 . 0 ", GitCommit : "
2e55dbef1fd b5fdb96b75 ff144a3394 89417b146b ", GitTreeState : "
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
imagePullSecrets:
- name: regsecret
```

where:

- **imagePullSecrets** 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 serviceaccounts default -o yaml > ./ sa .
yaml

cat sa . yaml

apiVersion : v1
kind : ServiceAccount
metadata :
  creationTimestamp : 2015 - 08 - 07T22 : 02 : 39Z
  name : default
  namespace : default
  resourceVersion : " 243024 "          ## Pay attention to
  this item
  selfLink : / api / v1 / namespaces / default / serviceaccounts /
  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 : ServiceAccount
metadata :
  creationTimestamp : 2015 - 08 - 07T22 : 02 : 39Z
  name : default
  namespace : default
  selfLink : / api / v1 / namespaces / default / serviceaccounts /
  default
  uid : 052fb0f4 - 3d50 - 11e5 - b066 - 42010af0d7 b6
  secrets :
- name : default - token - uudge
  imagePullSecrets :          ## Add this item
- name : regsecret
```

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

```
kubectl replace serviceaccount default -f ./ sa . yaml
```

```
serviceaccount "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 :
    matchLabels :
      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 :
            - containerPort : 8080containerPort : 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 :
  imagePullSecrets :
    - name: registry
```

## 7 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](#).

## 8 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-application-controller.yml` command to deploy alicloud-application-controller. In `alicloud-application-controller.yml`, enter the following orchestration template:

```
apiVersion : extensions / v1beta1
kind : Deployment
metadata :
  name : alicloud - application - controller
  labels :
    owner : aliyun
    app : alicloud - application - controller
    namespace : kube - system
spec :
  replicas : 1
  selector :
    matchLabels :
      owner : aliyun
      app : alicloud - application - controller
  template :
    metadata :
      labels :
        owner : aliyun
        app : alicloud - application - controller
    annotations :
      scheduler.alpha.kubernetes.io/critical-pod : ''
    spec :
      tolerations :
        - effect : NoSchedule
          operator : Exists
          key : node-role.kubernetes.io/master
        - effect : NoSchedule
          operator : Exists
          key : node.cloudprovider.kubernetes.io/uninitialized
      containers :
        - name : alicloud - application - controller
          image : registry.cn-hangzhou.aliyuncs.com/acs/aliyun-app-lifecycle-manager:0.1-c8d5da8
```

```
imagePullPolicy : IfNotPresent
serviceAccount : admin
```

## 9 Can I manually set a security group for a Kubernetes cluster?

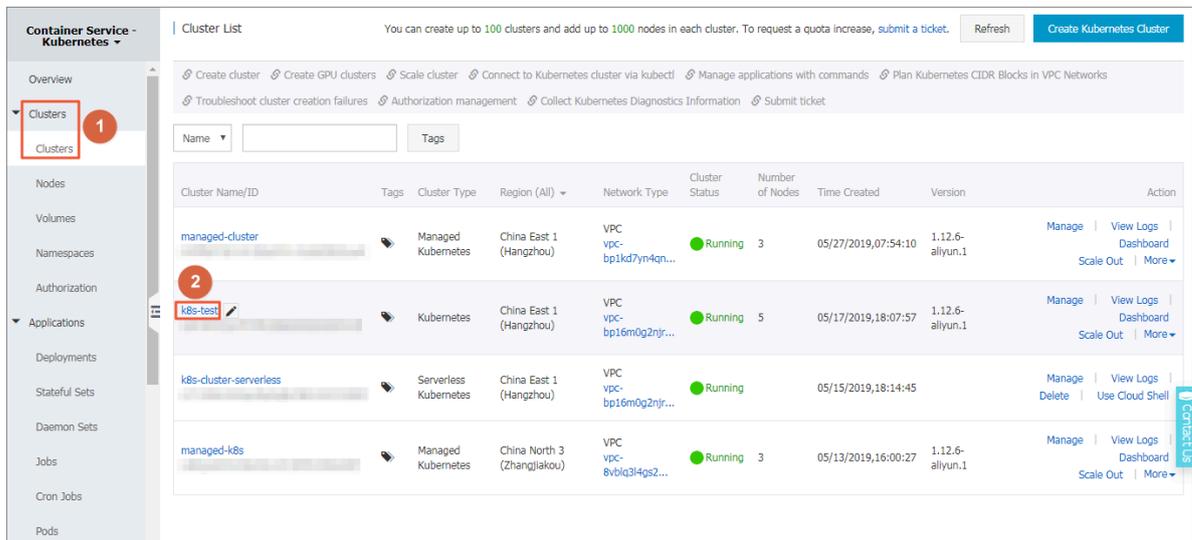
---

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.

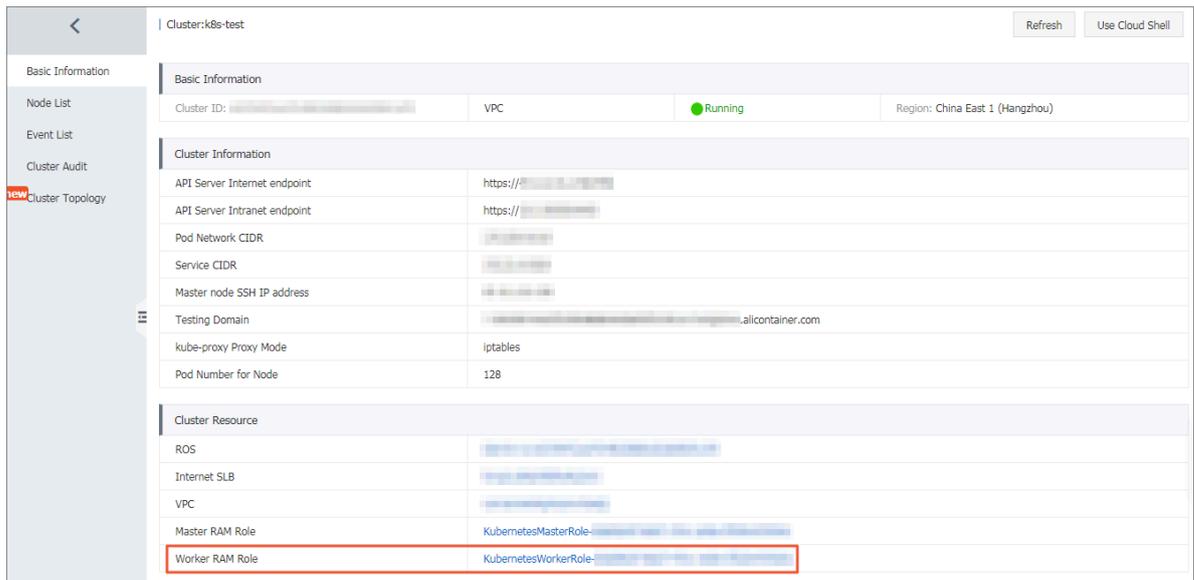
# 10 How do I customize a RAM role for a Kubernetes cluster?

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.

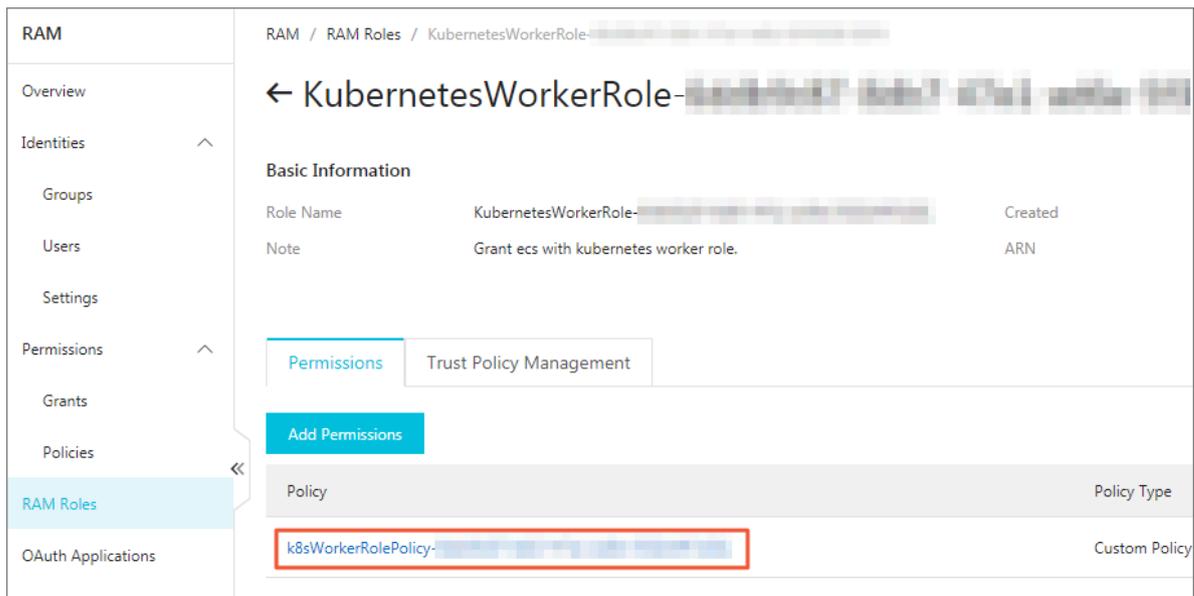
1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Container Service-Kubernetes, choose **Clusters > Clusters**.
3. Click the target cluster name.



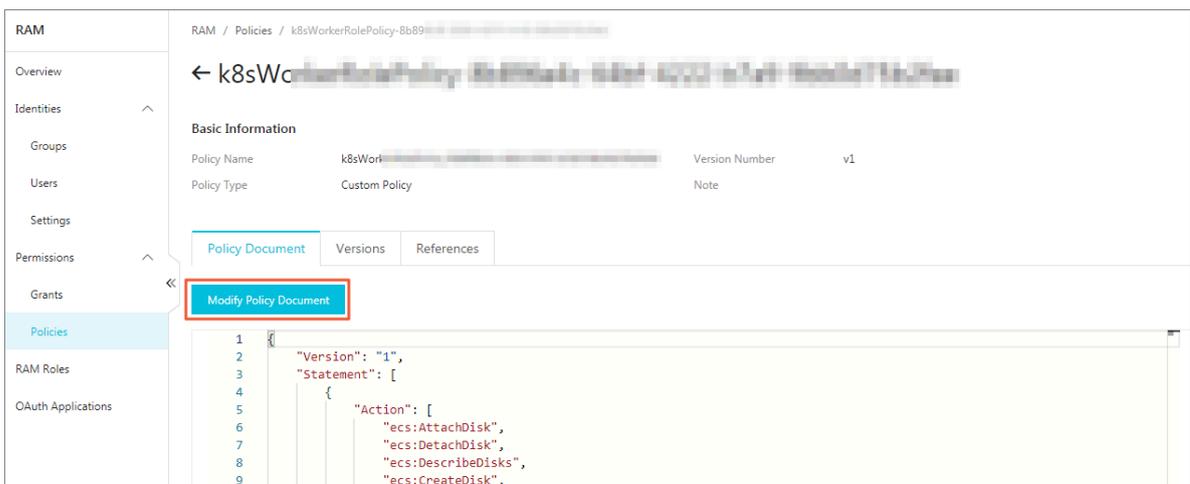
4. In the Cluster Resources area, click Worker RAM Role.



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



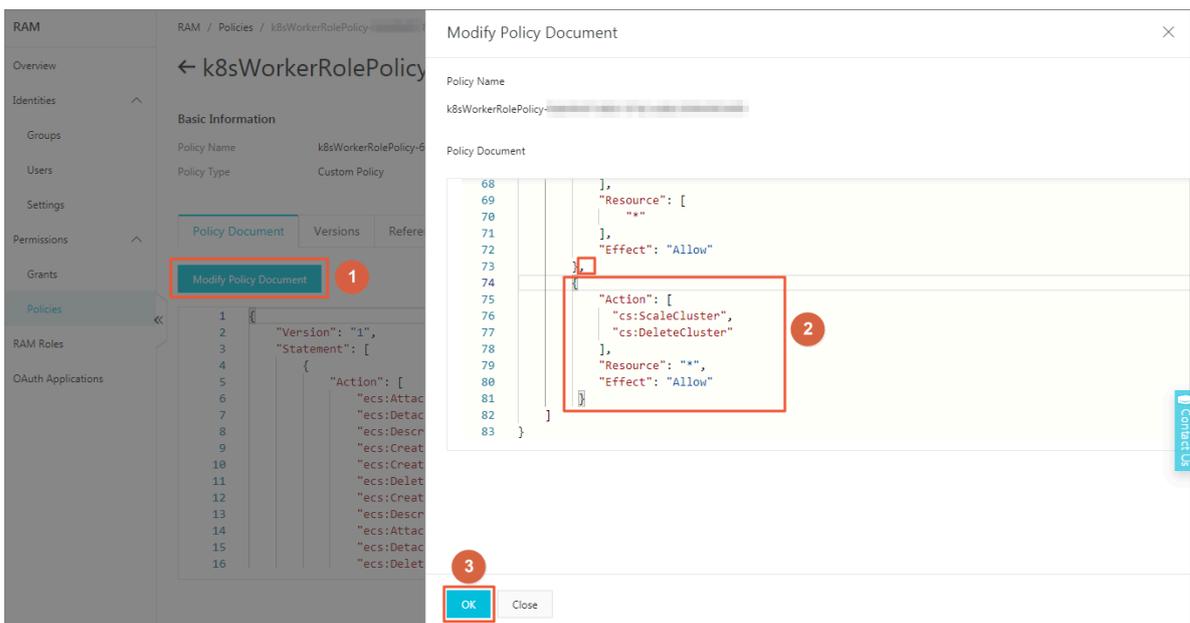
6. On the Policy Document tab page, click Modify Policy Document.



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 [#unique\\_17/grants\\_17\\_Connect\\_42\\_table\\_pzw\\_5s2\\_xdb](#).

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



## 11 How do I use a RAM user account to grant RBAC permissions to any other RAM user?

To grant RBAC permissions to any other RAM user, follow these steps:

1. Set the preset role `Administrator` or the custom role `cluster-admin` for your RAM user account in the target Kubernetes cluster or namespace, and grant corresponding RAM permissions to your RAM user account.



说明:

The following RAM permissions are required:

- View any other RAM user accounts.
  - Grant RAM permissions to any other RAM user accounts.
  - View RBAC permissions that are granted to any other RAM user accounts.
  - Grant RBAC permissions to any other RAM user accounts.
2. Use your RAM user account to log on to the RAM console to grant RAM permissions to the target RAM user. For more information, see [#unique\\_17](#).

The following is a sample policy of RAM permissions:

```
{
  "Statement": [{
    "Action": [
      "ram: Get *",
      "ram: List *",
      "cs: GetUserPermissions",
      "cs: GetSubUsers",
      "cs: GrantPermission"
    ],
    "Resource": "*",
    "Effect": "Allow"
  },
  {
    "Action": [
      "ram: AttachPolicyToUser",
      "ram: AttachPolicy"
    ],
    "Effect": "Allow",
    "Resource": [
      "acs:ram:*:*:policy/xxxxxx",
      "acs:*:*:*:user/*"
    ]
  }
],
  "Version": "1"
}
```



说明:

xxxxxx indicates the name of a RAM permission policy. For example, if you replace xxxxxx with \*, the target RAM user account then can grant all the RAM permissions to any other RAM users.

3. Use your RAM user account to grant RBAC permissions to the target RAM user. For more information, see [#unique\\_19](#).

## 12 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 [#unique\\_10](#).
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 172.19.1.1 47.111.1.1 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 "kubernetes.do.not.delete:" substr("a"$1$2$3$4$5, 1, 32)}'
```

- 
4. Log on to the [SLB console](#)~~SLB console~~. Then, select the corresponding region to which the target SLB instance belongs, and use the IP address 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 [#unique\\_21](#).

## 13 Failed to start a container on the GPU node

### Symptom

When a container fails to start after you restart your Kubelet and Docker on the GPU node, the following codes are displayed:

```
# service kubelet stop
Redirectin g to / bin / systemctl stop kubelet . service
# service docker stop
Redirectin g to / bin / systemctl stop docker . service
# service docker start
Redirectin g to / bin / systemctl start docker . service
# service kubelet start
Redirectin g to / bin / systemctl start kubelet . service

# docker ps
CONTAINER ID          IMAGE               COMMAND             PORTS
CREATED              NAMES
```

### Cause

The type of `Cgroup Driver` of Docker is `cgroupfs` (which is incorrect).



#### Note:

You can run the following command to check the type of `Cgroup Driver` :

```
docker info | grep -i cgroup
Cgroup Driver : cgroupfs
```

### Solution

1. Back up the `/ etc / docker / daemon . json` file and run the following command to update the file:

```
cat >/ etc / docker / daemon . json <<- EOF
{
  " default - runtime ": " nvidia ",
  " runtimes ": {
    " nvidia ": {
      " path ": "/ usr / bin / nvidia - container - runtime ",
      " runtimeArg s ": []
    }
  },
  " exec - opts ": [" native . cgroupdriv er = systemd "],
  " log - driver ": " json - file ",
  " log - opts ": {
    " max - size ": " 100m ",
    " max - file ": " 10 "
  },
  " oom - score - adjust ": - 1000 ,
  " storage - driver ": " overlay2 ",
```

```
" storage - opts ":[" overlay2 . override_k ernel_chec k =
true "],
" live - restore ": true
}
EOF
```

## 2. Run the following command to restart Docker and Kubelet:

```
# service kubelet stop
Redirectin g to / bin / systemctl stop kubelet . service
# service docker restart
Redirectin g to / bin / systemctl restart docker . service
# service kubelet start
Redirectin g to / bin / systemctl start kubelet . service
```

## 3. Run the following command to verify that the type of `Cgroup Driver` is

`systemd` .

```
# docker info | grep - i cgroup
Cgroup Driver : systemd
```

## 14 How do I combine kubeconfig files of multiple clusters?

---

If your Alibaba Cloud account has access to multiple clusters, you can run the following command to combine kubeconfig files of different clusters into one credential file:

```
KUBECONFIG = file1 : file2 : file3  kubectl  config  view  --  
merge  --  flatten  >  ~/.  kube  /  all  -  config  
export  KUBECONFIG  =~/.  kube  /  all  -  config
```

- Run the following command to view the available contexts:

```
kubectl  config  get  -  contexts
```

- Run the following command to view the help information:

```
kubectl  config  --  help
```

- Run the following command to change the contexts:

```
kubectl  config  use  -  context  {  your  -  contexts  }
```

For more information about kubeconfig files, see [Configure access to multiple clusters](#).



### Note:

You can obtain kubeconfig access credentials of the current user on the cluster list page of the Container Service console. You can also call an [API](#) action to obtain a specific kubeconfig file.