

Alibaba Cloud Aliyun Container for Kubernetes

User Guide

Issue: 20190614

Legal disclaimer

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

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

or trust in this document, even if Alibaba Cloud has been notified of the possibility of such a loss.

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

Generic conventions

Table -1: Style conventions

Style	Description	Example
	This warning information indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	 Danger: Resetting will result in the loss of user configuration data.
	This warning information indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	 Warning: Restarting will cause business interruption. About 10 minutes are required to restore business.
	This indicates warning information, supplementary instructions, and other content that the user must understand.	 Notice: Take the necessary precautions to save exported data containing sensitive information.
	This indicates supplemental instructions, best practices, tips, and other content that is good to know for the user.	 Note: You can use Ctrl + A to select all files.
>	Multi-level menu cascade.	Settings > Network > Set network type
Bold	It is used for buttons, menus, page names, and other UI elements.	Click OK.
Courier font	It is used for commands.	Run the <code>cd / d C :/ windows</code> command to enter the Windows system folder.
<i>Italics</i>	It is used for parameters and variables.	<code>bae log list --instanceid Instance_ID</code>
[] or [a b]	It indicates that it is a 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 Kubernetes cluster.....	1
1.1 Security bulletins.....	1
1.1.1 Vulnerability fix: <i>CVE - 2019 - 5736</i> in <i>runc</i>	1
1.1.2 Vulnerability fix: <i>CVE - 2018 - 18264</i> for Kubernetes dashboard.....	3
1.1.3 Vulnerability fix: <i>CVE - 2018 - 1002105</i>	5
1.2 Introduction.....	6
1.2.1 Overview.....	6
1.2.2 Alibaba Cloud Kubernetes vs. self-built Kubernetes.....	7
1.3 Authorization management.....	11
1.3.1 Role authorization.....	11
1.3.2 Use the Container Service console as a RAM user.....	15
1.3.3 Create custom authorization policies.....	17
1.3.4 Grant Kubernetes permissions to a RAM user.....	20
1.4 Cluster management.....	29
1.4.1 View cluster overview.....	29
1.4.2 Create a Kubernetes cluster.....	31
1.4.3 Configure a Kubernetes GPU cluster to support GPU scheduling....	41
1.4.4 Upgrade the NVIDIA driver of a Kubernetes cluster GPU node.....	47
1.4.5 Create a multi-zone Kubernetes cluster.....	53
1.4.6 Connect to a Kubernetes cluster by using <i>kubectl</i>	62
1.4.7 Use <i>kubectl</i> commands in Cloud Shell to manage a Kubernetes cluster.....	63
1.4.8 Use a ServiceAccount token to access a managed Kubernetes cluster.....	65
1.4.9 Access a Kubernetes cluster by using SSH.....	68
1.4.10 Set the SSH key pair logon mode for a Kubernetes cluster.....	70
1.4.11 Create a managed Kubernetes cluster.....	72
1.4.12 Upgrade a Kubernetes cluster.....	78
1.4.13 Upgrade a system component.....	80
1.4.14 Update the Kubernetes cluster certificates that are about to expire.....	83
1.4.15 Scale a Kubernetes cluster.....	86
1.4.16 Autoscale a Kubernetes cluster.....	87
1.4.17 Delete a Kubernetes cluster.....	98
1.4.18 Upgrade the Heapster components to the metrics-server component.....	100
1.4.19 Create a Kubernetes cluster that supports Windows.....	107

1.5 Node management.....	116
1.5.1 Add an existing ECS instance.....	116
1.5.2 View node list.....	120
1.5.3 Node monitoring.....	121
1.5.4 Manage node labels.....	122
1.5.5 Set node scheduling.....	123
1.5.6 Remove a node.....	125
1.5.7 Use Alibaba Cloud Kubernetes GPU node labels for scheduling....	128
1.5.8 View resource request and limit on nodes.....	134
1.5.9 Mount a disk to a Kubernetes cluster node.....	135
1.5.10 Mount a disk to the Docker data directory.....	140
1.6 Namespace management.....	145
1.6.1 Create a namespace.....	145
1.6.2 Set resource quotas and limits for a namespace.....	148
1.6.3 Edit a namespace.....	152
1.6.4 Delete a namespace.....	153
1.7 Service catalog management.....	154
1.7.1 Overview.....	154
1.7.2 Enable service catalog function.....	155
1.8 Application management.....	157
1.8.1 Create a deployment application by using an image.....	157
1.8.2 Create a StatefulSet application by using an image.....	179
1.8.3 Create a Job application by using an image.....	203
1.8.4 Create an application in Kubernetes dashboard.....	216
1.8.5 Create a Linux application by using an orchestration template....	218
1.8.6 Create a Windows application by using an orchestration template.....	222
1.8.7 Manage applications by using commands.....	227
1.8.8 Simplify Kubernetes application deployment by using Helm.....	227
1.8.9 Use an application trigger.....	236
1.8.10 View pods.....	238
1.8.11 Create a service.....	239
1.8.12 Scale a service.....	244
1.8.13 View a service.....	246
1.8.14 Update a service.....	248
1.8.15 Delete a service.....	252
1.8.16 Change container configurations.....	253
1.8.17 Schedule a pod to a specified node.....	254
1.8.18 View image list.....	256
1.8.19 Use an image Secret.....	257
1.8.20 Pull a private image without a password.....	261
1.9 Network management.....	266
1.9.1 Networks supported by Alibaba Cloud Container Service for Kubernetes.....	266
1.9.2 Access services by using Server Load Balancer.....	267

1.9.3 Support for Ingress.....	290
1.9.4 Configure Ingress monitoring.....	295
1.9.5 Ingress configurations.....	298
1.9.6 Create an Ingress in the Container Service console.....	301
1.9.7 View Ingress details.....	312
1.9.8 Update an Ingress.....	313
1.9.9 Deleting a route.....	315
1.9.10 Terway network plugin.....	316
1.9.11 Allocate an ENI to a pod.....	318
1.9.12 Use a network policy.....	320
1.10 Service mesh.....	326
1.10.1 Manage traffic.....	326
1.11 Server Load Balancer and Ingress management.....	335
1.11.1 Overview.....	335
1.12 Config Map and Secret management.....	335
1.12.1 Create a Config Map.....	336
1.12.2 Use a config map in a pod.....	340
1.12.3 View a ConfigMap.....	345
1.12.4 Update a config map.....	346
1.12.5 Delete a config map.....	350
1.12.6 Create a Secret.....	352
1.12.7 View secret details.....	355
1.12.8 Update a secret.....	356
1.12.9 Delete a secret.....	357
1.13 Storage management.....	358
1.13.1 Overview.....	358
1.13.2 Install the plug-in.....	359
1.13.3 Use Alibaba Cloud cloud disk volumes.....	363
1.13.4 Use NAS file systems of Alibaba Cloud.....	369
1.13.5 Use Alibaba Cloud OSS volumes.....	378
1.13.6 Create a Persistent Volume Claim.....	383
1.13.7 Use a persistent volume claim.....	386
1.14 Log management.....	387
1.14.1 Application log management.....	387
1.14.2 View cluster logs.....	388
1.14.3 Use Log Service to collect Kubernetes cluster logs.....	388
1.14.4 A solution to log collection problems of Kubernetes clusters by using log-pilot, Elasticsearch, and Kibana.....	400
1.14.5 Configure Log4jAppender for Kubernetes and Log Service.....	407
1.15 Monitoring management.....	412
1.15.1 Deploy the Prometheus monitoring system.....	412
1.15.2 Group-based monitoring and alarms.....	417
1.15.3 Integration and usage with CloudMonitor.....	422
1.15.4 Use Grafana to display monitoring data.....	428
1.15.5 Use an HPA auto scaling container.....	436

1.15.6 Monitor a Kubernetes cluster and send alarm notifications by using DingTalk.....	440
1.16 Security management.....	447
1.16.1 Security.....	448
1.16.2 Kube-apiserver audit logs.....	449
1.16.3 Implement secure access through HTTPS in Kubernetes.....	463
1.17 Release management.....	474
1.17.1 Manage a Helm-based release.....	474
1.17.2 Use batch release on Alibaba Cloud Container Service for Kubernetes.....	479
1.18 Istio management.....	483
1.18.1 Overview.....	483
1.18.2 Deploy Istio.....	485
1.18.3 Update Istio.....	499
1.18.4 Delete Istio.....	501
1.18.5 Upgrade Istio components.....	504
1.19 Template management.....	507
1.19.1 Create an orchestration template.....	507
1.19.2 Edit an orchestration template.....	511
1.19.3 Save an existing orchestration template as a new one.....	513
1.19.4 Download an orchestration template.....	514
1.19.5 Delete an orchestration template.....	515
1.20 App catalog management.....	516
1.20.1 App catalog overview.....	516
1.20.2 View app catalog list.....	517
1.21 Auto Scaling.....	518
1.21.1 Autoscale a Kubernetes cluster.....	518
1.21.2 Deploy a virtual node.....	529
2 Serverless Kubernetes cluster.....	536
2.1 Overview.....	536
2.2 Kubernetes supported functions.....	537
2.3 Cluster management.....	538
2.3.1 Create a serverless Kubernetes cluster.....	538
2.3.2 Connect to a Kubernetes cluster by using kubectl.....	540
2.3.3 Delete a cluster.....	540
2.4 Application management.....	541
2.4.1 Manage applications by using commands.....	541
2.4.2 Create an application by using an image.....	542
2.4.3 Create a service.....	545
2.4.4 Delete a service.....	549
2.4.5 View pods.....	549
2.4.6 View services.....	550
2.5 Config Map.....	551
2.5.1 Create a ConfigMap.....	551
2.5.2 Delete a config map.....	552

2.5.3 Modify a config map.....	552
2.6 Server Load Balancer management.....	553
2.6.1 Server Load Balancer.....	553
2.6.2 Use Ingress to provide Layer-7 service access.....	557
2.7 Log management.....	565
2.7.1 Overview.....	566
2.7.2 View cluster logs.....	566
2.7.3 Collect logs by using Alibaba Cloud Log Service.....	567
2.8 Service discovery based on Alibaba Cloud DNS Private Zone.....	569

1 Kubernetes cluster

1.1 Security bulletins

1.1.1 Vulnerability fix: *CVE - 2019 - 5736* in runc

The security vulnerability *CVE - 2019 - 5736* in runc has been fixed for Alibaba Cloud Container Service for Kubernetes. This topic describes the impacts of this vulnerability and how to remove it.

Background

The security vulnerability may occur with Docker, containerd, or any other containers that use runc. This vulnerability gives attackers the ability to use a specific container image or run the `exec` command to obtain the file handle used by the running host runc. Attackers can overwrite the host runc binary file, then obtain root permission to access the host, and execute commands as with root permission.

For more information about security vulnerability *CVE - 2019 - 5736*, see <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2019-5736>.

Affected clusters

- Alibaba Cloud Container Service clusters affected by the vulnerability:
 - All Docker Swarm clusters from versions earlier than Docker v18.09.02.
 - All Kubernetes clusters except for Serverless Kubernetes clusters.
- Self-built Docker and Kubernetes clusters affected by the vulnerability:
 - All clusters that use Docker versions earlier than v18.09.2.
 - All clusters that use runc v1.0-rc6 or earlier.



Note:

For both of the preceding vulnerability cases, we recommend that you consult your Docker or runc vendor for a solution.

Resolution

- **Method 1:** Create new Kubernetes clusters of v1.11 or v1.12. Kubernetes clusters of these two versions run the latest version of Docker, which are protected from this vulnerability.
- **Method 2:** Upgrade the Docker version of all existing clusters to v18.09.2 or later. Using this method will interrupt your cluster containers and services.
- **Method 3:** Only upgrade runc. This method is applicable to clusters running Docker v17.06. We recommend that you upgrade the runc binary file of each cluster node individually to avoid a service interruption caused by upgrading the Docker engine. To upgrade a runc binary file, complete the following steps:

1. Run the following command to locate docker-runc:



Note:

Usually, docker-runc is located in `/usr/bin/docker-runc`.

```
which docker - runc
```

2. Run the following command to back up the original runc:

```
mv /usr/bin/docker-runc /usr/bin/docker-runc.orig
.$(date +%Y%m%d%H%M%S)
```

3. Run the following command to download the fixed runc:

```
curl -o /usr/bin/docker-runc -sSL https://acs-public-mirror.oss-cn-hangzhou.aliyuncs.com/runc/docker-runc-17.06-amd64
```

4. Run the following command to set permission availability for docker-runc:

```
chmod +x /usr/bin/docker-runc
```

5. Run the following command to test whether runc works normally:

```
docker-runc -v
# runc version 1.0.0-rc3
# commit : fc48a25bde 6fb041aae0 977111ad81 41ff396438
# spec : 1.0.0-rc5
docker run -it --rm ubuntu echo OK
```

6. To upgrade the runc binary file of a Kubernetes cluster GPU node, you must also install nvidia-runtime by completing the following steps:

- a. Run the following command to locate nvidia-container-runtime:



Note:

Usually, `nvidia-container-runtime` is located in `/usr/bin/nvidia-container-runtime`.

which `nvidia-container-runtime`

- b. Run the following command to back up the original `nvidia-container-runtime`:

```
mv /usr/bin/nvidia-container-runtime /usr/bin/nvidia-container-runtime.orig.$(date +%Y%m%d%H%M%S)
```

- c. Run the following command to download the fixed `nvidia-container-runtime`:

```
curl -o /usr/bin/nvidia-container-runtime -sL https://acs-public-mirror.oss-cn-hangzhou.aliyuncs.com/runc/nvidia-container-runtime-17.06-amd64
```

- d. Run the following command to set permission availability for `nvidia-container-runtime`:

```
chmod +x /usr/bin/nvidia-container-runtime
```

- e. Run the following command to test whether `nvidia-container-runtime` works normally:

```
nvidia-container-runtime -v
# runc version 1.0.0-rc3
# commit : fc48a25bde 6fb041aae0 977111ad81 41ff396438 - dirty
# spec : 1.0.0-rc5

docker run -it --rm -e NVIDIA_VISIBLE_DEVICES=all ubuntu nvidia-smi -L
# GPU 0 : Tesla P100-PCIE-16GB (UUID : GPU-122e199c-9aa6-5063-0fd2-da009017e6dc)
```



Note:

This test is performed on a node of the GPU P100 model. Test outputs vary by GPU models.

1.1.2 Vulnerability fix: *CVE - 2018 - 18264* for Kubernetes dashboard

Alibaba Cloud Container Service for Kubernetes has fixed dashboard vulnerability *CVE - 2018 - 18264*. This topic describes the dashboard versions affected by the vulnerability and how to fix the vulnerability. The Kubernetes dashboards that are built in Alibaba Cloud Container Service for Kubernetes are not affected by this

vulnerability because they work in the hosted form and their security settings were upgraded before the vulnerability occurred.

Background information

A security vulnerability, that is, *CVE - 2018 - 18264* , was discovered in Kubernetes dashboards of V1.10 and earlier versions. This vulnerability allowed attackers to bypass identity authentication and read secrets within the cluster by using the dashboard logon account.

The Kubernetes dashboards that are built in Alibaba Cloud Container Service for Kubernetes are not affected by this vulnerability because they work in the hosted form and their security settings were upgraded before the vulnerability occurred.

For more information about security vulnerability *CVE - 2018 - 18264* , see:

- <https://github.com/kubernetes/dashboard/pull/3289>
- <https://github.com/kubernetes/dashboard/pull/3400>
- <https://github.com/kubernetes/dashboard/releases/tag/v1.10.1>

Conditions required to determine that a Kubernetes dashboard is vulnerable

Your dashboard is vulnerable if you have independently deployed Kubernetes dashboard V1.10 or earlier versions (V1.7.0 to V1.10.0) that supports the logon function in your Kubernetes cluster, and you have used custom certificates.

Resolution

- If you do not need a dashboard that is deployed independently, run the following command to remove the Kubernetes dashboard from your cluster:

```
kubectl -- namespace kube - system delete deployment  
kubernetes - dashboard
```

- If you need an independently deployed dashboard, upgrade your dashboard to V1.10.1. For more information, see <https://github.com/kubernetes/dashboard/releases/tag/v1.10.1>.
- If you use the dashboard hosted by Alibaba Cloud Container Service for Kubernetes, you can continue to use your dashboard in the Container Service console because the dashboard was upgraded before the vulnerability occurred.

1.1.3 Vulnerability fix: *CVE - 2018 - 1002105*

Alibaba Cloud has fixed system vulnerability *CVE - 2018 - 1002105* . This topic describes the impacts of this vulnerability and how to remove it.

This vulnerability does not affect Serverless Kubernetes clusters. Serverless Kubernetes was upgraded before the vulnerability occurred.

Background information

Engineers of the Kubernetes community have found security vulnerability *CVE - 2018 - 1002105* . Kubernetes users can gain access to the backend service by forging the request and escalating the permission on the established API Server connection. Alibaba Cloud has fixed this vulnerability. To remove the vulnerability, you need to log on to the Container Service console and upgrade Kubernetes to the latest version.

For more information about the vulnerability *CVE - 2018 - 1002105* , see <https://github.com/kubernetes/kubernetes/issues/71411>.

Affected Kubernetes versions:

- Kubernetes v1.0.x-1.9.x
- Kubernetes v1.10.0-1.10.10 (fixed in v1.10.11)
- Kubernetes v1.11.0-1.11.4 (fixed in v1.11.5)
- Kubernetes v1.12.0-1.12.2 (fixed in v1.12.3)

Affected configurations:

- Kubernetes cluster, which runs on Container Service and uses an extension API server. Furthermore, the extension API server network is directly accessible to the cluster component, kube-apiserver.
- Kubernetes cluster, which runs on Container Service and has opened permissions to interfaces such as pod exec, attach, and portforward. Then, users can use the vulnerability to obtain permissions to access all kubelet APIs of the cluster.

Cluster configuration of Alibaba Cloud Container Service for Kubernetes

- The API server of a Kubernetes cluster that runs on Container Service has RBAC enabled by default. That is, the API server denies anonymous user access through primary account authorization. Furthermore, the starting parameter of Kubelet is

`anonymous - auth = false` , providing security access control against external attacks.

- If your Kubernetes cluster has multiple RAM users, the RAM users may gain unauthorized access to the backend service through interfaces such as `pod exec` , `attach` , and `portforward` . If your cluster has no RAM users, you do not need to worry about the vulnerability.
- RAM users do not have access to aggregate API resources by default without custom authorization from the primary account.

Solution

Log on to the Container Service console to upgrade your cluster. For more information, see [Upgrade a Kubernetes cluster](#).

- If your cluster is V1.11.2, upgrade it to V1.11.5.
- If your cluster is V1.10.4, upgrade it to V1.10.11 or V1.11.5.
- If your cluster is V1.9 or earlier, upgrade it to V1.10.11 or V1.11.5. When you upgrade the cluster from V1.9 to V1.10 or V1.11, upgrade the flexvolume plugin through the console if your cluster uses cloud disk volumes.



Note:

In the Container Service console, select the target cluster and choose **More > Addon Upgrade**. In the Addon Upgrade dialog box, select **flexvolume** and click **Upgrade**.

1.2 Introduction

1.2.1 Overview

Kubernetes is a popular open-source container orchestration technology. To allow you to use Kubernetes to manage container applications in Alibaba Cloud, Alibaba Cloud Container Service provides support for Kubernetes clusters.

You can create a safe and high-availability Kubernetes cluster in the Container Service console. The Kubernetes cluster integrates with the virtualization, storage , network, and security capabilities of Alibaba Cloud to provide scalable, high-performance container application management, simplify cluster creation and

expansion, and focus on the development and management of containerized applications.

Kubernetes supports the deployment, expansion, and management of containerized applications, and provides the following features:

- Elastic expansion and self-reparation.
- Service discovery and server load balancing.
- Service release and rollback.
- Secret and configuration management.

Limits

- Currently, Kubernetes clusters only support Linux containers. The support for Kubernetes Windows containers is in the works.
- Currently, Kubernetes clusters only support Virtual Private Cloud (VPC). You can select to create a VPC or use an existing VPC when creating a Kubernetes cluster.

Related open-source projects

- Alibaba Cloud Kubernetes Cloud Provider: <https://github.com/AliyunContainerService/kubernetes>.
- Alibaba Cloud VPC network drive for Flannel: <https://github.com/coreos/flannel/blob/master/Documentation/alicloud-vpc-backend.md>.

If you have any questions or suggestions regarding a specific project, you are welcome to raise an issue or pull a request in the community.

1.2.2 Alibaba Cloud Kubernetes vs. self-built Kubernetes

Advantages of Alibaba Cloud Kubernetes

Easy to use

- Supports creating a Kubernetes cluster 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 1.8.6, 1.9.4, and 1.10 in the future. Upgrading clusters each time brings you great adjustments and Operation & Maintenance (O&M) costs. 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 expanding or contracting 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

Function	Description
Network	<ul style="list-style-type: none">• High-performance Virtual Private Cloud (VPC) network plug-in.• Supports network policy and flow control. <p>Container Service provides you with continuous network integration and the best network optimization.</p>
Server Load Balancer	<p>Supports creating Internet or intranet Server Load Balancer instances.</p> <p>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.</p>

Function	Description
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>
O&M	<ul style="list-style-type: none">· Integrates with Alibaba Cloud Log Service and CloudMonitor.· Supports auto scaling.
Image repository	<ul style="list-style-type: none">· High availability. Supports high concurrency.· Supports speeding up the pull of images.· Supports P2P distribution. <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 Container Service, which reduces the O&M burden and upgrade pressure.</p>

Function	Description
Stability	<ul style="list-style-type: none">• The dedicated team guarantees the stability of the container.• Each Linux version and Kubernetes version are provided to you after strict tests. <p>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 hang, network problems, and kernel compatibility, Container Service provides you with the best practices.</p>
High availability	<ul style="list-style-type: none">• Supports multiple zones.• Supports backup and disaster recovery.
Technical support	<ul style="list-style-type: none">• 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.

Costs and risks of self-built Kubernetes

- Building clusters is complicated

You must manually configure the components, configuration files, certificates, keys, plug-ins, and tools related to Kubernetes. It takes several days or weeks for professional personnel to build the cluster.

- For public cloud, it takes you significant costs to integrate with cloud products.

You must devote your own money to integrate with other products of Alibaba Cloud, such as Log Service, monitoring service, and storage management.

- The container is a systematic project, involving network, storage, operating system, orchestration, and other technologies, which requires the devotion of professional personnel.

- The container technology is continuously developing with fast version iteration, which requires continuous upgrade and test.

1.3 Authorization management

1.3.1 Role authorization

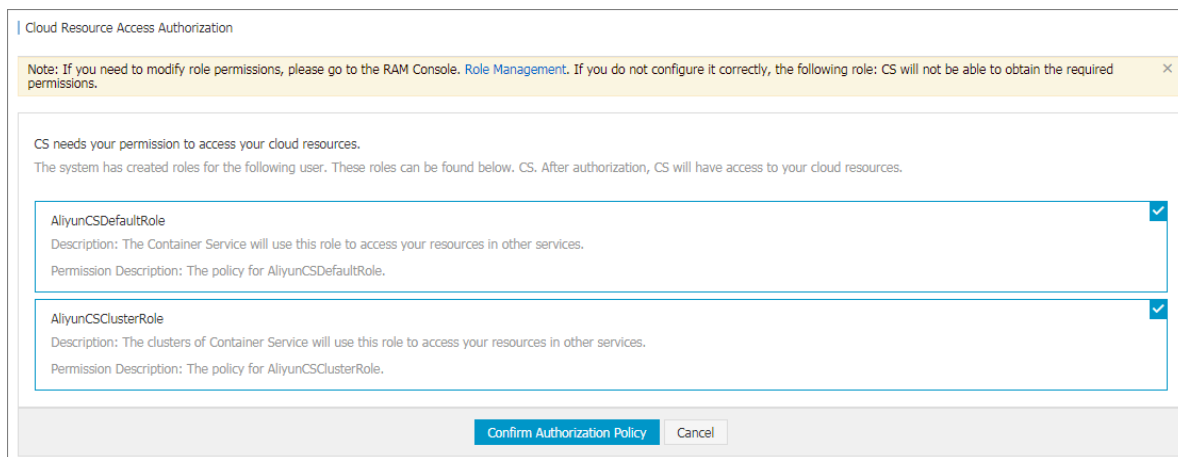
Grant the system default roles AliyunCSDefaultRole and AliyunCSClusterRole to the service account when you activate Container Service. Only after the roles are correctly granted, Container Service can normally call services such as Elastic Compute Service (ECS), Object Storage Service (OSS), Network Attached Storage (NAS), and Server Load Balancer (SLB), create clusters, and store logs.

Instructions

- If you have used Container Service before 15 January 2018, the system completes the role authorization by default. For the detailed granted permissions, see the following Default role permissions section. If you used Container Service with a Resource Access Management (RAM) user before, upgrade the authorization policy for the RAM user. For more information, see [Create custom authorization policies](#).
- On 15 January 2018, Container Service is fully accessed to the cross-service authorization. New users who use the primary account can use Container Service only after having the cross-service authorization completed. If new users need to authorize RAM users to use Container Service, go to the RAM console to authorize the RAM users. For more information, see [Use the Container Service console as a RAM user](#).

Procedure

1. If you have not granted the default roles to the service account correctly, the Cloud Resource Access Authorization page appears after you log on to the Container Service console. Click **Confirm Authorization Policy**.



Cloud Resource Access Authorization

Note: If you need to modify role permissions, please go to the RAM Console. [Role Management](#). If you do not configure it correctly, the following role: CS will not be able to obtain the required permissions. X

CS needs your permission to access your cloud resources.
The system has created roles for the following user. These roles can be found below. CS. After authorization, CS will have access to your cloud resources.

AliyunCSDefaultRole Description: The Container Service will use this role to access your resources in other services. Permission Description: The policy for AliyunCSDefaultRole.	✓
AliyunCSClusterRole Description: The clusters of Container Service will use this role to access your resources in other services. Permission Description: The policy for AliyunCSClusterRole.	✓

Confirm Authorization Policy Cancel



Note:

Container Service has configured the default role permissions. To modify the role permissions, go to the User Management page of the RAM console. Note that incorrect configurations might cause Container Service cannot obtain the required permissions.

2. After completing the authorization, refresh the Container Service console and then perform the operations.

To view the policy details of the roles AliyunCSDefaultRole and AliyunCSClusterRole, log on to the [RAM console](#).

Default role permissions

For more information about permissions of each role, see the API documents of each product.

AliyunCSDefaultRole permissions

The default role AliyunCSDefaultRole contains the following main permissions:

- ECS-related permissions

Action	Description
ecs:RunInstances	Query ECS instance information.
ecs:RenewInstance	Renew ECS instances.

Action	Description
ecs:Create*	Create ECS-related resources, such as instances and disks.
ecs:AllocatePublicIpAddress	Allocate public IP addresses.
ecs:AllocateEipAddress	Allocate Elastic IP (EIP) addresses.
ecs>Delete*	Delete ECS instances.
ecs:StartInstance	Start ECS-related resources.
ecs:StopInstance	Stop ECS instances.
ecs:RebootInstance	Restart ECS instances.
ecs:Describe*	Query ECS-related resources.
ecs:AuthorizeSecurityGroup	Configure inbound security group rules.
ecs:RevokeSecurityGroup	Revoke security group rules.
ecs:AuthorizeSecurityGroupEgress	Configure outbound security group rules.
ecs:AttachDisk	Add disks.
ecs:DetachDisk	Clean up disks.
ecs:AddTags	Add tags.
ecs:ReplaceSystemDisk	Change system disks of ECS instances.
ecs:ModifyInstanceAttribute	Modify ECS instance attributes.
ecs:JoinSecurityGroup	Add ECS instances to specified security groups.
ecs:LeaveSecurityGroup	Remove ECS instances from specified security groups.
ecs:UnassociateEipAddress	Unbind EIP addresses.
ecs:ReleaseEipAddress	Release EIP addresses.

· Virtual Private Cloud (VPC)-related permissions

Permission name (Action)	Permission description
vpc:Describe*	Query information of VPC-related resources.
vpc:DescribeVpcs	Query VPC information.
vpc:AllocateEipAddress	Allocate EIP addresses.
vpc:AssociateEipAddress	Associate with EIP addresses.

Permission name (Action)	Permission description
vpc:UnassociateEipAddress	Do not associate with EIP addresses.
vpc:ReleaseEipAddress	Release EIP addresses.
vpc:CreateRouteEntry	Create router interfaces.
vpc>DeleteRouteEntry	Delete router interfaces.

· SLB-related permissions

Action	Description
slb:Describe*	Query information related to Server Load Balancer.
slb:CreateLoadBalancer	Create Server Load Balancer instances.
slb>DeleteLoadBalancer	Delete Server Load Balancer instances.
slb:RemoveBackendServers	Unbind Server Load Balancer instances.
slb:StartLoadBalancerListener	Start specified listeners.
slb:StopLoadBalancerListener	Stop specified listeners.
slb:CreateLoadBalancerTCPLListener	Create TCP-based listening rules for Server Load Balancer instances.
slb:AddBackendServers	Add backend servers.

AliyunCSClusterRole permissions

The default role AliyunCSClusterRole contains the following main permissions:

· OSS-related permissions

Action	Description
oss: PutObject	Upload file or folder objects.
oss: GetObject	Get file or folder objects.
oss: ListObjects	Query file list information.

· NAS-related permissions

Action	Description
nas:Describe*	Return NAS-related information.
nas:CreateAccessRule	Create permission rules.

· SLB-related permissions

Action	Description
slb:Describe*	Query information related to Server Load Balancer.
slb:CreateLoadBalancer	Create Server Load Balancer instances.
slb>DeleteLoadBalancer	Delete Server Load Balancer instances.
slb:RemoveBackendServers	Unbind Server Load Balancer instances.
slb:StartLoadBalancerListener	Start specified listeners.
slb:StopLoadBalancerListener	Stop specified listeners.
slb:CreateLoadBalancerTCPLListener	Create TCP-based listening rules for Server Load Balancer instances.
slb:AddBackendServers	Add backend servers.
slb>DeleteLoadBalancerListener	Delete listening rules of Server Load Balancer instances.
slb:CreateVServerGroup	Create VServer groups and add backend servers.
slb:ModifyVServerGroupBackendServers	Change backend servers in VServer groups.
slb:CreateLoadBalancerHTTPListener	Create HTTP-based listeners for Server Load Balancer instances.
slb:SetBackendServers	Configure backend servers and set the weight for a group of ECS instances at the Server Load Balancer instance backend.
slb:AddTags	Add tags for Server Load Balancer instances.

1.3.2 Use the Container Service console as a RAM user

You can log on to and perform operations in the Container Service console as a RAM user.

Before you can log on to the Container Service console and perform operations as a RAM user, you must grant related permissions to the RAM user.

Step 1: Create a RAM user and enable console login

1. Log on to the [RAM console](#).
2. In the left-side navigation bar, click Users. Then, click Create User.

3. Enter a user name for the RAM user and then click OK.
4. On the Users page, select the created RAM user and click Manage.
5. In the Web Console Logon Management area, click Enable Console Logon.
6. Enter a logon password and click OK.

Step 2: Grant the RAM user permissions to access Container Service

1. On the Users page, select the created RAM user and click Authorize.

RAM	User Management Create User Refresh				
Dashboard	<div> <div>User Name</div> <div>Search by User Name</div> <div>Search</div> </div>				
Users					
Groups					
Policies					
Roles					
Settings					
	User Name/Display Name	Description	Created At	Actions	
	zhongguo... zhongguo...		2017-11-01 11:26:17	Manage	Authorize
	AliyunOSSTokenGeneratorUser		2017-11-27 11:55:21	Manage	Authorize
	yangxue... yangxue...		2017-11-01 11:24:54	Manage	Authorize

2. Select the required policies to attach them to the RAM user.

Edit User-Level Authorization

Members added to this group have all the permissions of this group. A member cannot be added to the same group more than once.

Available Authorization Policy Names	Type
CS	
EcsRamRoleDocument...	
AliyunACSResourcesAccess_yangx...	Custom
aliyun container s...	
AliyunCSReadOnlyAccess	System
Provides read-only...	
AliyunCSFullAccess	System
Provides full acce...	

Selected Authorization Policy Name	Type
AdministratorAccess	System
Provides full acce...	
AliyunACSResourcesAccess_xingy...	Custom
aliyun container s...	

OK

Close

You can use the following system policies:

- **AliyunCSFullAccess:** Provides full access to Container Service.
- **AliyunCSReadOnlyAccess:** Provides read-only access to Container Service.

You can also create custom policies as you need and attach them to the RAM user.

For more information, see [Create custom authorization policies](#).

Step 3: Log on to the Container Service console as a RAM user

- If you have granted the `AliyunCSDefaultRole` and `AliyunCSClusterRole` roles to the Alibaba Cloud account, you can log on to the Container Service console and perform operations as a RAM role directly.

Log on to the [Container Service console](#) as a RAM user.

- If you have not granted the `AliyunCSDefaultRole` and `AliyunCSClusterRole` roles to the Alibaba Cloud account, you must log on to the Container Service console using the account credentials and

click **Confirm Authorization Policy** on the authorization page to grant the account the following permissions.

Cloud Resource Access Authorization

Note: If you need to modify role permissions, please go to the RAM Console. [Role Management](#). If you do not configure it correctly, the following role: CS will not be able to obtain the required permissions. X

CS needs your permission to access your cloud resources.
The system has created roles for the following user. These roles can be found below. CS. After authorization, CS will have access to your cloud resources.

AliyunCSDefaultRole Description: The Container Service will use this role to access your resources in other services. Permission Description: The policy for AliyunCSDefaultRole.	✓
AliyunCSClusterRole Description: The clusters of Container Service will use this role to access your resources in other services. Permission Description: The policy for AliyunCSClusterRole.	✓

Confirm Authorization Policy Cancel

After you grant the preceding permissions to the account, you can log on to the Container Service and perform related operations as a RAM user.

1.3.3 Create custom authorization policies

The authorization granularity of the system authorization policies provided by Container Service is coarse. If these authorization policies with coarse granularity cannot satisfy your requirements, create the custom authorization policies. For example, to control the permissions to a specific cluster, you must use the custom authorization policy to meet the requirements with fine granularity.

Create custom authorization policies

Get to know the basic structure and syntax of the authorization policy language before creating custom authorization policies. For more information, see [Authorization policy language descriptions](#).

This document introduces how to grant Resource Access Management (RAM) users permissions to query, expand, and delete clusters.

Procedure

1. Log on to the [RAM console](#) with the primary account.
2. Click Policies in the left-side navigation pane. Click Create Authorization Policy in the upper-right corner.
3. Select a template. Enter the authorization policy name and the policy content.

Create Authorization Policy

Step 1: Select an authorization policy

Step 2: Edit permissions and submit.

Policy creation complete.

* Authorization Policy

Name :

clusterpolicy

Names must be 1-128 characters long. They may only contain the letters A-Z, numbers 0-9, and hyphens.

Description :

Policy Content :

```

1 {
2   "Statement": [{
3     "Action": [
4       "cs:Get*",
5       "cs:ScaleCluster",
6       "cs>DeleteCluster"
7     ],
8     "Effect": "Allow",
9     "Resource": [
10      "acs:cs:*:*:cluster/cb2f4c
11    ]
12  }],
13  "Version": "1"
14 }

```

Authorization Policy Format

Previous

Create Authorization Policy

Cancel

```

{
  "Statement": [{
    "Action": [
      "cs : Get *",
      "cs : ScaleClust er ",
      "cs : DeleteClus ter "
    ],
    "Effect": " Allow ",
    "Resource": [
      " acs : cs :*:*: cluster / cluster ID "
    ]
  }],
  "Version": " 1 "
}

```

```
}

```

where:

- **Action** : Enter the permission that you want to grant.



Note:

All the Actions support wildcards.

- **Resource** supports the following configuration methods.

- Grant permissions of a single cluster

```
" Resource ": [
  " acs : cs :*:*: cluster / cluster ID "
]
```

- Grant permissions of multiple clusters

```
" Resource ": [
  " acs : cs :*:*: cluster / cluster ID ",
  " acs : cs :*:*: cluster / cluster ID "
]
```

- Grant permissions of all your clusters

```
" Resource ": [
  "*"
]
```

You must replace `cluster ID` with your actual cluster ID.

4. Click Create Authorization Policy after completing the configurations.

Table 1-1: Container Service RAM action

Action	Description
CreateCluster	Create clusters.
AttachInstances	Add existing Elastic Compute Service (ECS) instances to clusters.
ScaleCluster	Expand clusters.
GetClusters	View cluster list.
GetClusterById	View cluster details.
ModifyClusterName	Modify cluster names.
DeleteCluster	Delete clusters.
UpgradeClusterAgent	Upgrade cluster Agent.

Action	Description
GetClusterLogs	View cluster operation logs.
GetClusterEndpoint	View cluster access point.
GetClusterCerts	Download cluster certificate.
RevokeClusterCerts	Revoke cluster certificate.
BindSLB	Bind Server Load Balancer instances to clusters.
UnBindSLB	Unbind Server Load Balancer instances from clusters.
ReBindSecurityGroup	Rebind security groups to clusters.
CheckSecurityGroup	Check existing security group rules of clusters.
FixSecurityGroup	Fix cluster security group rules.
ResetClusterNode	Reset cluster nodes.
DeleteClusterNode	Delete cluster nodes.
CreateAutoScale	Create node auto scaling rules.
UpdateAutoScale	Update node auto scaling rules.
DeleteAutoScale	Delete node auto scaling rules.
GetClusterProjects	View applications in clusters.
CreateTriggerHook	Create triggers for applications.
GetTriggerHook	View application trigger list.
RevokeTriggerHook	Delete application triggers.
CreateClusterToken	Create tokens.

1.3.4 Grant Kubernetes permissions to a RAM user

This topic describes how to configure Kubernetes Resource Access Management (RAM) cluster permissions and the corresponding Kubernetes RBAC application permissions within the cluster by using the Container Service console.

Before you begin

- Make sure that you have an Alibaba Cloud account and that you have created one or more RAM users.

- If you grant a RAM user the corresponding permissions for a Container Service Kubernetes cluster, you must use an account to complete the task if further operations in the RAM console is required.
- If a RAM user is granted the preset admin role or the custom cluster-admin role in a cluster or namespace, then the RAM user can grant permissions to other RAM users in the same cluster or namespace.
- If RAM authorization is involved in configuring permissions through the Container Service console, you must manually perform authorization in the RAM console for the target RAM user according to the reference policy and operation instructions on the page due the security restrictions of RAM.

Authorization policy upgrade notice

Container Service will upgrade the cluster authorization policy (in April 2019) to enhance the security of Kubernetes clusters. After the upgrade, any RAM users in a Kubernetes cluster that are not granted the required permissions will not be able to access the cluster resources.

Therefore, we recommend that you grant required permissions to the RAM users in each of your Kubernetes clusters. After you complete this process, your managed RAM users will only have the specified permissions to access the their corresponding authorized cluster.

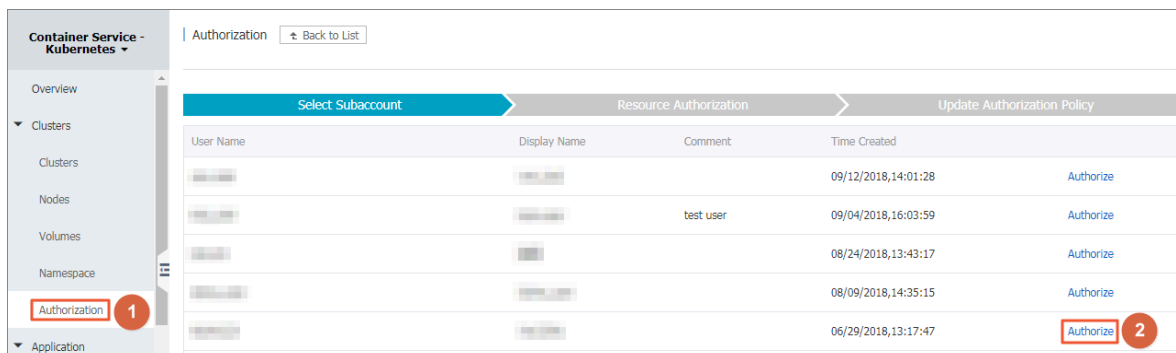
Procedure

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Container Service - Kubernetes, choose Clusters > Authorization.
3. On the right of the target RAM user, click Authorize.



Note:

If you want to use a RAM user for the preceding operations, make sure that the target RAM user has been granted the corresponding admin permissions or the cluster-admin role.



- In the upper-left corner of the Resource Authorization tab page, click the plus sign, set the cluster and namespace where the permissions to be granted apply, set a role for the RAM user, and then click Next Step.



Note:

- You can grant one preset role and multiple custom roles to a RAM user in a cluster or a namespace.

- You can click the minus sign to remove a group of permission settings (that is, the settings of the cluster, the namespace, and the role).

Authorization [Back to List](#)

Select Subaccount [Resource Authorization](#) [Update Authorization Policy](#)

Cluster/Namespace	Role
Clusters: k8s-cluster Namespace: all namespace	<input checked="" type="radio"/> Admin <input type="radio"/> Operation <input type="radio"/> Developer <input type="radio"/> Restricted User <input type="radio"/> Custom
Clusters: k8s-cluster Namespace: all namespace	<input type="radio"/> Admin <input type="radio"/> Operation <input type="radio"/> Developer <input type="radio"/> Restricted User <input checked="" type="radio"/> Custom alibaba-loc-controller
Clusters: k8s-cluster Namespace: all namespace	<input type="radio"/> Admin <input type="radio"/> Operation <input type="radio"/> Developer <input type="radio"/> Restricted User <input checked="" type="radio"/> Custom alicloud-disk-controller-runner
Clusters: k8s-cluster Namespace: default	<input type="radio"/> Developer <input checked="" type="radio"/> Restricted User <input type="radio"/> Custom

Permission description

	Cluster management permissions	Application management permissions
Admin	Read and write permissions of clusters. Can delete, scale clusters, add nodes	Read and write permissions of resources in all namespaces, Read and write permissions of nodes, volumes, namespaces and quotas
Operation	Read and write permissions of clusters. Can delete, scale clusters, add nodes	Read and write permissions of resources in all namespaces, Read only permissions of nodes, volumes, namespaces and quotas
Developer	Read only permissions of clusters	Read and write permissions of resources in all namespaces or specified namespace
Restricted User	Read only permissions of clusters	Read only permissions of resources in all namespaces or specified namespace
Custom	The permissions are determined by the specified cluster roles. Before you authorize the RAM user, make sure that you are aware of all the resource access permissions of the selected cluster roles. This will prevent the RAM user from obtaining unnecessary permissions.	

[Prev Step](#) [Next Step](#)

The role permissions that apply to clusters and namespaces are shown in the following table.

Table 1-2: Roles and permissions

Role	Cluster management permissions	Application management permissions
Admin	The read and write permissions of clusters . Specifically, the permissions to delete clusters, scale clusters, and add nodes.	The read and write permissions of resources in all namespaces, and those of nodes, volumes, namespaces, and quotas.
Operation	The read and write permissions of clusters . Specifically, the permissions to delete clusters, scale clusters, and add nodes.	The read and write permissions of resources in all namespaces, and the ready permission s of nodes, volumes, namespaces, and quotas.

Role	Cluster management permissions	Application management permissions
Developer	The read permissions of clusters.	The read and write permissions of the resources in all namespaces or specified namespaces.
Restricted User	The read permissions of clusters.	The read permissions of resources in all namespaces or specified namespaces.
Custom	The read and write permissions of clusters. Specifically, the permissions to delete clusters, scale clusters, and add nodes.	The permissions of the RAM user depend on the cluster role you select. Confirm the permissions that your selected cluster role has on resources before authorization, to avoid inappropriate permissions granted to the RAM user.

5. If Authorized success is displayed, this means that you have granted the target RAM user the permissions. If Add authorization policies is displayed, follow these steps:

- If you are using a RAM user, you need to log out, and then use an account to log on to the console to grant required permissions to the target RAM user.
- If you are using a primary account, you can directly perform the next step.

6. Click Copy, and then click Policy Management.

Container Service - Kubernetes

Authorization [Back to List](#)

Select Subaccount > Resource Authorization > Update Authorization Policy

Add authorization policies

To keep your account secure, you need to manually add authorization policies to the sub-account by accessing the RAM console.

- 1 Open [Policy Management](#) to add or update custom policies starting with AliyunACSRessourcesAccess-mengguoqiang. The policy name and content are as follows:

Policy Name	Policy Content
AliyunACSRessourcesAccess-mengguoqiang	<pre>{ "Statement": [{ "Action": ["ram:Get*", "ram:List*", "cs:GetUserPermissions", "cs:GetSubUsers", "cs:GrantPermission"], "Resource": "*" }] }</pre>

[Copy](#)

- 2 Open [User Management](#), and perform the "authorization" operation on the sub-account mengguoqiang to authorize the policy created above.
- 3 Authorized success

7. Choose Permissions > Policies, and then click Create Policy.

RAM / Policies

Policies

A policy describes a permission set. Alibaba Cloud uses a simple language specification to describe permission sets. For more information, see [Policy syntax structure](#).

RAM supports two types of policies: system policies managed by Alibaba Cloud and custom policies managed by you.

- You can use but cannot modify the system policies managed by Alibaba Cloud. Alibaba Cloud maintains and updates the system policy versions.
- You can create, modify, or delete the custom policies. In addition, you need to maintain the policy versions by yourself.

[Create Policy](#) Enter a policy name or note Policy Type

Policy Name	Note	Policy Type	Used Times	Actions
...	Provides full access to Alibaba Cloud services and resources.	System Policy	9	
...	Provides full access to Object Storage Service(OSS) via Management Console.	System Policy	0	

8. Enter a Policy Name, select the Script configuration mode, use the hot key Ctrl+V to paste the content that was copied in step 6 in the Policy Document area, and then click OK. For more information, see [Create custom authorization policies](#).

RAM / Policies / Create Custom Policy

← Create Custom Policy

Policy Name
AliyunCS-001

Note

Configuration Mode
☐ Visualized
☒ Script

Policy Document

Import an existing system policy

```

1  {
2    "Statement": [
3      {
4        "Action": [
5          "ram:Get*",
6          "ram:List*",
7          "cs:GetUserPermissions",
8          "cs:GetSubUsers",
9          "cs:GrantPermission"
10         ],
11        "Resource": "*"
12      }
13    ]
14  }
  
```

OK Back

9. In the left-side navigation pane, choose Identities > Users. On the right of the target user, click Add Permissions.

RAM / Users

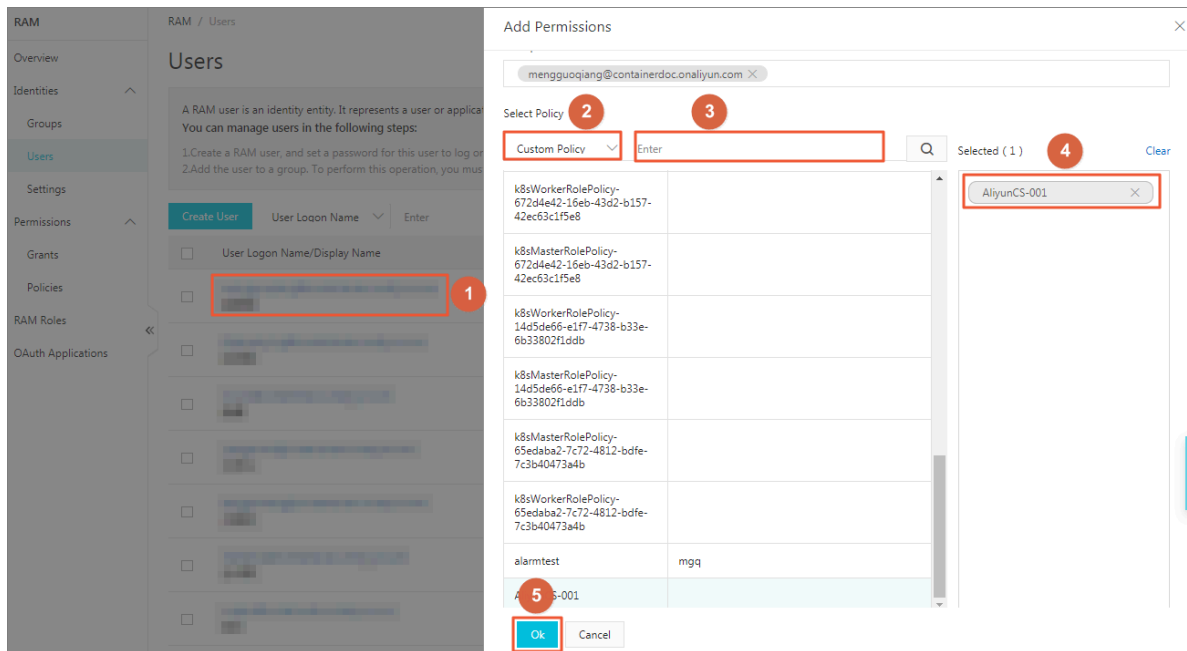
Users

A RAM user is an identity entity. It represents a user or application in your organization that needs to access cloud resources.
 You can manage users in the following steps:
 1. Create a RAM user, and set a password for this user to log on to the console or create an AccessKey for the application to call APIs.
 2. Add the user to a group. To perform this operation, you must have created a group and granted permissions to it.

Create User User Logon Name ▾ Enter 🔍

	User Logon Name/Display Name	Note	Created	Actions
<input type="checkbox"/>	[Redacted]		Feb 18, 2019, 16:03:11	Add to Group Add Permissions Delete
<input type="checkbox"/>	[Redacted]		Feb 15, 2019, 14:21:00	Add to Group Add Permissions Delete
<input type="checkbox"/>	[Redacted]		Jan 30, 2019, 17:09:53	Add to Group Add Permissions Delete
<input type="checkbox"/>	[Redacted]		Jan 22, 2019, 13:06:46	Add to Group Add Permissions Delete

10. Select Custom Policy, search for or manually look for the customized policy, click the policy name to add the policy to the Selected area on the right, and then click OK.



Note:

After you complete these steps, you can use the target RAM user to log on to the Container Service console and perform the operations allowed by the granted permissions. Typically, if the target RAM user is granted the admin permissions, you can use it to grant permissions to other RAM users.

Custom permissions

Alibaba Cloud Container Service offers four types of permissions by pre-setting four types of roles: Admin, Operation, Developer, and Restricted User. These types of permissions can meet the needs of most users in the Container Service console. However, if you want to customize the access permissions to clusters, you can also use the custom permissions.



Note:

Alibaba Cloud Container Service provides several custom permission. Among them, the cluster-admin permission is a super administrator permission with the permissions to access and operate on all resources.

Authorization [Back to List](#)

Select Subaccount **Resource Authorization** Update Authorization Policy

Cluster/Namespace Role

Clusters Namespace

☐ Admin
 ☐ Operation
 ☐ Developer
 ☐ Restricted User
 ☒ Custom

admin

admin

alibaba-log-controller

alicloud-disk-controller-runner

cluster-admin

edit

flannel

view

Permission description

	Cluster management permissions	Application management permissions
Admin	Read and write permissions of clusters. Can delete, scale clusters, add nodes	Read and write permissions of resources in all namespaces, Read and write permissions of nodes, volumes, namespaces and quotas
Operation	Read and write permissions of clusters. Can delete, scale clusters, add nodes	Read and write permissions of resources in all namespaces, Read only permissions of nodes, volumes, namespaces and quotas
Developer	Read only permissions of clusters	Read and write permissions of resources in all namespaces or specified namespace
Restricted User	Read only permissions of clusters	Read only permissions of resources in all namespaces or specified namespace
Custom	The permissions are determined by the specified cluster roles. Before you authorize the RAM user, make sure that you are aware of all the resource access permissions of the selected cluster roles. This will prevent the RAM user from obtaining unnecessary permissions.	

Prev Step Next Step

You can log on to the cluster Master node and run the following command to view the details of the custom permissions.



Note:

Only some of the cluster roles are displayed.

```
# kubectl get clusterrole
NAME
AGE
admin
13d
alibaba-log-controller
13d
alicloud-disk-controller-runner
13d
cluster-admin
13d
cs:admin
13d
edit
13d
flannel
13d
kube-state-metrics
22h
node-exporter
22h
prometheus-k8s
22h
prometheus-operator
22h
system:aggregate-to-admin
13d
```

```
....
system : volume - scheduler
      13d
view
      13d
```

To view the permission details of the super administrator cluster-admin, run the following command.



Note:

After the RAM user is granted the cluster-admin role, the RAM user can be regarded as a super administrator that has the same privileges as the Alibaba Cloud account, and it can perform operations on any resources in the cluster. Execute caution when you grant the cluster-admin role.

```
# kubectl get clusterrole cluster-admin -o yaml
apiVersion : rbac.authorization.k8s.io/v1
kind : ClusterRole
metadata :
  annotations :
    rbac.authorization.kubernetes.io/autoupdate : "true"
  creationTimestamp : 2018-10-12T08:31:15Z
  labels :
    kubernetes.io/bootstrapping : rbac-defaults
  name : cluster-admin
  resourceVersion : "57"
  selfLink : /apis/rbac.authorization.k8s.io/v1/clusterroles/cluster-admin
  uid : 2f29f9c5-cdf9-11e8-84bf-00163e0b2f97
rules :
- apiGroups :
  - '*'
  resources :
  - '*'
  verbs :
  - '*'
- nonResourceURLs :
  - '*'
  verbs :
  - '*'
```

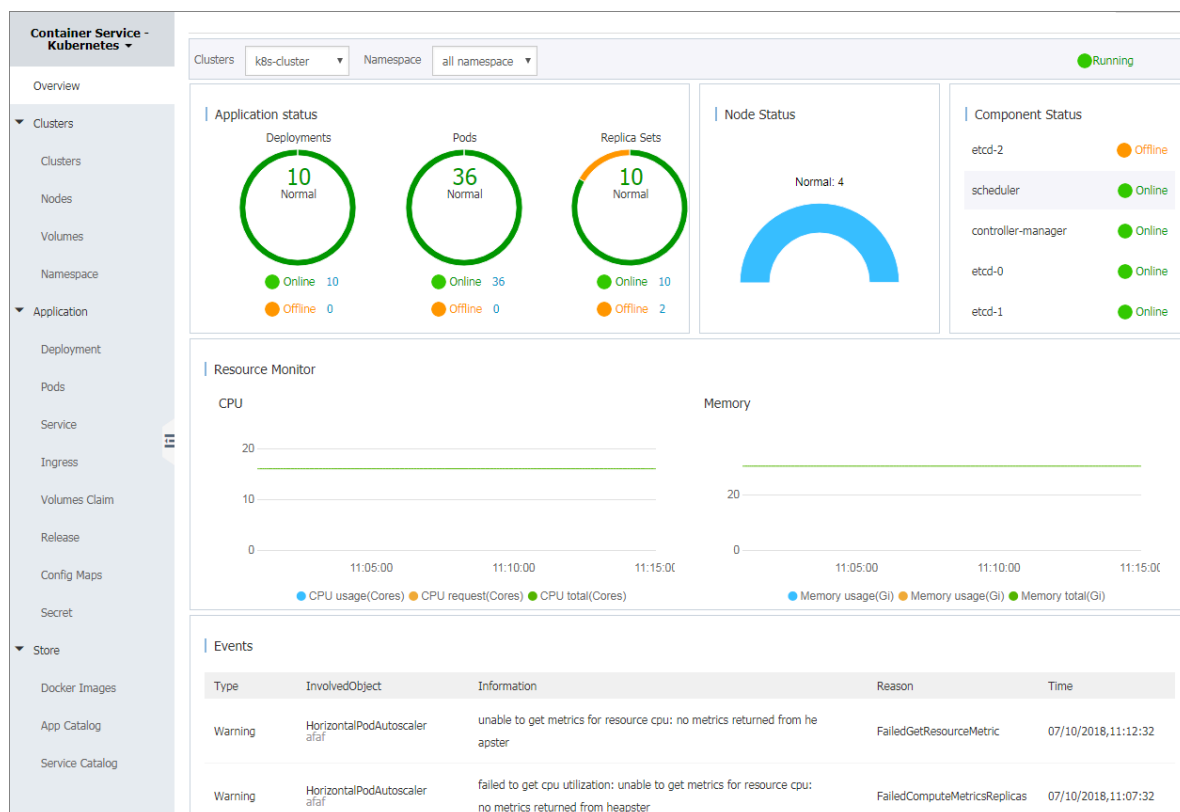
1.4 Cluster management

1.4.1 View cluster overview

This topic describes how to view the health status of a Kubernetes cluster. Alibaba Cloud Container Service for Kubernetes provides each Kubernetes cluster with the application status, component status, and resource monitoring charts.

Procedure

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Kubernetes, click Overview.
3. Select the target cluster and namespace. You can view the application status, component status, and resource monitoring charts.
 - **Application status:** The status of deployments, pods, and replica sets that are currently running. Green indicates the normal status and orange indicates an exception.
 - **Node status:** Displays the node status of the current cluster.
 - **Component status:** The components of Kubernetes clusters are generally deployed under the kube-system namespace, including the core components such as scheduler, controller-manager, and etcd.
 - **Resource monitor:** Provides the monitoring charts of CPU and memory. CPU is measured in cores and is accurate to three decimal places. The minimum unit is millicores, that is, one thousandth of one core. Memory is measured in G and is accurate to three decimal places. For more information, see [Meaning of CPU](#) and [Meaning of memory](#).
 - **Event:** Displays event information of the cluster, such as warnings and error events.



1.4.2 Create a Kubernetes cluster

This topic describes how to quickly create a Kubernetes cluster in the Container Service console.

Detail analysis

During cluster creation, Container Service performs the following operations:

- Creates Elastic Compute Service (ECS) instances, sets the public key used for SSH logon from the management node to other nodes, and installs and configures the Kubernetes cluster by using cloud-init.
- Creates a security group to allow inbound access to all ICMP ports in a Virtual Private Cloud (VPC).
- Creates a new VPC and VSwitch (if no existing VPC is selected), and creates an SNAT entry for the VSwitch.
- Creates VPC routing rules.
- Creates a NAT gateway and a shared bandwidth package or Elastic IP (EIP).
- Creates a Resource Access Management (RAM) user and an AccessKey. The RAM user has the permissions to query, create, and delete an ECS instance, add and delete a cloud disk, and all permissions to perform required actions on Server Load Balancer (SLB) instances, CloudMonitor, VPC, Log Service, and the Network Attached Storage (NAS) service. The Kubernetes cluster dynamically creates SLB instances, cloud disks, and VPC routing rules according to your configurations.
- Creates an intranet SLB instance and allows access to port 6443.
- Creates an Internet SLB instance and allows access to port 6443. (If you enable the SSH logon for Internet access when creating the cluster, port 22 is opened. Otherwise, port 22 is closed.)

Prerequisites

You have activated the following services: Container Service, Resource Orchestration Service (ROS), RAM, and Auto Scaling service.

If you have not activated one or some of these services, log on to the [Container Service console](#), [ROS console](#), [RAM console](#), and [Auto Scaling console](#) to activate the corresponding service or services.



Note:

The deployment of Kubernetes clusters in Alibaba Cloud Container Service for Kubernetes is dependent on the application deployment capabilities of Alibaba Cloud ROS. Therefore, you need to activate ROS before creating a Kubernetes cluster.

Limits

- The SLB instances created with the cluster support only the Pay-As-You-Go billing method.
- The Kubernetes cluster supports only the VPC network type.
- By default, each account has a set quota of cloud resources that can be used to create clusters. If you exceed this number when creating a cluster, the cluster creation fails. However, you can open a ticket to increase your quota of cloud resources.
 - By default, each account can create up to 5 clusters across all regions and add up to 40 nodes to each cluster. You can open a ticket to increase your quota to create more clusters or nodes.



Note:

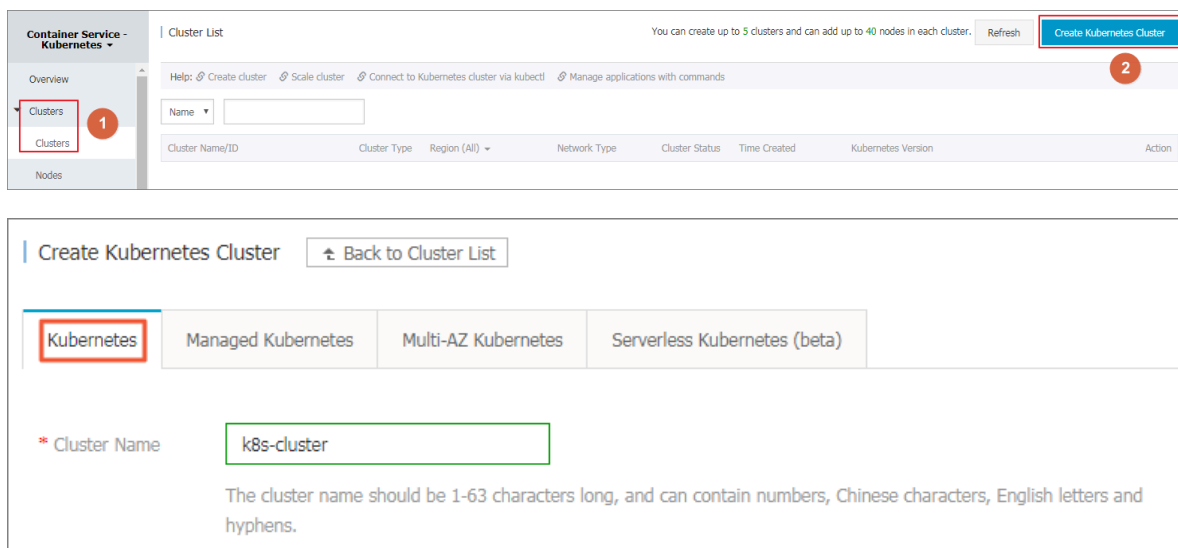
In a Kubernetes cluster, the maximum number of default routes a VPC can have is 48, which means the number of nodes the Kubernetes cluster can use is also a maximum of 48. To increase the number of nodes, you need to first open a ticket for the target VPC to increase the number of VPC routes, and then open a ticket for Container Service to apply the increase.

- By default, each account can create up to 100 security groups.
- By default, each account can create up to 60 Pay-As-You-Go SLB instances.
- By default, each account can create up to 20 EIPs.
- The limits for ECS instances are as follows:
 - Only the CentOS operating system is supported.
 - Both Pay-As-You-Go and Subscription ECS instances can be created.

Procedure

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

3. In the upper-right corner, click Create Kubernetes Cluster.



The screenshot shows the 'Cluster List' page in the Container Service - Kubernetes console. In the upper right corner, there is a 'Create Kubernetes Cluster' button, which is highlighted with a red box and a red circle with the number 2. On the left sidebar, the 'Clusters' tab is highlighted with a red box and a red circle with the number 1. Below the sidebar, there is a 'Create Kubernetes Cluster' section with a 'Back to Cluster List' link. Under this section, there are four tabs: 'Kubernetes' (highlighted with a red box), 'Managed Kubernetes', 'Multi-AZ Kubernetes', and 'Serverless Kubernetes (beta)'. Below the tabs, there is a form for creating a cluster. The 'Cluster Name' field is labeled with a red asterisk and contains the text 'k8s-cluster'. Below the field, there is a note: 'The cluster name should be 1-63 characters long, and can contain numbers, Chinese characters, English letters and hyphens.'

4. Enter the cluster name.

The cluster name must be 1 to 63 characters in length and can contain letters, numbers, Chinese characters, letters, and hyphens (-).

5. Select the region and the zone where you want to locate the cluster.



The screenshot shows the region and zone selection interface. The 'Region' dropdown is set to 'China North 2 (Beijing)'. The 'Zone' dropdown is set to 'China North 2 Zone A'. Below the dropdowns, there is a table of available regions and zones.

Region	China North 2 (Beijing)	China North 3 (Zhangjiakou)	China East 1 (Hangzhou)	China East 2 (Shanghai)	China South 1 (Shenzhen)	Hong Kong	Asia Pacific SE 1 (Singapore)	Asia Pacific SE 3 (Kuala Lumpur)	Asia Pacific SE 5 (Jakarta)	Asia Pacific SOU 1 (Mumbai)
US West 1 (Silicon Valley)	US East 1 (Virginia)	Middle East 1 (Dubai)	EU Central 1 (Frankfurt)							

6. Set the cluster network type. Kubernetes clusters support only the VPC network type.

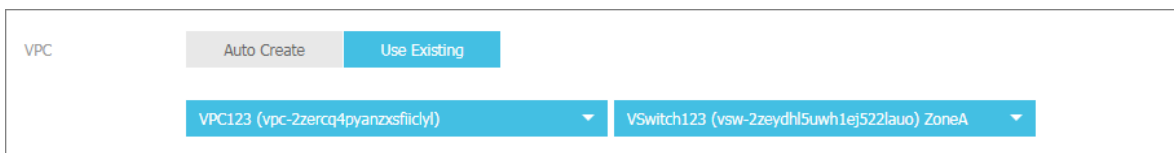
VPC: You can click Auto Create to create a VPC together with the Kubernetes cluster, or click Use Existing to use an existing VPC. If you click Use Existing, you can select a VPC and a VSwitch from the two displayed drop-down lists.

- If you click Auto Create, the system automatically creates a NAT gateway for your VPC when creating the cluster.
- If you click Use Existing and the selected VPC has an existing NAT gateway, the system then uses the already created NAT gateway. Otherwise, the system automatically creates a NAT gateway by default. Alternatively, if you do not want the system to automatically create a NAT gateway, deselect the Configure SNAT for VPC check box.



Note:

If you set the system to not automatically create a NAT gateway, you need to manually set a NAT gateway, or set an SNAT entry to ensure that your selected VPC is accessible to the Internet. Otherwise, instances in the VPC cannot access the Internet, which result in a cluster creation failure.



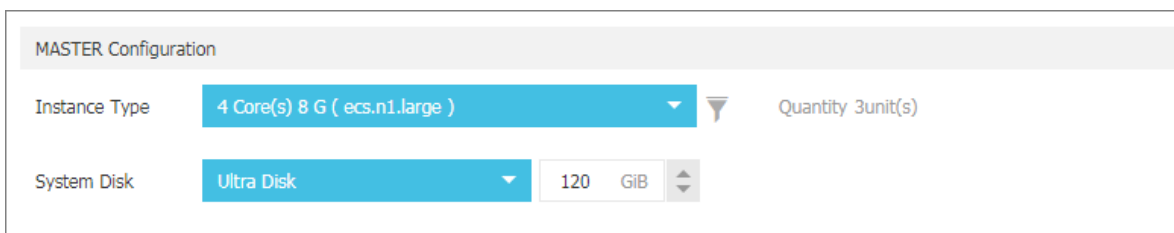
VPC configuration interface showing two buttons: 'Auto Create' and 'Use Existing'. Below the buttons are two dropdown menus. The first dropdown is labeled 'VPC123 (vpc-2zercq4pyanzsficlyf)' and the second dropdown is labeled 'VSwitch123 (vsw-2zeydh15uwh1ej522lauo) ZoneA'.

7. Set the required billing method for the node type. Pay-As-You-Go and Subscription are supported.
8. Set the Master node configuration by selecting the Master node instance type and the system disk.



Note:

- Only the CentOS operating system is supported.
- By default, the system creates three Master nodes for each cluster. This number cannot be modified.
- A system disk is attached to Master nodes by default. Available system disks are SSD disks and Ultra disks.



MASTER Configuration interface showing two rows. The first row is 'Instance Type' with a dropdown menu showing '4 Core(s) 8 G (ecs.n1.large)' and a 'Quantity 3unit(s)' label. The second row is 'System Disk' with a dropdown menu showing 'Ultra Disk' and a '120 GiB' label.

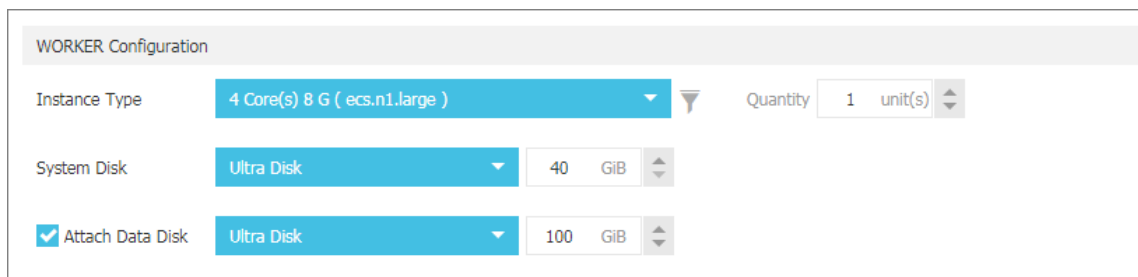
9. Set Worker nodes. You can create Worker node instances or add existing Worker node instances.



Note:

- Only the CentOS operating system is supported.
- Each cluster can contain up to 37 Worker nodes. To create more Worker nodes, open a ticket.
- A system disk is attached to Worker nodes by default. Available system disks are SSD disks and Ultra disks.

- You can manually attach a data disk to Worker nodes. Available data disks are SSD disks, Ultra disks, and basic disks.
- To create Worker node instances, select the instance type and set the number of Worker nodes. In this example, one Worker node is created.



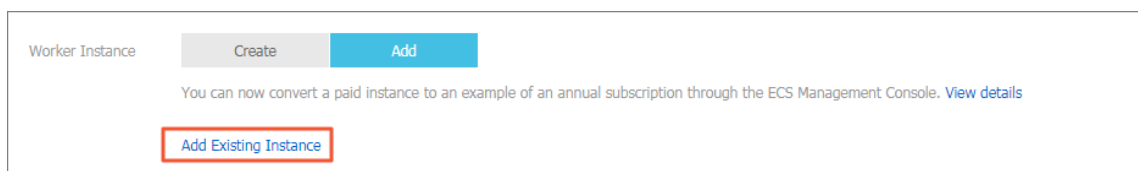
WORKER Configuration

Instance Type: 4 Core(s) 8 G (ecs.n1.large) Quantity: 1 unit(s)

System Disk: Ultra Disk 40 GiB

☒ Attach Data Disk: Ultra Disk 100 GiB

- To add existing Worker node instances, you must create ECS instances in the target region in advance.



Worker Instance

Create Add

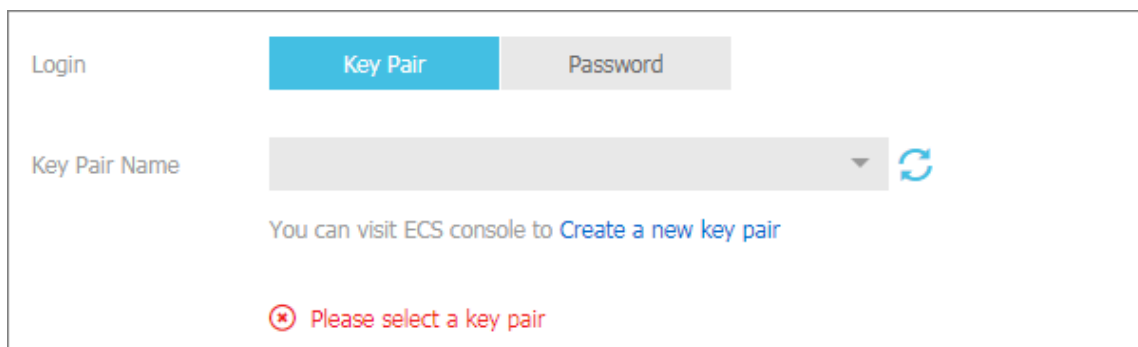
You can now convert a paid instance to an example of an annual subscription through the ECS Management Console. [View details](#)

Add Existing Instance

10.Set the logon mode.

- Set a key pair.

Click Create a new key pair to create a key pair in the ECS console, and then set the key pair as the credentials for logging on to the cluster. For more information, see [Create an SSH key pair](#).



Login

Key Pair Password

Key Pair Name

You can visit ECS console to [Create a new key pair](#)

Please select a key pair

- Set a password.
 - Logon Password: Set the node logon password.
 - Confirm Password: Confirm your node logon password.

11.Set the Pod Network CIDR and Service CIDR.



Note:

You need to set these two parameters only if you choose to use an existing VPC.

The two Classless Inter-Domain Routing (CIDR) blocks cannot overlap with each other, or with your selected VPC, or with the CIDR blocks used by any other existing Kubernetes clusters in the selected VPC. Furthermore, they cannot be modified after the cluster is created. For more information, see [Plan Kubernetes CIDR blocks under VPC](#).

12. Select a Kubernetes version as needed. The system displays the Kubernetes version and the Docker versions.

Docker Version	17.06.2-ce-3
Kubernetes Version	<div>1.11.2</div> <div>1.10.4</div>

13. Configure an SNAT gateway for the VPC.



Note:

We recommend that you configure this feature. Otherwise, instances in the VPC cannot access the Internet, which will result in a cluster creation failure.

- If you set the system to automatically create a VPC, you must configure an SNAT gateway for the VPC.
- If you set the system to use an existing VPC, you can perform either of the following operations:
 - Set the system to automatically configure an SNAT gateway.
 - Manually configure a NAT gateway for the VPC or manually configure an SNAT gateway for the VPC.

Configure SNAT	<input checked="" type="checkbox"/> Configure SNAT for VPC
If the VPC you choose does not have access to Internet, NAT gateway and EIP will be used to configure SNAT for the VPC. During this period, NAT gateway, EIP, and other resources may be created.	

14.Enable the Use Public SLB to Expose API Server function.

The API server provides add, delete, edit, check, watch, and other HTTP REST interfaces for resource objects, such as pods and services.

- If you enable this function by selecting the check box, an Internet SLB is created and the Master node port (namely, port 6443) is opened. In this case, you can use kubeconfig to connect to and operate the cluster through the Internet.



Note:

The API server uses the Master node port.

- If you do not enable this function, no Internet SLB is created. In this case, you can only use kubeconfig to connect to and operate the cluster within the VPC.

Public SLB

☒ Expose API SERVER with public SLB

When the selection is not open, the cluster API SERVER can not be accessed out of the VPC.

15.Enable the Internet SSH login.



Note:

To enable this function, you must enable the Use Public SLB to Expose API Server function.

- If you enable this function, you can use SSH to access the cluster.
- If you do not enable this function, the cluster cannot be accessed by using SSH or kubectl. To access the cluster by using SSH in such a case, you must manually associate an EIP to the ECS instance, set a security group rule, and open the SSH port (port 22). For more information, see [Access a Kubernetes cluster by using SSH](#).

SSH Login

☒ Enable SSH access for Internet

If this function is disabled, see [SSH access to Kubernetes clusters](#) to manually enable SSH access.

16.Install cloud monitoring plug-ins on your ECS node instances.

If you install cloud monitoring plug-ins on the ECS nodes, you can view the ECS instance monitoring information in the CloudMonitor console.

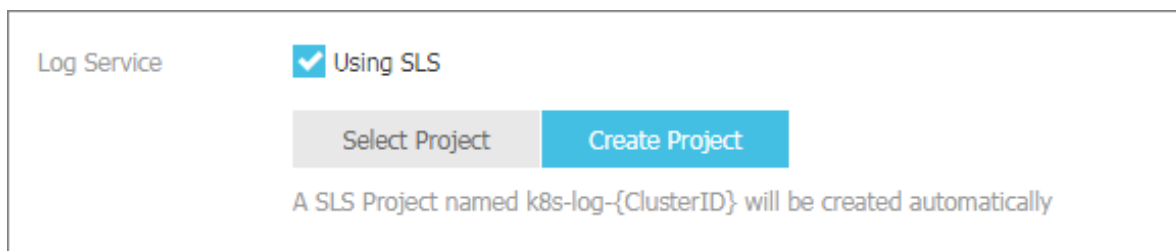
Monitoring Plug-in

☒ Install cloud monitoring plug-in on your ECS.

Installing a cloud monitoring plug-in on the node allows you to view the monitoring information of the created ECS instance in the CloudMonitor console

17.Enable Log Service. You can select an existing project or create a new project.

If you select the Using SLS check box, a Log Service plugin is automatically installed in the cluster. Then, when you create an application in the cluster, you can immediately use Log Service with only a few configurations required. For more information, see [Use Log Service to collect Kubernetes cluster logs](#).

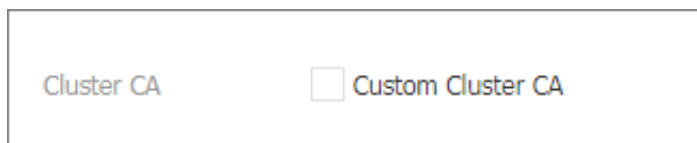


18.Set advanced configurations.

- Select a network plugin. Available network plugins are Flannel and Terway. For more information, see [Do I select the Terway or Flannel plugin for my Kubernetes cluster network?](#).
- Set the number of pods for a node. This parameter specifies the maximum number of pods that can be run by a single node. We recommend that you retain the default setting.



- Enable the Custom Cluster CA function. A CA certificate ensures secure information exchanges between the server and the client. To enable this function, select the check box to add the CA certificate to the Kubernetes cluster.



19.Click Create in the upper-right corner.

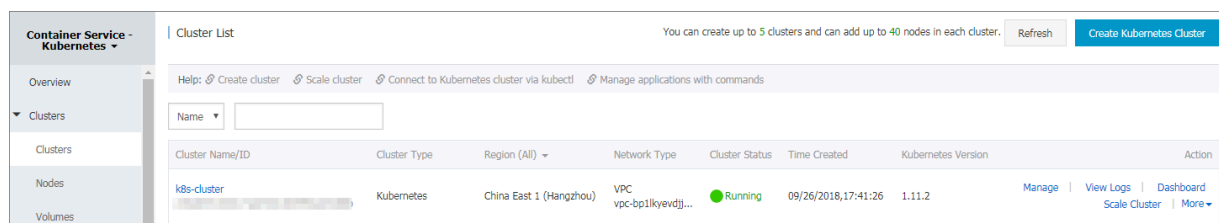


Note:

A Kubernetes cluster that contains multiple nodes typically takes ten minutes to be created.

View the cluster deployment result

After the cluster is created, you can view the cluster in the cluster list of the Container Service console.



- To view the cluster logs, click **View Logs** on the right of the cluster. To view more details, click **Stack Events**.

Detailed resource deployment logs: Stack Events	
Time	Information
06/22/2018,11:13:50	c06eadb6387ec4c6080d495e243649a7b Start to DescribeK8sUserCertConfig
06/22/2018,11:03:39	c06eadb6387ec4c6080d495e243649a7b Set up k8s DNS configuration successfully
06/22/2018,11:02:30	c06eadb6387ec4c6080d495e243649a7b Stack CREATE completed successfully;o
06/22/2018,11:02:30	c06eadb6387ec4c6080d495e243649a7b Start describeStackInfo
06/22/2018,11:02:29	c06eadb6387ec4c6080d495e243649a7b Start describeStackInfo
06/22/2018,10:46:55	c06eadb6387ec4c6080d495e243649a7b Successfully to CreateStack
06/22/2018,10:46:55	c06eadb6387ec4c6080d495e243649a7b Start to wait stack ready
06/22/2018,10:46:53	c06eadb6387ec4c6080d495e243649a7b Start to create cluster task
06/22/2018,10:46:53	c06eadb6387ec4c6080d495e243649a7b Start to CreateK8sCluster
06/22/2018,10:46:53	c06eadb6387ec4c6080d495e243649a7b Start to CreateStack
06/22/2018,10:46:50	c06eadb6387ec4c6080d495e243649a7b Start to validateCIDR
06/22/2018,10:46:41	c06eadb6387ec4c6080d495e243649a7b Start create cluster certificate

- To view the basic information, the connection information, and other information, return to the cluster list page and click Manage in the action column of the cluster.

Basic Information			
Cluster ID:	VPC	Running	Region: China East 1 (Hangzhou)
Connection Information			
API Server Internet endpoint	https://:6443		
API Server Intranet endpoint	https://:6443		
Master node SSH IP address			
Service Access Domain	n-hangzhou.alicontainer.com		
Cluster resource			
ROS	k8s-for-cs		
Internet SLB	lb-		
VPC	vpc-		
NAT Gateway	ngw-		
Connect to Kubernetes cluster via kubectl <div>1. Download the latest kubectl client from the Kubernetes Edition page . 2. Install and set up the kubectl client. For more information, see Installing and Setting Up kubectl 3. Configure the cluster credentials:</div> <div><div>KubeConfig</div><div>SSH</div></div>			

Cluster information

- **API Server Internet endpoint:** the IP address and the port through which the Kubernetes API server provides services to the Internet. It enables you to manage the cluster by using kubectl or other tools on your terminal.
- **API Server Intranet endpoint:** the IP address and the port through which the Kubernetes API server provides services within the cluster. This IP address is

the SLB instance IP address. The three Master nodes on its backend provide services.

- Master node SSH IP address: You can directly log on to Master nodes by using SSH to maintain the cluster.
- Service Access Domain: provides the services in the cluster with access domain name for testing. The service access domain name has the following suffix: `< cluster_id >.< region_id >.alicontainer.com`.

For example, you can log on to the Master node by using SSH, and run the `kubectl get node` command to view the cluster nodes.

```
login as: root
root@iZbp1d7yvpa3j183u0url1Z's password:

Welcome to Alibaba Cloud Elastic Compute Service !

[root@iZbp1d7yvpa3j183u0url1Z ~]# kubectl get node
NAME                                STATUS    ROLES    AGE     VERSION
cn-hangzhou.i-XXXXXXXXXXXX          Ready    <none>    17m     v1.8.4
cn-hangzhou.i-XXXXXXXXXXXX          Ready    master    19m     v1.8.4
cn-hangzhou.i-XXXXXXXXXXXX          Ready    master    24m     v1.8.4
cn-hangzhou.i-XXXXXXXXXXXX          Ready    master    22m     v1.8.4
[root@iZbp1d7yvpa3j183u0url1Z ~]#
```

In this example, the system displays four nodes, which are the three Master nodes and one Worker node.

1.4.3 Configure a Kubernetes GPU cluster to support GPU scheduling

From version 1.8, Kubernetes will support hardware acceleration devices such as NVIDIA GPU, InfiniBand, and FPGA, by using [device plugins](#). Furthermore, GPU solutions of Kubernetes open source communities will be deprecated in version 1.10, and removed from the master code in version 1.11.

We recommend that you use an Alibaba Cloud Kubernetes cluster combined with GPU to run highly dense computational tasks such as machine learning and image processing. With this method, you can implement one-click deployment, elastic scaling, and other functions, without needing to install NVIDIA drivers or Compute Unified Device Architecture (CUDA) beforehand.

Background information

During cluster creation, Container Service performs the following operations:

- Creates Elastic Compute Service (ECS) instances, sets the public key used for SSH logon from the management node to other nodes, and installs and configures the Kubernetes cluster by using CloudInit.
- Creates a security group to allow inbound access to all ICMP ports in a VPC.
- Creates a new VPC and VSwitch if you do not use the existing VPC, and also creates an SNAT entry for the VSwitch.
- Creates VPC routing rules.
- Creates a NAT gateway and Elastic IP (EIP).
- Creates a Resource Access Management (RAM) user and AccessKey (AK). This RAM user has the permissions to query, create, and delete ECS instances, add and delete cloud disks, and all relevant access permissions for Server Load Balancer (SLB) instances, CloudMonitor, VPC, Log Service, and Network Attached Storage (NAS) services. The Kubernetes cluster dynamically creates the SLB instances, cloud disks, and VPC routing rules according to your configurations.
- Creates an intranet SLB instance and exposes port 6443.
- Creates an Internet SLB instance and exposes ports 6443, 8443, and 22. (If you enable the SSH logon for Internet access when creating the cluster, port 22 is exposed. Otherwise, port 22 is not exposed.)

Prerequisites

You have activated Container Service, Resource Orchestration Service (ROS), and RAM.

You have logged on to the [Container Service console](#), [ROS console](#), and [RAM console](#) to activate the corresponding services.



Note:

The deployment of Container Service Kubernetes clusters depends on the application deployment capabilities of Alibaba Cloud ROS. Therefore, you need to activate ROS before creating a Kubernetes cluster.

Limits

- The SLB instance created with the Kubernetes cluster only supports the Pay-As-You-Go billing method.
- The Kubernetes cluster supports only Virtual Private Cloud (VPC).

- By default, each account has a specified quota of the number of cloud resources that it can create. If the number of cloud resources has reached the quota limit, the account cannot create a cluster. Make sure you have sufficient resource quota to create a cluster. You can open a ticket to increase your quota.
- By default, each account can create up to 5 clusters across all regions and add up to 40 nodes to each cluster. You can open a ticket to create more clusters or nodes.
- By default, each account can create up to 100 security groups.
- By default, each account can create up to 60 Pay-As-You-Go SLB instances.
- By default, each account can create up to 20 EIPs.
- The limits for ECS instances are as follows:
 - Only the CentOS operating system is supported.
 - Only Pay-As-You-Go ECS instances can be created.

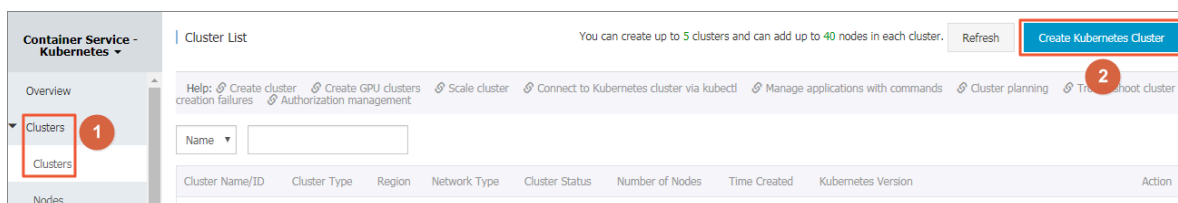


Note:

After creating an instance, you can [Switch from Pay-As-You-Go to Subscription billing](#) in the ECS console.

Create a GN5 Kubernetes cluster

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Kubernetes, click Clusters.
3. Click Create Kubernetes Cluster in the upper-right corner.



By default, the Create Kubernetes Cluster page is displayed.



Note:

Worker nodes are set to use GPU ECS instances to create a GPU cluster. For information about other parameter settings, see [Create a Kubernetes cluster](#).

4. Set the Worker nodes. In this example, the gn5 GPU instance type is selected to set Worker nodes as GPU working nodes.

a. If you choose to create Worker instances, you must select the instance type and the number of Worker nodes. In this example, two GPU nodes are created.

b. If you choose to add existing instances, you need to have already created GPU cloud servers in the same region where the cluster is to be created.

5. After you have completed all required settings, click Create to start cluster deployment.

6. After the cluster is created, choose Clusters > Nodes in the left-side navigation pane.

7. To view the GPU devices mounted to either of the created nodes, select the created cluster from the clusters drop-down list, select one of the created Worker nodes, and choose More > Details in the action column.

Create a GPU experimental environment to run TensorFlow

Jupyter is a popular tool used by data scientists for the experimental environment TensorFlow. This topic describes an example of how to deploy a Jupyter application.

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Kubernetes, choose Applications > Deployments.

3. Click Create by Template in the upper-right corner.
4. Select the target cluster and namespace and then select a sample template or the custom template from the resource type drop-down list. After you orchestrate your template, click DEPLOY.

Deploy templates

Only Kubernetes versions 1.8.4 and above are supported. For clusters of version 1.8.1, you can perform "upgrade cluster" operation in the cluster list

Clusters: xuntest2

Namespace: default

Resource Type: Custom

Template

```

1 ---
2 # Define the tensorflow deployment
3 apiVersion: apps/v1
4 kind: Deployment
5 metadata:
6   name: tf-notebook
7   labels:
8     app: tf-notebook
9 spec:
10   replicas: 1
11   selector: # define how the deployment finds the pods it manages
12     matchLabels:
13       app: tf-notebook
14   template: # define the pods specifications
15     metadata:
16       labels:
17         app: tf-notebook
18     spec:
19       containers:
20         - name: tf-notebook
21           image: tensorflow/tensorflow:1.4.1-gpu-py3
22           resources:
23             limits:
24               nvidia.com/gpu: 1
25           ports:
26             - containerPort: 8888
27               hostPort: 8888
28           env:
29             - name: PASSWORD

```

Add Deployment

Deploy with exist template

Save Template DEPLOY

In this example, a Jupyter application template is orchestrated. The template includes a deployment and a service.

```

---
# Define the tensorflow deployment
apiVersion : apps / v1
kind : Deployment
metadata :
  name : tf - notebook
  labels :
    app : tf - notebook
spec :
  replicas : 1
  selector : # define how the deployment finds the
pods it manages
  matchLabels :
    app : tf - notebook
  template : # define the pods specifications
    metadata :
      labels :
        app : tf - notebook
    spec :
      containers :
        - name : tf - notebook
          image : tensorflow / tensorflow : 1 . 4 . 1 - gpu - py3
          resources :
            limits :

```

```

      nvidia . com / gpu : 1 # specify
the number of NVIDIA GPUs that are called by
the application
  ports :
    - containerPort : 8888
      hostPort : 8888
    env :
      - name : PASSWORD # specify
the password used to access the Jupyter service .
You can modify the password as needed .
      value : mypassw0rd

# Define the tensorflow service
---
apiVersion : v1
kind : Service
metadata :
  name : tf - notebook
spec :
  ports :
    - port : 80
      targetPort : 8888
      name : jupyter
  selector :
    app : tf - notebook
  type : LoadBalancer # set Alibaba
Cloud SLB service for the application so that
its services are accessible from the Internet .

```

If you use a GPU deployment solution of Kubernetes earlier than 1.9.3, you must define the following volumes in which the NVIDIA drivers reside:

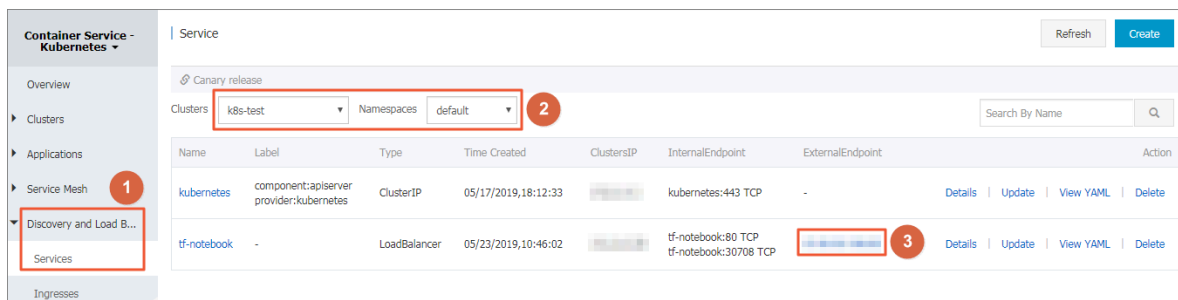
```

volumes :
  - hostPath :
      path : / usr / lib / nvidia - 375 / bin
      name : bin
  - hostPath :
      path : / usr / lib / nvidia - 375
      name : lib

```

When you orchestrate your deployment template in a cluster by using the GPU deployment solution of Kubernetes earlier than 1.9.3, your template must be highly dependent on the cluster. As a result, portability of the template is not achievable. However, in Kubernetes version 1.9.3 and later, you do not need to specify these hostPaths because the NIVEA plugins automatically discover the library links and execution files required by the drivers.

5. In the left-side navigation pane under Container Service-Kubernetes, choose **Discovery and Load Balancing > Services**. Then, select the target cluster and namespace, and then view the external endpoint of the tf-notebook service.

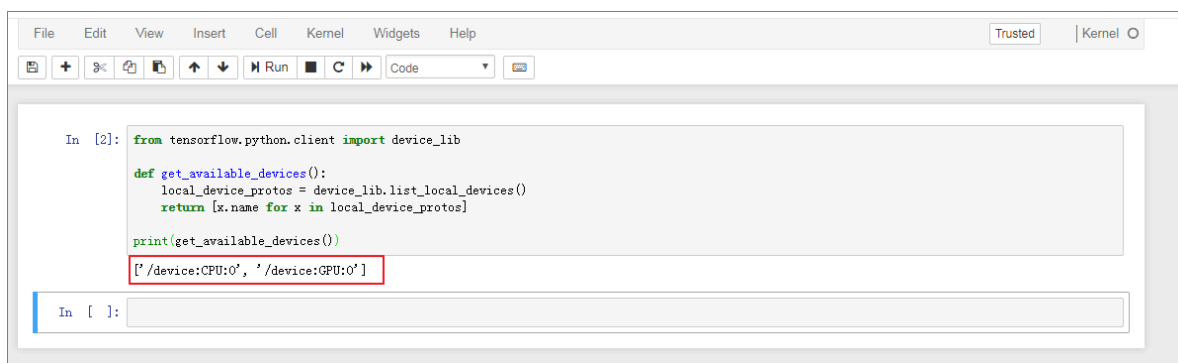


6. Access the Jupyter application in a browser. The access address is `http://EXTERNAL_IP`. You need to enter the password set in the template.
7. By running the following program, you can verify that this Jupyter application can use GPU, and the program is able to list all devices that can be used by Tensorflow:

```
from tensorflow.python.client import device_lib

def get_available_devices():
    local_device_protos = device_lib.list_local_devices()
    return [x.name for x in local_device_protos]

print(get_available_devices())
```



1.4.4 Upgrade the NVIDIA driver of a Kubernetes cluster GPU node

This topic describes how to upgrade the NVIDIA driver of a Kubernetes cluster GPU node where services are running or no service runs.

Prerequisites

- You have created a Kubernetes GPU cluster. For more information, see [Configure a Kubernetes GPU cluster to support GPU scheduling](#).

- You have connected to the Kubernetes GPU cluster by using `kubectl`, see [Connect to a Kubernetes cluster by using kubectl](#).

Upgrade the NVIDIA driver of a GPU node where services are running

1. Run the following command to disable scheduling for the target GPU node:

```
kubectl cordon node-name
```



Note:

- Only the NVIDIA drivers of Worker nodes can be upgraded.
- The `node-name` parameter must be in the format of `your-region-name.node-id`.
 - `your-region-name` indicates the name of the region where your cluster is located.
 - `node-id` indicates the ID of the ECS instance where the target node is located.

You can run the following command to view `node-name`:

```
kubectl get node
```

```
[root@gpu-test ~]# kubectl cordon cn-hangzhou.i-  
node/cn-hangzhou.i- already cordoned
```

2. Run the following command to migrate the pods on the target node to other nodes:

```
kubectl drain node-name --grace-period=120 --ignore-daemonsets=true
```

```
[root@gpu-test ~]# kubectl drain cn-hangzhou.i-  
node/cn-hangzhou.i- --grace-period=120 --ignore-daemonsets=true  
node/cn-hangzhou.i- cordoned  
WARNING: Ignoring DaemonSet-managed pods: flexvolume-  
pod/domain-nginx- evicted  
pod/old-nginx- evicted  
pod/new-nginx- evicted  
pod/old-nginx- evicted
```

3. Run the following command to log on to the target node:

```
ssh root @ xxx . xxx . x . xx
```

4. On the target node, run the following command to view the driver version:

```
nvidia - smi
```

```
[root@ ~]# nvidia-smi
Fri Jan 18 16:44:52 2019
+-----+
| NVIDIA-SMI 384.111                Driver Version: 384.111 |
+-----+-----+
| GPU Name Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |
+-----+-----+
| 0 Tesla P4             On   | 00000000:00:08.0 Off |           0         |
| N/A   24C    P8      6W / 75W | 0MiB / 7606MiB |      0%      Default |
+-----+-----+
+-----+
| Processes:                         GPU Memory |
|   GPU       PID    Type    Process name      Usage   |
+-----+-----+
| No running processes found |
+-----+
```

5. Run the following commands to remove the existing driver:



Note:

- If the existing driver is v 384 . 111 , you can directly run the commands in this step.
- If the existing driver is not v 384 . 111 , you must download the correct driver version from NVIDIA Website before running the commands in this step.

```
cd / tmp
```

```
curl -O https://cn.download.nvidia.cn/tesla/384.111/NVIDIA-Linux-x86_64-384.111.run
```

```
chmod u+x NVIDIA-Linux-x86_64-384.111.run
```

```
./NVIDIA-Linux-x86_64-384.111.run --uninstall -a -s -q
```

6. Run the following command to restart the target node:

```
reboot
```

7. Download the driver version that you want from the NVIDIA Website. This example uses v 410 . 79 .

8. Run the following command to install the downloaded NVIDIA driver in the directory where the driver is downloaded:

```
sh . / NVIDIA - Linux - x86_64 - 410 . 79 . run - a - s - q
```

9. Run the following commands to add the following settings to the NVIDIA driver:

```
nvidia - smi - pm 1 || true
```

```
nvidia - smi - acp 0 || true
```

10. Run the following command to update two device plugins:

```
mv / etc / kubernetes / manifests / nvidia - device - plugin . yml /
```

```
mv / nvidia - device - plugin . yml / etc / kubernetes / manifests /
```

11. In any path of the Master node, run the following command to enable scheduling for the target node:

```
kubectll uncordon node-name
```

Verify the results

Run the following command on the Master node. Then check the driver version for the target GPU node. The system displays that the driver is v 410 . 79 , indicating the node driver has been upgraded.



Note:

You need to replace the *node-name* parameter with your target node name.

```
kubectll exec -n kube-system -t nvidia-device-plugin-node-name nvidia-smi
```

```
[root@gpu-test ~]# kubectll exec -n kube-system -t nvidia-device-plugin-cn- nvidia-smi
Mon Jan 21 03:14:48 2019
```

NVIDIA-SMI 410.79 Driver Version: 410.79 CUDA Version: N/A									
GPU	Name	Persistence-M	Bus-Id	Disp.A	Volatile	Uncorr.	ECC		
Fan	Temp	Perf	Pwr:Usage/Cap	Memory-Usage	GPU-Util	Compute	M.		
0	Tesla P4	On	00000000:00:08.0	Off	0%		0		
N/A	21C	P8	6W / 75W	0MiB / 7611MiB					

Processes:					GPU Memory
GPU	PID	Type	Process name		Usage
No running processes found					

Upgrade the NVIDIA driver of a GPU node where no service runs

1. Run the following command to log on to the target GPU node:

```
ssh root @ xxx . xxx . x . xx
```

2. On the target node, run the following command to view the driver version:

```
nvidia - smi
```

```
[root@ ~]# nvidia-smi
Fri Jan 18 16:44:52 2019
+-----+
| NVIDIA-SMI 384.111                Driver Version: 384.111 |
+-----+-----+
| GPU  Name            Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf    Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |
+-----+-----+
|   0   Tesla P4             On      | 00000000:00:08.0 Off |   0%        0      |
| N/A   24C    P8             6W / 75W | 0MiB / 7606MiB |   0%        Default |
+-----+-----+
+-----+
| Processes:                         GPU Memory |
|   GPU       PID    Type    Process name      Usage    |
+-----+
| No running processes found |
+-----+
```

3. Run the following commands to remove the existing driver:



Note:

- If the existing driver is v 384 . 111 , you can directly run the commands in this step.

- If the existing driver is not v 384 . 111 , you must download the correct driver version from NVIDIA Website before running the commands in this step.

```
cd / tmp
```

```
curl -O https://cn.download.nvidia.cn/tesla/384.111/NVIDIA-Linux-x86_64-384.111.run
```

```
chmod u+x NVIDIA-Linux-x86_64-384.111.run
```

```
./NVIDIA-Linux-x86_64-384.111.run --uninstall -a -s -q
```

4. Run the following command to restart the target node:

```
reboot
```

5. Download the driver version that you want from the NVIDIA Website. This example uses v 410 . 79 .

6. Run the following command to install the downloaded NVIDIA driver in the directory where the driver is downloaded:

```
sh ./NVIDIA-Linux-x86_64-410.79.run -a -s -q
```

7. Run the following commands to add the following settings to the NVIDIA driver:

```
nvidia-smi -pm 1 || true
```

```
nvidia-smi -acp 0 || true
```

Verify the results

Run the following command on the Master node. Then check the driver version for the target GPU node. The system displays that the driver is v 410 . 79 , indicating the node driver has been upgraded.



Note:

You need to replace the `node-name` parameter with your target node name.

```
kubectl exec -n kube-system -t nvidia-device-plugin-node-name nvidia-smi
```

```
[root@gpu-test ~]# kubectl exec -n kube-system -t nvidia-device-plugin-cn- nvidia-smi
Mon Jan 21 03:14:48 2019

+-----+
| NVIDIA-SMI 410.79      | Driver Version: 410.79      | CUDA Version: N/A      |
+-----+-----+
| GPU   Name           Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf    Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |
+-----+-----+
|  0 Tesla P4             0n      | 00000000:00:08.0 Off  |          0          |
| N/A   21C    P8             6W / 75W |  0MiB / 7611MiB |      0%      Default  |
+-----+-----+

+-----+
| Processes:             | GPU Memory |
| GPU       PID    Type   Process name          Usage       |
+-----+-----+
| No running processes found |
+-----+
```

1.4.5 Create a multi-zone Kubernetes cluster

You can create a multi-zone Kubernetes cluster to guarantee high availability.

Prerequisites

- You have activated the following services: Container Service, Resource Orchestration Service (ROS), Resource Access Management (RAM), and Auto Scaling service.

Log on to the [Container Service console](#), [ROS console](#), [RAM console](#), and [Auto Scaling console](#) to activate the corresponding services.



Note:

The deployment of Container Service Kubernetes clusters depends on the application deployment capabilities of Alibaba Cloud ROS. Therefore, activate ROS before creating a Kubernetes cluster.

- You must create a Virtual Private Cloud (VPC) and create at least three VSwitches in the VPC. To achieve high availability, we recommend that you create VSwitches in different availability zones.
- You need to manually configure SNAT for each VSwitch in the VPC. Otherwise, instances in the VPC cannot access the Internet normally.

Context

You can create Kubernetes clusters with ECS instances in different availability zones by using the Container Service console to achieve high availability.

Context

During cluster creation, Container Service performs the following operations:

- Create Elastic Compute Service (ECS) instances and configure to log on to other nodes from management nodes with the SSH public key. Install and configure the Kubernetes cluster by using CloudInit.
- Create a security group. This security group allows the VPC inbound access of all the ICMP ports.
- Create a RAM user and an AccessKey. The RAM user has the permissions for querying, creating, and deleting ECS instances, the permissions for adding and deleting cloud disks, and all permissions for the operations on Server Load Balancer (SLB), CloudMonitor, VPC, Log Service, and Network Attached Storage (NAS). The Kubernetes cluster dynamically creates SLB instances, cloud disks, and VPC routing rules according to your configurations.
- Create an intranet SLB instance and expose the port 6443.
- Create an Internet SLB instance and expose the port 6443. (If you enable the SSH logon for Internet when creating the cluster, port 22 is exposed. Otherwise, port 22 is not exposed.)

Limits

- The SLB instances created with the cluster support only the Pay-As-You-Go billing method.
- Kubernetes clusters support only the Virtual Private Cloud (VPC) network type.
- By default, each account has a specified quota for the cloud resources it can create. If the number of cloud resources exceeds the quota, the account cannot create a cluster. Make sure you have enough quota before creating a cluster. To increase your quota, open a ticket.
 - By default, each account can create up to 5 clusters in all regions and add up to 40 nodes to each cluster. To create more clusters or nodes, open a ticket.
 - By default, each account can create up to 100 security groups.
 - By default, each account can create up to 60 Pay-As-You-Go SLB instances.
- The limits for ECS instances are as follows:
 - Only the CentOS operating system is supported.
 - The Pay-As-You-Go and Subscription ECS instances can be created.

Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Clusters in the left-side navigation pane. Click Create Kubernetes Cluster in the upper-right corner.
3. On the Create Kubernetes Cluster page, click Multi-AZ Kubernetes.
4. Enter the cluster name.

The cluster name can be 1–63 characters long and contain numbers, Chinese characters, English letters, and hyphens (-).

5. Select the region where the cluster is located.
6. Select a VPC.

Select a VPC from the existing VPC drop-down list and select three VSwitches under the VPC. To achieve high availability, we recommend that you select the VSwitches located in different zones.

VPC

VPC123 (vpc-2zercq4pyanzxsfiidyl)

VSwitch

Select three VSwitches. To ensure high availability, switches in different zones are recommended.

	Name	ID	Zone	CIDR
<input checked="" type="checkbox"/>	vs-01-k8s	vsw-2zey8crxysjn3pdrbykd	China North 2 (Beijing) ZoneB	192.168.35.0/24
<input checked="" type="checkbox"/>	test	vsw-2zeva6cj8lotckx1b9fc1	China North 2 (Beijing) ZoneC	192.168.18.0/24
<input checked="" type="checkbox"/>	VSwitch123	vsw-2zeydhl5uwh1ej522lauo	China North 2 (Beijing) ZoneA	192.168.0.0/24

7. Configure the Master nodes and Worker nodes.
 - a) Select a node payment type from Pay-As-You-Go and Subscription.
 - b) Select instance types of the Master nodes and Worker nodes, and set the number of Worker nodes.



Note:

- Currently, only the CentOS operating system is supported.
- Currently, only three Master nodes can be created.
- Each cluster can contain up to 37 Worker nodes. To create more nodes, open a ticket.
- System disks are attached to Master nodes and Worker nodes by default. Available system disks include SSD Cloud Disks and Ultra Cloud Disks.

- You can also manually attach a data disk to the Worker node. The data disk can be an Ultra Cloud Disk or an SSD Cloud Disk.

Node Type
Pay-As-You-Go
Subscription

Master Configuration

Instance Type

Zone	Type	Quantity
Zone E	2 Core(s) 4 G (ecs.sn1ne.large)	1 unit(s)
Zone D	2 Core(s) 4 G (ecs.n1.medium)	1 unit(s)
Zone A	4 Core(s) 16 G (ecs.i1.xlarge)	1 unit(s)

System Disk

Ultra Cloud Disk
120
GiB

Worker Configuration

Instance Type

Zone	Type	Quantity
Zone E	2 Core(s) 4 G (ecs.sn1ne.large)	1 unit(s)
Zone D	2 Core(s) 4 G (ecs.n1.medium)	1 unit(s)
Zone A	4 Core(s) 16 G (ecs.i1.xlarge)	1 unit(s)

System Disk

Ultra Cloud Disk
40
GiB

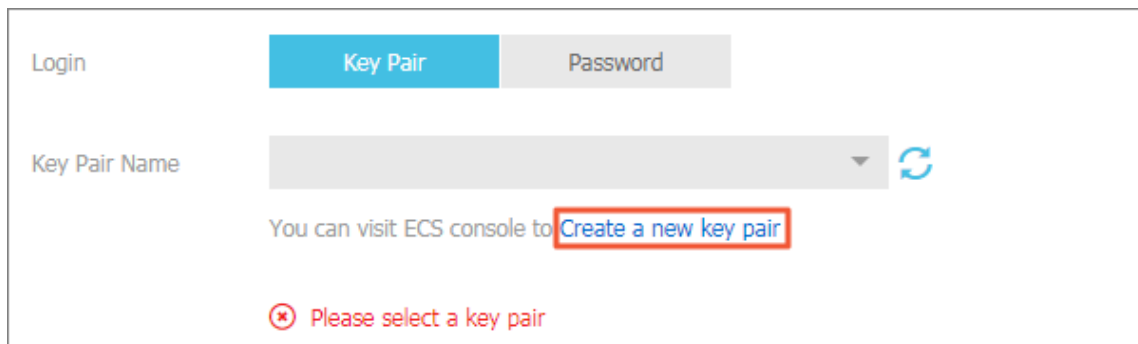
☒ Attach Data Disk
Ultra Cloud Disk
100
GiB

8. Configure the logon mode.

- Set the key pair.

When creating a cluster, select the key pair logon mode and click Create a new key pair. In the ECS console, create a key pair. For details, see [Create an SSH key](#)

[pair](#). After the key pair is created, set the key pair as the credentials for logging on to the cluster.



The screenshot shows a login interface with two tabs: 'Key Pair' (active) and 'Password'. Below the tabs is a 'Key Pair Name' dropdown menu. A red box highlights the text 'Create a new key pair' in the dropdown menu. Below the dropdown, there is a red error message: 'Please select a key pair'.

- Set the password.
 - Logon Password: Set the node logon password.
 - Confirm Password: Confirm your node logon password.

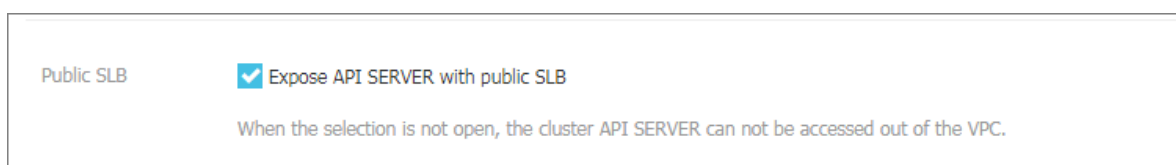
9. Specify the Pod Network CIDR and Service CIDR parameters.

Both of them cannot overlap with the Classless Inter-Domain Routing (CIDR) block used by VPC and the existing Kubernetes clusters in VPC. The values cannot be modified after the cluster is created. In addition, service CIDR cannot overlap with pod network CIDR. For more information about how to plan Kubernetes CIDR blocks, see [Plan Kubernetes CIDR blocks under VPC](#).

10. Select whether to enable Use Public SLB to Expose API Server.

API server provides add, delete, edit, check, watch, and other HTTP Rest interfaces for a variety of resource objects (such as pods and services).

- a. If you select to enable this option, the Internet SLB is created and the port 6443 of the Master nodes is exposed. The port corresponds to the API server. Then you can use kubeconfig to connect to and operate the clusters through the Internet.
- b. If you select not to enable this option, the Internet SLB is not created. You can only use kubeconfig to connect to and operate the clusters inside the VPC.



The screenshot shows a form with the label 'Public SLB'. There is a checked checkbox next to the text 'Expose API SERVER with public SLB'. Below this, there is a note: 'When the selection is not open, the cluster API SERVER can not be accessed out of the VPC.'

11. Select whether to enable SSH logon for Internet.



Note:

To enable SSH access for Internet, you must select Use Public SLB to Expose API Server.

- If you select to enable SSH access for Internet, you can use SSH to access a cluster.
- If you select not to enable SSH access for Internet, you cannot access a cluster by using SSH or connect to a cluster by using kubectl. To access a cluster instance by using SSH, manually bind an EIP to the ECS instance, configure security group rules, and open the SSH port (22). For details, see [Access Kubernetes clusters by using SSH](#).

SSH Login

☐ Enable SSH access for Internet

If you choose not to open it, please refer to [SSH access to Kubernetes cluster](#) to manually enable SSH access.

12. Select whether to install a cloud monitoring plug-in on your ECS.

You can install a cloud monitoring plug-in on the ECS node to view the monitoring information of the created ECS instances in the CloudMonitor console.

Monitoring Plug-in

☒ Install cloud monitoring plug-in on your ECS.

Installing a cloud monitoring plug-in on the node allows you to view the monitoring information of the created ECS instance in the CloudMonitor console

13. Select whether to use Log Service. You can select an existing project or create a project.

If you select Using SLS, the Log Service plug-in is automatically configured in the cluster. When creating an application, you can quickly use Log Service with a simple configuration. For details, see [Use Log Service to collect Kubernetes cluster logs](#).

Log Service

☒ Using SLS

Select Project

Create Project

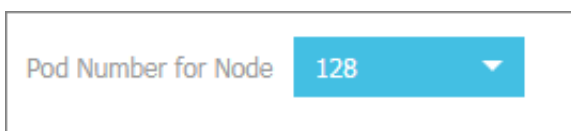
A SLS Project named k8s-log-{ClusterID} will be created automatically

14. Select whether to show advance config.

a. Select a network plug-in. Available network plug-ins are Flannel and Terway. For details, see [Do I select the Terway or Flannel plugin for my Kubernetes cluster network?](#).

- Flannel: a simple and stable community Flannel CNI plug-in. It provides only a few simple features. For example, it does not support the Kubernetes Network Policy.
- Terway: a network plug-in developed by Alibaba Cloud Container service. It 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.

b. Set the number of pods for a node, that is, the maximum number of pods that can be run by a single node.



The screenshot shows a form field with the label "Pod Number for Node" and a blue dropdown menu displaying the value "128".

c. Select whether to use Custom Image. The ECS instance installs the default CentOS version if no custom image is selected.

Currently, you can only select an image based on CentOS custom version to quickly deploy the environment you need.

d. Sets whether to use Custom Cluster CA. If this option is selected, the CA certificate can be added to the Kubernetes cluster, which enhances the security of information exchange between the server and client.



The screenshot shows a form field with the label "Cluster CA" and a checkbox labeled "Custom Cluster CA". The checkbox is currently unchecked.

15. Click **Create**, confirm the Internet access for VPC in the displayed dialog box, and click **OK** to start the deployment.

Confirm Cluster Configuration

Item	Status	Detail
Account status check	Success	
Products Activation Status	Success	
Product Quota Check	Success	
Internet access for VPC	<div>Confirm</div>	Ensure that the NAT gateway has been configured for the VPC, or configure SNAT manually. Otherwise, instances in the VPC can not access the Internet the cluster will fail. Check Again

OK

Cancel



Note:

A multi-node Kubernetes cluster typically takes 10 minutes to be created.

Result

View cluster deployment results.

After the cluster is successfully created, you can view the cluster in the Cluster List of the Container Service console.

Cluster List

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

Refresh

Create Kubernetes Cluster

Help: [Create cluster](#) [Scale cluster](#) [Connect to Kubernetes cluster via kubectl](#) [Manage applications with commands](#)

Name






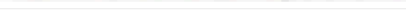
Cluster Name/ID	Cluster Type	Region (All)	Network Type	Cluster Status	Time Created	Kubernetes Version	Action
k8s-cluster		China East 1 (Hangzhou)	VPC vpc-bp1kyevdj...	<div>Running</div>	09/26/2018,17:41:26	1.11.2	Manage View Logs Dashboard Scale Cluster More

What's next

- Click View Logs at the right of the cluster to view the cluster logs. To view more detailed information, click Stack Events.

Detailed resource deployment logs: Stack Events	
Time	Information
06/22/2018,11:13:50	c06eadb6387ec4c6080d495e243649a7b Start to DescribeK8sUserCertConfig
06/22/2018,11:03:39	c06eadb6387ec4c6080d495e243649a7b Set up k8s DNS configuration successfully
06/22/2018,11:02:30	c06eadb6387ec4c6080d495e243649a7b Stack CREATE completed successfully:o
06/22/2018,11:02:30	c06eadb6387ec4c6080d495e243649a7b Start describeStackInfo
06/22/2018,11:02:29	c06eadb6387ec4c6080d495e243649a7b Start describeStackInfo
06/22/2018,10:46:55	c06eadb6387ec4c6080d495e243649a7b Successfully to CreateStack
06/22/2018,10:46:55	c06eadb6387ec4c6080d495e243649a7b Start to wait stack ready
06/22/2018,10:46:53	c06eadb6387ec4c6080d495e243649a7b Start to create cluster task
06/22/2018,10:46:53	c06eadb6387ec4c6080d495e243649a7b Start to CreateK8sCluster
06/22/2018,10:46:53	c06eadb6387ec4c6080d495e243649a7b Start to CreateStack
06/22/2018,10:46:50	c06eadb6387ec4c6080d495e243649a7b Start to validateCIDR
06/22/2018,10:46:41	c06eadb6387ec4c6080d495e243649a7b Start create cluster certificate

- You can also click Manage on the right of the cluster to view the basic information and connection information about this cluster.

Basic Information			
Cluster ID: 	VPC	 Running	Region: China East 1 (Hangzhou)
Connection Information			
API Server Internet endpoint			
API Server Intranet endpoint			
Master node SSH IP address			
Service Access Domain			

In the Cluster Info section:

- API Server Internet endpoint:** The IP address and port through which the Kubernetes API server provides services for the Internet. It enables you to manage the cluster by using kubectl or other tools on your terminal.
- API Server Intranet endpoint:** The IP address and port through which the Kubernetes API server provides services inside the cluster. This IP address is the

address of the SLB instance, and three Master nodes in the backend provide the services.

- **Master node SSH IP address:** You can directly log on to the Master nodes by using SSH to perform routine maintenance for the cluster.
- **Service Access Domain:** Provides the services in the cluster with access domain name for testing. The service access domain name suffix is `< cluster_id >.< region_id >.alicontainer.com`.

1.4.6 Connect to a Kubernetes cluster by using kubectl

To connect to a Kubernetes cluster from a client computer, use the Kubernetes command line client [kubectl](#).

Procedure

1. Download the latest kubectl client from the [Kubernetes release page](#).
2. Install and set the kubectl client.

For more information, see [Install and set kubectl](#).

3. Configure the cluster credentials.

You can use the `scp` command to safely copy the master node configurations from the `/etc/kubernetes/kube.conf` file on the master virtual machine of the Kubernetes cluster to the `$HOME/.kube/config` file (where the `kubectl` expected credentials reside) of the local computer.

- If you select Password in the Login field when creating the cluster, copy the kubectl configuration file in the following method:

```
mkdir $HOME/.kube
scp root@<master - public - ip>:/etc/kubernetes/kube.conf $HOME/.kube/config
```

- If you select Key Pair in the Login field when creating the cluster, copy the kubectl configuration file in the following method:

```
mkdir $HOME/.kube
scp -i [the storage path of the .pem private key file on the local
```

```
machine ] root @:/ etc / kubernetes / kube . conf $ HOME /.  
kube / config
```

You can check the cluster `master - public - ip` on the cluster information page.

- a) Log on to the [Container Service console](#).
- b) Under Kubernetes, click Clusters in the left-side navigation pane.
- c) Click Manage at the right of the cluster.

In the Connection Information section, view the Master node SSH IP address.

1.4.7 Use kubectl commands in Cloud Shell to manage a Kubernetes cluster

This topic describes how to use kubectl on Cloud Shell to manage a Kubernetes cluster after you log on to the console of Alibaba Cloud Container Service for Kubernetes.

Prerequisites

You have created a Kubernetes cluster. For more information, see [Create a Kubernetes cluster](#).

Context

If you want to use kubectl to manage a Kubernetes cluster of Container Service, you can download kubectl to your local host. For more information, see [Connect to a Kubernetes cluster by using kubectl](#). Additionally, you can also start Cloud Shell on the console of Container Service for Kubernetes, and then use kubectl on Cloud Shell to manage a Kubernetes cluster.

Procedure

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Kubernetes, click Clusters.

3. In the Action column of the target cluster, choose More > cos.cls.cloudshell.

Container Service - Kubernetes

Cluster List

You can create up to 100 clusters and can add up to 1000 nodes in each cluster. [Refresh](#) [Create Kubernetes Cluster](#)

Help: [Create cluster](#) [Create GPU clusters](#) [Scale cluster](#) [Connect to Kubernetes cluster via kubectl](#) [Manage applications with commands](#) [Cluster planning](#) [Troubleshoot cluster creation](#)

Name

Cluster Name/ID	Cluster Type	Region (All)	Network Type	Cluster Status	Number of Nodes	Time Created	Kubernetes Version	Action
kubernetes-test	Kubernetes	China East 1 (Hangzhou)	VPC vpc-bp1frndpc9b...	Running	6	12/18/2018,15:52:19	1.11.5	Manage View Logs Dashboard Scale Cluster More
k8s-delete-node	Kubernetes	China East 1 (Hangzhou)	VPC vpc-bp1nr6ohb0c...	Running	5	12/17/2018,10:26:28	1.11.5	Manage View Logs Dashboard Scale Cluster More
k8s-test	Kubernetes	China North 2 (Beijing)	VPC vpc-zzefxdn1wdl...	Running	6	12/05/2018,13:27:46	1.11.5	Manage View Logs Dashboard Scale Cluster More
k8s-managed-cluster	ManagedKubernetes	China East 1 (Hangzhou)	VPC vpc-bp1kd7yn4qn...	Running	3	11/01/2018,11:21:13	1.11.5	Delete Add Existing Instance Upgrade Cluster Automatic Scaling Add-on Upgrade Deploy Istio
test-mia	Kubernetes	China East 1 (Hangzhou)	VPC vpc-bp1kyevdj...	Running	7	09/17/2018,11:37:55	1.11.5	



Note:

Do the following:

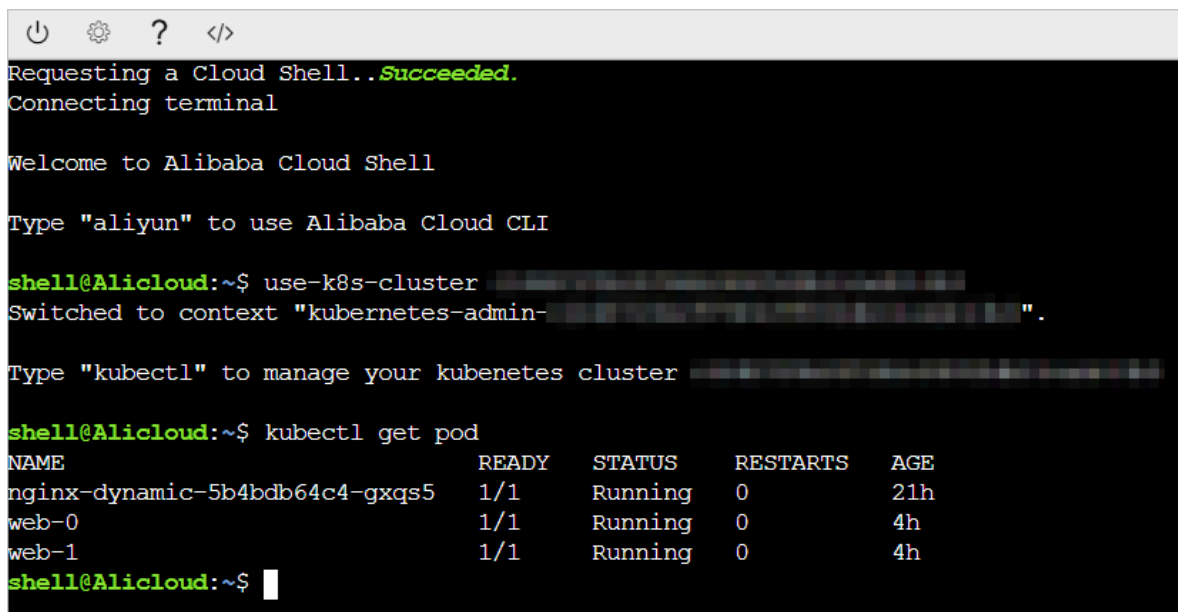
- On the Authorization page, click OK to obtain a temporary access key that expires within one hour.
- On the Storage Space page, click Create Now or Skip as needed.

4. On Cloud Shell, you can use kubectl to manage a Kubernetes cluster of Container Service.



Note:

When you start Cloud Shell associated with the Kubernetes cluster, the system loads the `kubeconfig` file of the cluster onto Cloud Shell. Then you can use `kubectl` to manage your cluster.



```

Requesting a Cloud Shell...Succeeded.
Connecting terminal

Welcome to Alibaba Cloud Shell

Type "aliyun" to use Alibaba Cloud CLI

shell@Alicloud:~$ use-k8s-cluster
Switched to context "kubernetes-admin-".

Type "kubectl" to manage your kubernetes cluster

shell@Alicloud:~$ kubectl get pod
NAME                                READY   STATUS    RESTARTS   AGE
nginx-dynamic-5b4bdb64c4-gxqs5     1/1     Running   0           21h
web-0                               1/1     Running   0           4h
web-1                               1/1     Running   0           4h
shell@Alicloud:~$
```

1.4.8 Use a ServiceAccount token to access a managed Kubernetes cluster

This topic describes how to use a ServiceAccount token to access a managed Kubernetes cluster.

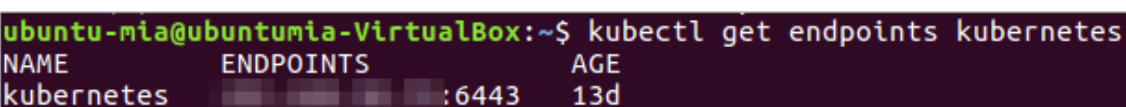
Context

- You have created a managed Kubernetes cluster. For more information, see [Create a managed Kubernetes cluster](#).
- You have connected to the managed Kubernetes cluster by using `kubectl`, see [Connect to a Kubernetes cluster by using kubectl](#).

Procedure

1. Run the following command to obtain the API server intranet endpoint:

```
$ kubectl get endpoints kubernetes
```



```

ubuntu-mia@ubuntumia-VirtualBox:~$ kubectl get endpoints kubernetes
NAME            ENDPOINTS                               AGE
kubernetes      :6443                                   13d
```

2. Create a file named `kubernetes - public - service . yaml` and set the `ip` parameter to the intranet endpoint obtained in step 1.

```
kind : Service
apiVersion : v1
metadata :
  name : kubernetes - public
spec :
  type : LoadBalancer
  ports :
    - name : https
      port : 443
      protocol : TCP
      targetPort : 6443
---
apiVersion : v1
kind : Endpoints
metadata :
  name : kubernetes - public
  namespace : default
subsets :
- addresses :
  - ip : < API Service address > # Set this parameter to
    the intranet endpoint obtained in step 1 .
  ports :
    - name : https
      port : 6443
      protocol : TCP
```

3. Run the following command to deploy the API server Internet endpoint:

```
$ kubectl apply -f kubernetes - public - service . yaml
```

4. Run the following command to obtain the Internet SLB address, namely, `EXTERNAL - IP` :

```
$ kubectl get service name
```



Note:

The `name` parameter in the command and the `name` parameter in the `kubernetes - public - service . yaml` file of step 2 must be set to the same value. In this example, this parameter is set to `kubernetes - public` .

```
ubuntu-mia@ubuntumia-VirtualBox:~$ kubectl get service kubernetes-public
NAME                TYPE          CLUSTER-IP      EXTERNAL-IP      PORT(S)          AGE
kubernetes-public   LoadBalancer  10.100.100.10    10.100.100.10    443:30080/TCP    7d
```

- Run the following command to view the corresponding secret of the ServiceAccount (in this example, the *namespace* parameter is set to default):

```
$ kubectl get secret --namespace=namespace
```

```
ubuntu-mia@ubuntumia-VirtualBox:~$ kubectl get secret --namespace=default
```

NAME	TYPE	DATA	AGE
aliyun-acr-credential-a	kubernetes.io/dockerconfigjson	1	13d
aliyun-acr-credential-b	kubernetes.io/dockerconfigjson	1	13d
	kubernetes.io/service-account-token	3	13d

6. Run the following command to obtain a token value:

```
$ kubectl get secret -n --namespace=namespace -o  
jsonpath={.data.token} | base64 -d
```



Note:

The *namespace* parameter in this command and the *namespace* parameter in step 5 must be set to the same value.

7. Run the following command to access the managed Kubernetes cluster:

```
$ curl -k -H 'Authorization: Bearer token' https://service-ip
```



Note:

- The value of `token` is the token value obtained in step 6.
- The value of `service - ip` is the Internet SLB address obtained in step 4, that is, `EXTERNAL - IP`.

Result

After you run the command, the following message is displayed, indicating that you have connected to the cluster.












```
ubuntu-mia@ubuntu-mia-VirtualBox:~$ curl -k -H 'Authorization: Bearer 'https://
{
  "kind": "Status",
  "apiVersion": "v1",
  "metadata": {
  },
  "status": "Failure",
  "message": "Forbidden: User \"system:serviceaccount:default:default\" cannot get path \"/\".",
  "reason": "Forbidden",
  "details": {
  },
  "code": 403
}
ubuntu-mia@ubuntu-mia-VirtualBox:~$ ^C
```

1.4.9 Access a Kubernetes cluster by using SSH

If you select not to enable SSH access for Internet when creating the Kubernetes cluster, you cannot access the Kubernetes cluster by using SSH or connect to the Kubernetes cluster by using kubectl. To access the cluster by using SSH after creating the cluster, manually bind Elastic IP (EIP) to the Elastic Compute Service (ECS) instance, configure security group rules, and open the SSH port (22).

Procedure

1. Log on to the [Container Service console](#).
2. Under the Kubernetes menu, click Clusters in the left-side navigation pane.
3. Click Manage at the right of the cluster.
4. In Cluster Resource, click the ID of the Internet SLB. Then, you are redirected to the Instance Details page of your Internet Server Load Balancer instance.

Cluster:k8s-cluster	
Basic Information	
Cluster ID: 	VPC   Running Region: China East 1 (Hangzhou)
Connection Information	
API Server Internet endpoint	
API Server Intranet endpoint	
Master node SSH IP address	
Service Access Domain	
Cluster Resource	
ROS	
Internet SLB	
VPC	
NAT Gateway	

5. Select Instances > Server Load Balancer, and click Add Listener.

6. Add the SSH listening rule.

- Front-end Protocol [Port]: Select TCP and enter 22.
- Backend Protocol [Port]: Enter 22.
- Turn on the Use Server Group switch and select VServer Group.
- Server Group ID: Select sshVirtualGroup.
- Click Next and then click Confirm to create the listener.

Front-end Protocol [Port]:*

TCP
:
22

Port range is 1-65535.

Backend Protocol [Port]:*

TCP
:
22

Port range is 1-65535.

Peak Bandwidth:

No Limits
Configure

Instances charged by traffic are not limited by peak bandwidth. Peak bandwidth range is 1-5000.

Scheduling Algorithm:

Weighted f

Use Server Group:

☒

Server Group Type:

☒ VServer Group
☐ Master-Slave Server Group

Server Group ID:

sshVirtualGn

Automatically Enable Listener After Creation:

☒ Enable

☐ Show Advanced Options

Next
Cancel

- Then, you can use the Server Load Balancer instance IP address to access your cluster by using SSH.

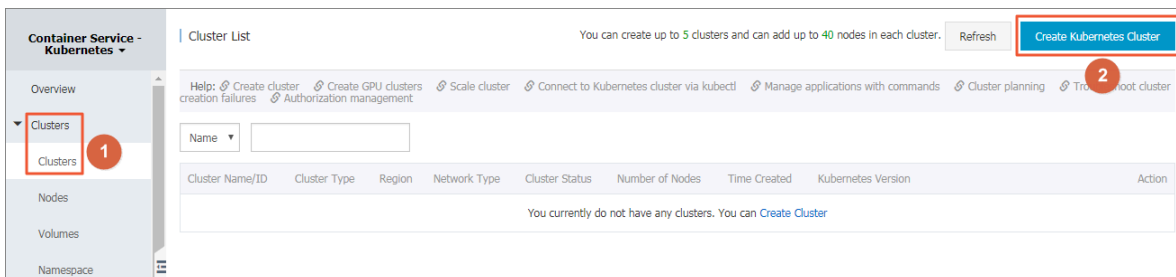
Basic Information	
Server Load Balancer ID: ilb-1459181135	Status: ● Running
Server Load Balancer Name: KubernetesCluster...	Region: China East 1 (Hangzhou)
Instance IP Type: Public IP	Zone: cn-hangzhou-b(Master)/cn-hangzhou-d(Slave)
Network Type: Classic Network	
Billing Information	
Billing Method: Pay by Traffic	Created At: 2018-01-24 11:13:01
Instance IP Address: 114.59.181.35 (Public IP)	Automatic Release Time: -

1.4.10 Set the SSH key pair logon mode for a Kubernetes cluster

This topic describes how to set the SSH key pair logon mode for a Kubernetes cluster, and how to log on to the cluster by using an SSH key pair. Alibaba Cloud Container Service for Kubernetes provides this mode to secure SSH logon to a Kubernetes cluster.

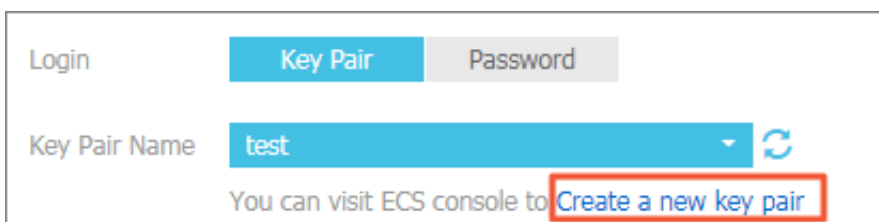
Procedure

- Log on to the [Container Service console](#).
- In the left-side navigation pane under Kubernetes, choose Clusters > Clusters.
- In the upper-right corner, click Create Kubernetes Cluster.

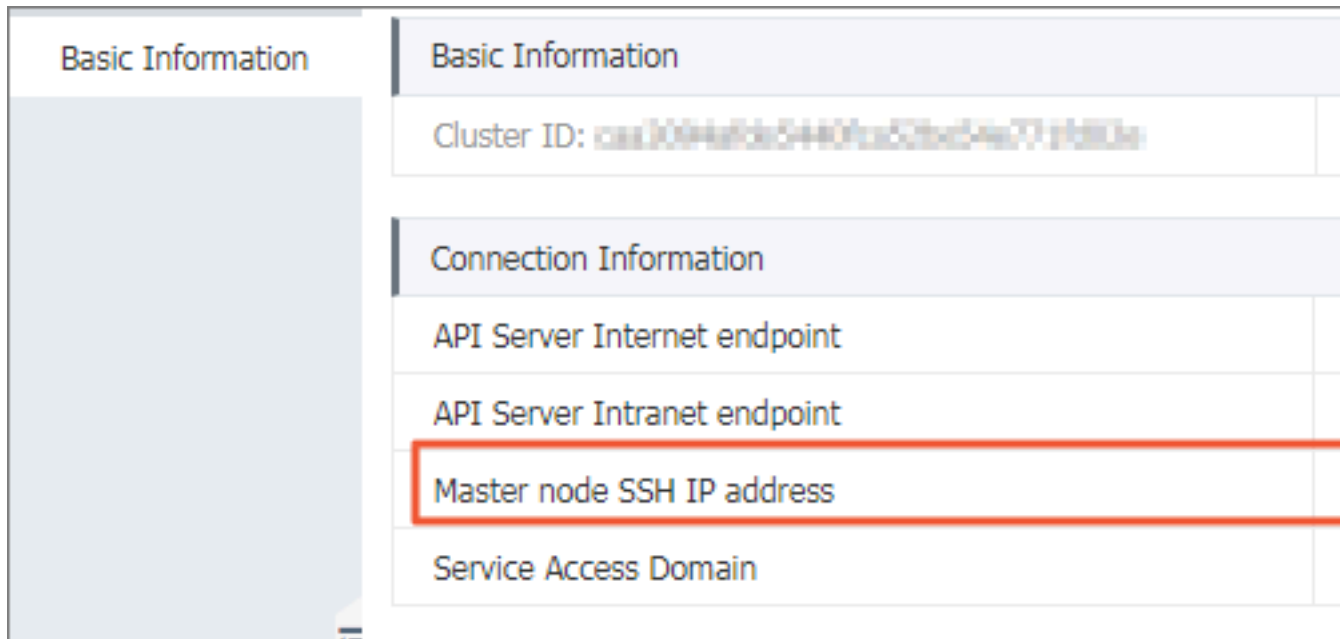


- Set the key pair logon mode for the cluster, and then set other cluster parameters. For more information, see [Create a Kubernetes cluster](#). Then, click Create.

- If you have created key pairs in the Elastic Compute Service (ECS) console, select a key pair from the drop-down list.
- If you have not created any key pairs, click Create a new key pair to create one in the ECS console. For more information, see [Create an SSH key pair](#).



5. After the cluster is created, it is displayed in the cluster list. Click Manage on the right of the target cluster and record the Master node SSH IP address in the cluster information area.



6. Download the `.pem` private key file and complete the cluster logon settings according to your operating system (which can be either Windows or Linux). This topic only describes one type of the cluster logon settings in a Linux operating system as follows:
 - a) Find the directory where you have stored the downloaded `.pem` private key file. For example, `/ root / xxx . pem`.
 - b) Run the following command to modify the private key file attribute: `chmod 400 [the directory where the .pem private key file is stored on your local host]`. In this example, `chmod 400 / root / xxx . pem` is run.
 - c) Run the following command to connect to the cluster: `ssh -i [the directory where the .pem private key file is stored on your local host] root @[Internet IP address]`. In the command, the Internet IP address is the Master node SSH IP address. For example, `ssh -i / root / xxx . pem root @ 10 . 10 . 10 . 100` is run in this example.

For more information about the cluster logon settings in a Window or Linux operating system, see [Connect to a Linux instance by using an SSH key pair](#).

1.4.11 Create a managed Kubernetes cluster

You can create a managed Kubernetes cluster quickly and easily in the Container Service console.

Prerequisites

You have activated the following services: Container Service, Resource Orchestration Service (ROS), Resource Access Management (RAM), and Auto Scaling service.

Log on to the [Container Service console](#), [ROS console](#), [RAM console](#), and [Auto Scaling console](#) to activate the corresponding services.



Note:

The deployment of Container Service managed Kubernetes clusters depends on the application deployment capabilities of Alibaba Cloud ROS. Therefore, you need to activate ROS before creating a managed Kubernetes cluster.

Context

- The SLB instances created with the cluster support only the Pay-As-You-Go billing method.
- Kubernetes clusters support only the Virtual Private Cloud (VPC) network type.
- By default, each account has a specified quota for the cloud resources it can create. If the number of cloud resources exceeds the quota, the account cannot create a cluster. Make sure you have enough quota before creating a cluster. To increase your quota, open a ticket.
 - By default, each account can create up to 100 security groups.
 - By default, each account can create up to 60 Pay-As-You-Go SLB instances.
 - By default, each account can create up to 20 EIPs.

Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Clusters in the left-side navigation pane. The Cluster List page is displayed. Click Create Kubernetes Cluster in the upper-right corner.
3. On the Create Kubernetes Cluster page, click Managed Kubernetes (beta).

4. Enter the cluster name.

The cluster name can be 1–63 characters long and contain numbers, Chinese characters, English letters, and hyphens (-).

* Cluster Name

cluster-managed

The cluster name should be 1-63 characters long, and can contain numbers, Chinese characters, English letters and hyphens.

5. Select the region and zone where the cluster is located.

Region

China North 2
(Beijing)

China East 1
(Hangzhou)

Asia Pacific SE 1
(Singapore)

Zone

China North 2 Zone A

6. Set the cluster network type.



Note:

Kubernetes clusters support only the VPC network type.

VPC: You can select Auto Create to create a VPC together with the Kubernetes cluster, or select Use Existing to use an existing VPC. If you select Use Existing, you can select a VPC and VSwitch from the two displayed drop-down lists.

- **Auto Create:** The system automatically creates a NAT Gateway for your VPC when a cluster is created.
- **Use Existing:** If the selected VPC has a NAT Gateway, Container Service uses the NAT Gateway. Otherwise, the system automatically creates a NAT Gateway by default. If you do not want the system to automatically create a NAT Gateway, deselect the Configure SNAT for VPC check box.



Note:

If you deselect the check box, configure the NAT Gateway on your own to implement the VPC Internet environment with secure access, or manually

configure the SNAT. Otherwise, instances in the VPC cannot access the Internet normally, which leads to cluster creation failure.

VPC

Auto Create
Use Existing

vpc-k8s-for-cs-c2ddc901be4a74ff68ef76b5edb32c4c2...
(vsw-2zeezdtsxxwf2omlsdw23) ZoneA

7. Set the node type.



Note:

Pay-As-You-Go and Subscription types are supported.

Node Type
Pay-As-You-Go

8. Configure the instance.



Note:

- Currently, only the CentOS operating system is supported.
- Each cluster contains at least two nodes.
- Each cluster contains up to 48 nodes. To create more nodes, open a ticket.
- System disks are attached to the instances by default. Available system disks are Ultra Disks and SSD Disks.
- You can attach a data disk to the instances. The data disk can be an Ultra Disk or an SSD Disk.

Instance Configuration

Instance Type
4 Core(s) 8 G (ecs.sn1ne.xlarge)

Quantity
3
unit(s)

System Disk
Ultra Disk
40
GiB

☒ Attach Data Disk
Ultra Disk
100
GiB

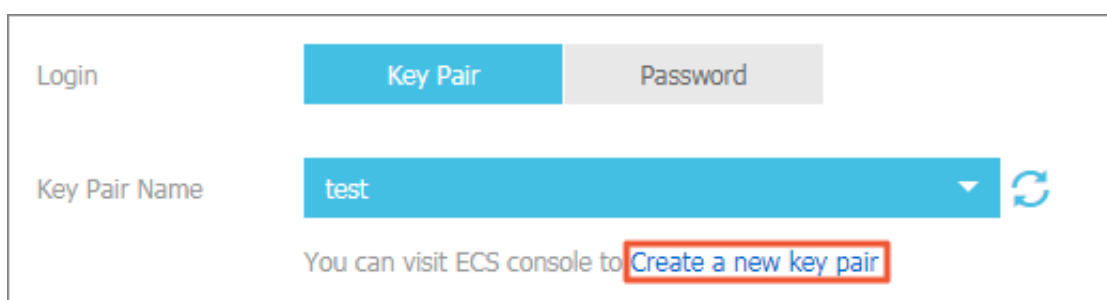
9. Set the logon mode.

- Set the key pair.

When creating a cluster, select the key pair logon mode and click Create a new key pair. In the ECS console, create a key pair. For details, see [Create an SSH key pair](#). After the key pair is created, set the key pair as the credentials for logging on to the cluster.

- Set the password.

- Logon Password: Set the node logon password.
- Confirm Password: Confirm your node logon password.

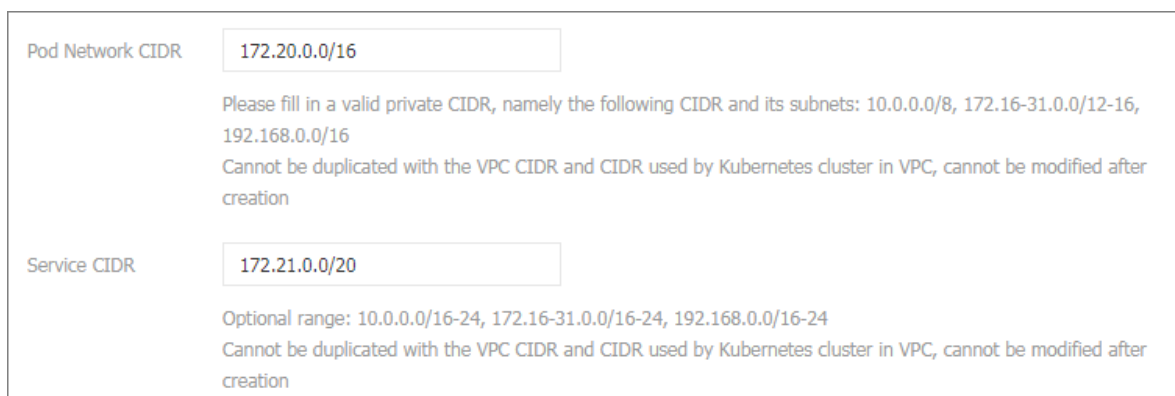


10. Set the Pod Network CIDR and Service CIDR parameters.



Note:

- These two parameters are available only when you select to Use Existing VPC.
- Both Pod Network CIDR and Service CIDR cannot overlap with the Classless Inter-Domain Routing (CIDR) block used by the VPC and the existing Kubernetes clusters in the VPC. The values cannot be modified after the cluster is created. In addition, service CIDR cannot overlap with pod network CIDR. For more information about how to plan Kubernetes CIDR blocks, see [Plan Kubernetes CIDR blocks under VPC](#).



Pod Network CIDR: 172.20.0.0/16

Please fill in a valid private CIDR, namely the following CIDR and its subnets: 10.0.0.0/8, 172.16-31.0.0/12-16, 192.168.0.0/16
Cannot be duplicated with the VPC CIDR and CIDR used by Kubernetes cluster in VPC, cannot be modified after creation

Service CIDR: 172.21.0.0/20

Optional range: 10.0.0.0/16-24, 172.16-31.0.0/16-24, 192.168.0.0/16-24
Cannot be duplicated with the VPC CIDR and CIDR used by Kubernetes cluster in VPC, cannot be modified after creation

11. Select whether to configure a SNAT Gateway for the VPC.



Note:

- If you select Auto Create, you must configure a SNAT Gateway.
- If you select Use Existing, you can select whether to automatically configure a SNAT Gateway. If you select not to automatically configure a SNAT Gateway, you can configure a NAT Gateway for VPC instances to securely access the Internet, or you can configure a SNAT Gateway manually. Otherwise, the instances in the VPC cannot access the Internet, and the cluster fails to be created.

Configure SNAT
☒ Configure SNAT for VPC

If the VPC you choose does not have access to Internet, NAT gateway and EIP will be used to configure SNAT for the VPC. During this period, NAT gateway, EIP, and other resources may be created.

12. Select whether to install a cloud monitoring plug-in on your ECS.

You can install a cloud monitoring plug-in on the ECS node to view the monitoring information of the created ECS instances in the CloudMonitor console.

Monitoring Plug-in
☒ Install cloud monitoring plug-in on your ECS.

Installing a cloud monitoring plug-in on the node allows you to view the monitoring information of the created ECS instance in the CloudMonitor console

13. Select a network plug-in. Available network plug-ins are Flannel and Terway. For details, see [Do I select the Terway or Flannel plugin for my Kubernetes cluster network?](#).

- Flannel: a simple and stable community Flannel plug-in. It provides only a few simple features. For example, it does not support the Kubernetes Network Policy .
- Terway: a network plug-in developed by Alibaba Cloud Container service. It 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.

Network Plugin

Flannel
Terway

14. Set the RDS whitelist.

Add the IP addresses of the ECS instances to the RDS instance whitelist.



Note:

This option is available only when you select to Use Existing VPC.

RDS Whitelist

Select RDS Instances

15. Click Create in the upper-right corner.



Note:

Normally, it takes five minutes to create a multi-node Kubernetes cluster.

Result

After the cluster is successfully created, you can view the cluster on the Cluster List page of the Container Service console.

Container Service - Kubernetes		Cluster List							You can create up to 5 clusters and can add up to 40 nodes in each cluster. Refresh Create Kubernetes Cluster	
Overview		Help: Create cluster Create GPU clusters Scale cluster Connect to Kubernetes cluster via kubectl Manage applications with commands Cluster planning Troubleshoot cluster creation failures Authorization management								
Clusters		Name								
Nodes										
Volumes										
Namespace										
		Cluster Name/ID	Cluster Type	Region (All)	Network Type	Cluster Status	Number of Nodes	Time Created	Kubernetes Version	Action
		managed-cluster	ManagedKubernetes	China North 2 (Beijing)	VPC vpc-2zewej129a...	Running	3	11/16/2018,11:19:26	1.11.2	Manage View Logs Dashboard Scale Cluster More

Click View Logs on the right of the cluster to view the cluster logs on the Cluster Logs page. To view more information, click Stack Events.

Cluster Logs: managed-cluster		Back to Cluster List	Refresh
Detailed resource deployment logs: Stack Events			
Time	Information		
11/16/2018,11:24:54	Set up k8s DNS configuration successfully		
11/16/2018,11:24:54	start to update cluster status CREATE_COMPLETE		
11/16/2018,11:24:54	Successfully to create managed kubernetes cluster		
11/16/2018,11:22:50	Stack CREATE completed successfully:		
11/16/2018,11:20:19	Successfully to CreateStack with response &ros.CreateStackResponse{Id: , Name:"k8s-for-cs-"}.		
11/16/2018,11:20:19	Start to wait stack ready		
11/16/2018,11:20:18	Start to CreateStack		
11/16/2018,11:19:32	Successfully to startLoadBalancerListener (lb-dj17u17byskq0vne2vrbs)		

On the Cluster List page, find the created cluster and click Manage to view the basic information and connection information about this cluster.

Basic Information			
Cluster ID: [REDACTED]	VPC	Running	Region: China North 2 (Beijing)
Cluster Info			
API Server Internet endpoint	https://[REDACTED].k8s-g1.cn-beijing.aliyuncs.com:6443		
Pod Network CIDR	[REDACTED]/16		
Service CIDR	[REDACTED]/20		
Service Access Domain	[REDACTED].cn-beijing.alicontainer.com		
Cluster Resource			
ROS	k8s-for-cs-[REDACTED]		
VPC	vpc-[REDACTED]		

Connect to Kubernetes cluster via kubectl

1. Download the latest kubectl client from the [Kubernetes Edition](#) page .

2. Install and set up the kubectl client. For more information, see [Installing and Setting Up kubectl](#)

3. Configure the cluster credentials:

In the Cluster Info section:

- **API Server Internet endpoint:** The IP address and port through which the Kubernetes API server provides services for the Internet. With the API Server Internet endpoint, you can manage the cluster by using kubectl or other tools on your terminal.
- **Service Access Domain:** Provides the services in the cluster with access domain name for testing. The service access domain name suffix is `< cluster_id >.< region_id >.alicontainer.com`.

You can see [Connect to a Kubernetes cluster by using kubectl](#) and run `kubectl get node` to view the node information of the cluster.

```
ubuntu-mia@ubuntumia-VirtualBox: ~
ubuntu-mia@ubuntumia-VirtualBox:~$ kubectl get nodes
NAME                                STATUS    ROLES    AGE   VERSION
cn-hangzhou.i-[REDACTED]           Ready     <none>   5h    v1.11.2
cn-hangzhou.i-[REDACTED]           Ready     <none>   5h    v1.11.2
cn-hangzhou.i-[REDACTED]           Ready     <none>   5h    v1.11.2
ubuntu-mia@ubuntumia-VirtualBox:~$
```

1.4.12 Upgrade a Kubernetes cluster

This topic describes how to upgrade the Kubernetes version of your cluster in the Alibaba Cloud Container Service for Kubernetes console.

On the cluster list page, you can view the Kubernetes version of your cluster.

Name	Cluster Name/ID	Region	Network Type	Cluster Status	Time Created	Kubernetes Version	Action
test	test	China East 1 (Hangzhou)	VPC	Running	04/24/2018,09:58:28	1.9.3	Manage View Logs Dashboard Scale Cluster More

Limits

- The cluster node instances must be able to access the Internet so that the system can download the required software package.
- We recommend that you create a snapshot for your cluster to guarantee your data security before upgrading the cluster. For more information, see [Create a snapshot](#).
- If you are upgrading a Kubernetes cluster of version 1.8.1 or 1.8.4 to version 1.9.3, all cluster pods will be restarted. This means that applications running on the cluster will be affected. If you are upgrading a Kubernetes cluster of another version, cluster applications will not be affected. However, if a cluster application is highly dependent on the API server, the application may be temporarily affected by the upgrade.
- When a cluster is being upgraded, the network is reset, which means OSS volumes will be remounted to the cluster. As a result, you need to re-create the pods that use OSS volumes after the upgrade.

Preparations

You must make sure that the target cluster is in the healthy status before the upgrade. To do so, follow these steps:

1. Log on to the Master node. For more information, see [Access Kubernetes clusters by using SSH](#) and [Connect to a Kubernetes cluster by using kubectl](#).
2. Run the `kubectl get cs` command to verify that all cluster modules are in the healthy status.

NAME	STATUS	MESSAGE	ERROR
scheduler	Healthy	ok	
controller - manager	Healthy	ok	
etcd - 0	Healthy	{" health ": " true "}	
etcd - 1	Healthy	{" health ": " true "}	
etcd - 2	Healthy	{" health ": " true "}	

3. Run the `kubectl get nodes` command to verify that all nodes are in the ready status.



Note:

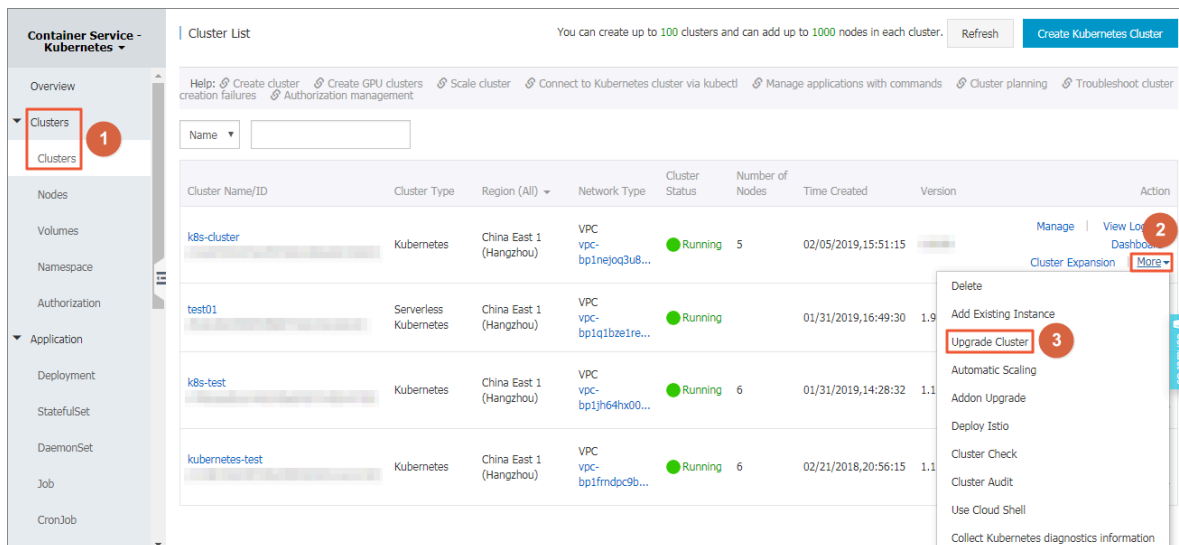
All nodes must be in the ready status.

```
kubectl get nodes
NAME                                STATUS    ROLES    AGE
VERSION
cn-shanghai.i-xxxxxx              Ready    master   38d
v1.9.3
cn-shanghai.i-xxxxxx              Ready    < none > 38d
v1.9.3
cn-shanghai.i-xxxxxx              Ready    < none > 38d
v1.9.3
cn-shanghai.i-xxxxxx              Ready    < none > 38d
v1.9.3
cn-shanghai.i-xxxxxx              Ready    master   38d
v1.9.3
cn-shanghai.i-xxxxxx              Ready    master   38d
v1.9.3
```

If a node is abnormal, you can either repair it manually or open a ticket.

Procedure

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Container Service-Kubernetes, choose Clusters > Clusters.
3. On the right of the target cluster, choose More > Upgrade Cluster.



4. In the displayed dialog box, click Upgrade.

1.4.13 Upgrade a system component

This topic describes how to upgrade a system component.

Prerequisites

You have created a Kubernetes cluster. For more information, see [Create a Kubernetes cluster](#).

Context

The following procedure is for if you need to independently upgrade one or multiple system components of a Kubernetes cluster even if the cluster is of the latest version.

Procedure

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Kubernetes, click Clusters.
3. In the Action column of the target cluster, choose More > Addon Upgrade.

The screenshot shows the 'Container Service - Kubernetes' console. The left navigation pane has 'Clusters' selected. The main area displays a 'Cluster List' table with columns: Cluster Name/ID, Cluster Type, Region (All), Network Type, Cluster Status, Number of Nodes, Time Created, Kubernetes Version, and Action. Five clusters are listed: 'kubernetes-test', 'k8s-delete-node', 'k8s-test', 'k8s-managed-cluster', and 'test-mia'. All are in 'Running' status. The 'test-mia' cluster's 'Action' menu is open, showing options like 'Delete', 'Add Existing Instance', 'Upgrade Cluster', 'Automatic Scaling', 'Addon Upgrade' (highlighted), and 'Deploy Istio'.

Cluster Name/ID	Cluster Type	Region (All)	Network Type	Cluster Status	Number of Nodes	Time Created	Kubernetes Version	Action
kubernetes-test	Kubernetes	China East 1 (Hangzhou)	VPC vpc-bp1fmdpc9b...	Running	6	12/18/2018,15:52:19	1.11.5	Manage View Logs Dashboard Scale Cluster More
k8s-delete-node	Kubernetes	China East 1 (Hangzhou)	VPC vpc-bp1nr6ohb0c...	Running	5	12/17/2018,10:26:28	1.11.5	Manage View Logs Dashboard Scale Cluster More
k8s-test	Kubernetes	China North 2 (Beijing)	VPC vpc-2zefxdn1wdl...	Running	6	12/05/2018,13:27:46	1.11.5	Manage View Logs Dashboard Scale Cluster More
k8s-managed-cluster	ManagedKubernetes	China East 1 (Hangzhou)	VPC vpc-bp1kd7yn4qn...	Running	3	11/01/2018,11:21:13	1.11.5	Manage View Logs Dashboard Scale Cluster More
test-mia	Kubernetes	China East 1 (Hangzhou)	VPC vpc-bp1kyevdj...	Running	7	09/17/2018,11:37:55	1.11.5	Manage View Logs Dashboard Scale Cluster More

4. Select the target system component, and click Upgrade in the Action column.

Upgrading is then displayed in the Status column.

Addon Upgrade

Component	Current Version	Upgradeable Version	Consistency Check	Action	Status
alicloud-application-controller	v0.1.0.1-f832bed-aliyun	v0.1.0.1-f832bed-aliyun	Success	Latest	
alicloud-disk-controller	v1.11.2.2-a390cfb-aliyun	v1.11.2.2-a390cfb-aliyun	Success	Latest	
Cloud Controller Manager Readme Version Information	v1.9.3.16-gcc144c7-aliyun	v1.9.3.16-gcc144c7-aliyun	Success	Latest	
flexvolume	v1.11.2.5-85c062f-aliyun	v1.11.2.32-af2d48c-aliyun	Success	<div>Upgrade</div>	
Nginx Ingress Controller Readme Version Information	v0.20.0.1-4597ce2-aliyun	v0.20.0.1-4597ce2-aliyun	Success	Latest	

Refresh

Close

Result

On the Addon Upgrade page, Latest is displayed in the Action column of the target system component.

Addon Upgrade ✕					
Component	Current Version	Upgradeable Version	Consistency Check	Action	Status
alicloud-application-controller	v0.1.0.1-f832bed-aliyun	v0.1.0.1-f832bed-aliyun	Success	Latest	
alicloud-disk-controller	v1.11.2.2-a390cfb-aliyun	v1.11.2.2-a390cfb-aliyun	Success	Latest	
Cloud Controller Manager Readme Version Information	v1.9.3.16-gcc144c7-aliyun	v1.9.3.16-gcc144c7-aliyun	Success	Latest	
flexvolume	v1.11.2.32-af2d48c-aliyun	v1.11.2.32-af2d48c-aliyun	Success	Latest	
Nginx Ingress Controller Readme Version Information	v0.20.0.1-4597ce2-aliyun	v0.20.0.1-4597ce2-aliyun	Success	Latest	
<div>Refresh Close</div>					

1.4.14 Update the Kubernetes cluster certificates that are about to expire

This topic describes how to update the Kubernetes cluster certificates that are about to expire through the Container Service console.

Prerequisites

You have created a Kubernetes cluster and the system has already prompted you to update the cluster certificates. For more information, see [Create a Kubernetes cluster](#).

Procedure

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Kubernetes, click Clusters.
3. Click Update Certificate on the right of the target cluster.






Note:

If cluster certificates are about to expire in about two months, the system displays the Update Certificate prompt for the cluster.

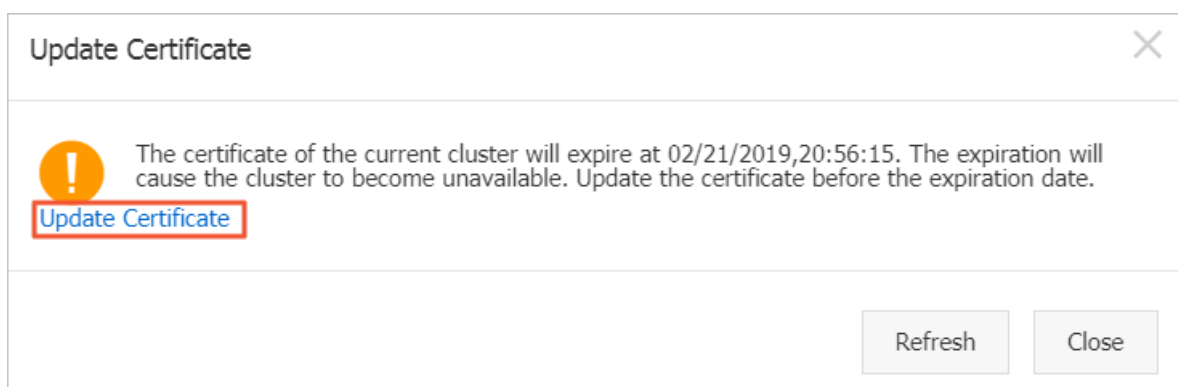
Cluster List

Help: [Create cluster](#) [Create GPU clusters](#) [Scale cluster](#) [Authorization management](#)

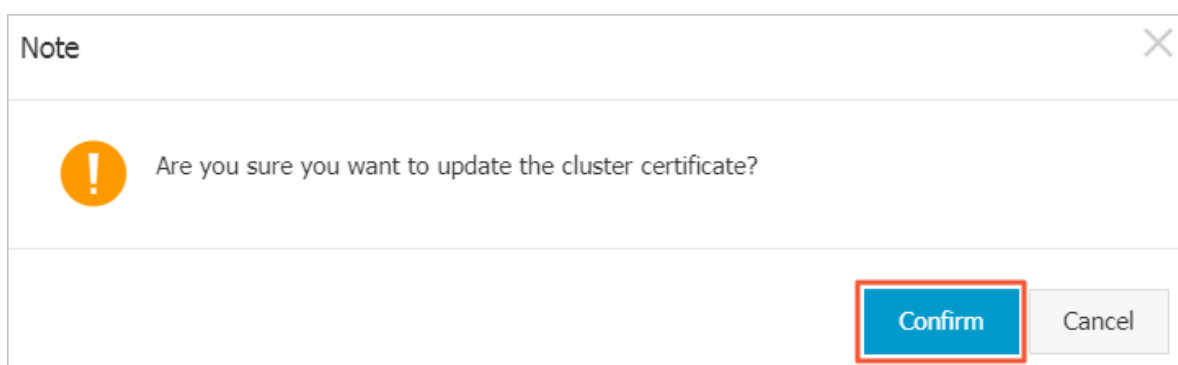
Name

Cluster Name/ID	Cluster Type	Region (All)
k8s-test 	Kubernetes	China North 2 (Beijing)
test-mia 	Kubernetes	China East 1 (Hangzhou)
kubernetes-test 	Kubernetes	China East 1 (Hangzhou)

4. Click Update Certificate.

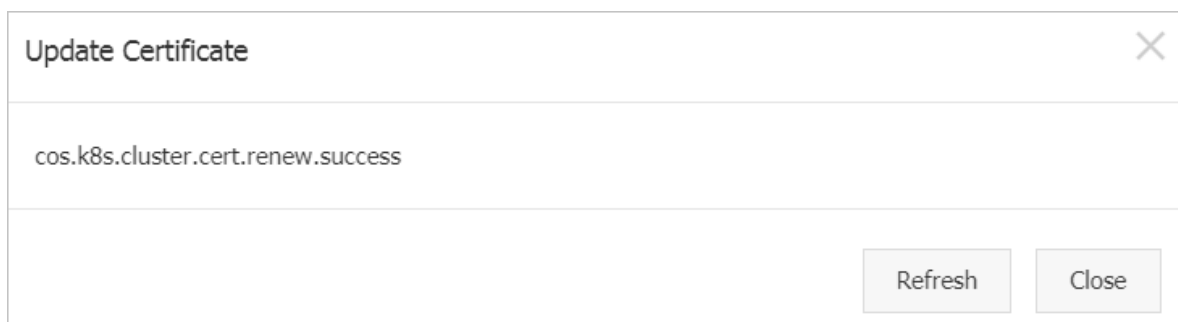


5. Click Confirm.



Result

- The Update Certificate page displays Success.



- On the Cluster List page, the Update Certificate prompt of the target cluster has been removed.

1.4.15 Scale a Kubernetes cluster

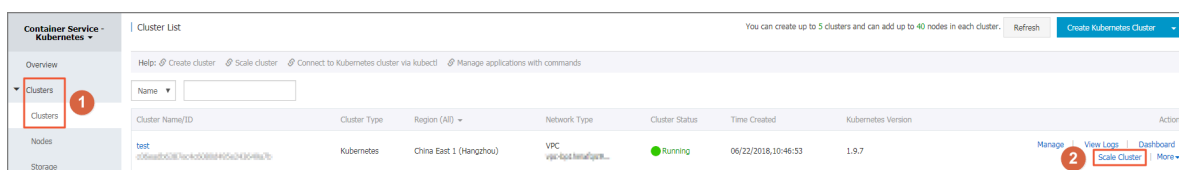
This topic describes how to scale a Kubernetes cluster (by increasing or decreasing the number of Worker nodes in the cluster) in the Alibaba Cloud Container Service for Kubernetes console.

Limits

- The number of Master nodes in any Kubernetes cluster cannot be changed.
- You can only decrease the number of Worker nodes that are added when you create or scale out the cluster. The number of Worker nodes that are added through the [Add an existing ECS instance](#) feature cannot be decreased.
- Worker nodes cannot be removed either through the `kubectl delete` command or through node removal operations in the Container Service console.
- When you scale in a cluster, Worker nodes are removed from the cluster in the sequence that they were added when you scaled out the cluster.
- You can scale in a cluster only if the cluster has more than one node that was created automatically.

Procedure

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Container Service-Kubernetes, choose Clusters > Clusters.
3. On the right of the target cluster, click Scale Cluster.



4. Click Scale out or Scale in and then set the number of Worker nodes.

In this example, the number of Worker nodes of the cluster is increased from 1 to 4.

The screenshot displays the scaling configuration interface for a Kubernetes cluster. The cluster name is 'k8s-test' and the region is 'China East 1 (Hangzhou) ZoneG'. The current number of existing nodes is 1. The 'Scale' section has two buttons: 'Scale out' (highlighted in blue) and 'Scale in' (greyed out). The 'Instance Type' is set to '2 Core(s) 4 G (ecs.n4.large)'. The 'Scaling Number' is set to 3, and the 'Number of workers after scaling' is 4. A red box highlights the 'Scaling Number' and 'Number of workers after scaling' fields. Below this, there is a 'Logon Password' field with a green border and a placeholder text 'Please enter the login password used when creating the cluster'. At the bottom, there is a 'Submit' button and a link for 'RDS Whitelist'.

5. Enter the node logon password.

6. Click Submit.

In the left-side navigation pane, choose Clusters > Node to verify that the current number of Worker nodes is changed to 4.

1.4.16 Autoscale a Kubernetes cluster

This topic describes how to autoscale a Kubernetes cluster to meet the requirements of your Kubernetes cluster workload. Alibaba Cloud Container Service for Kubernetes (ACK) provides the capability to autoscale a Kubernetes cluster through using the cluster autoscaler program.

Background information

You can set the the cluster autoscaler to add different ECS instance types to your Kubernetes cluster, such as the general, GPU, and preemptive instance types. You can set multiple zones, instance specifications, and autoscaling modes.

Cluster autoscaler overview

The cluster autoscaler changes the size of a Kubernetes cluster based on the use of resource in the nodes of a pod in a Kubernetes cluster. Resource usage is calculated based on the pod resource requests.

When a pod requests more resources than what the associated node can provide, the pod enters the pending status. At which time, the autoscaler calculates the change to the size of cluster. It does so by calculating the number of nodes necessary to provide the requested resource with regard to the resource specification and threshold that you set for an autoscaling group.

For example, if you set a low threshold value for the number of nodes in an autoscaling group, the cluster autoscaler deletes a node, which reduces the amount of resources that the pod can request.

Notes

- By default, your account can use up to 30 Pay-As-You-Go ECS instances in all your clusters, and the route table of one VPC can contain up to 50 entries. To increase the number of available ECS instances or entries in a route table of one VPC, open a ticket.
- For a single type of ECS instances, the number of ECS instances of one specification that is permitted at one time varies frequently. Therefore, we recommend that you set multiple ECS instance types of one ECS specification.
- When a node for which you set the fast scaling mode is shut down and reclaimed, it is in the `NotReady` status. When the node is reused by the cluster autoscaler, the node enters the `Ready` status.
- When a node for which you set the fast scaling mode is shut down and reclaimed, only the disks attached to the node are charged (except that the node uses local disks, for example, `ecs.d1ne.2xlarge`).

Enable cluster autoscaling

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

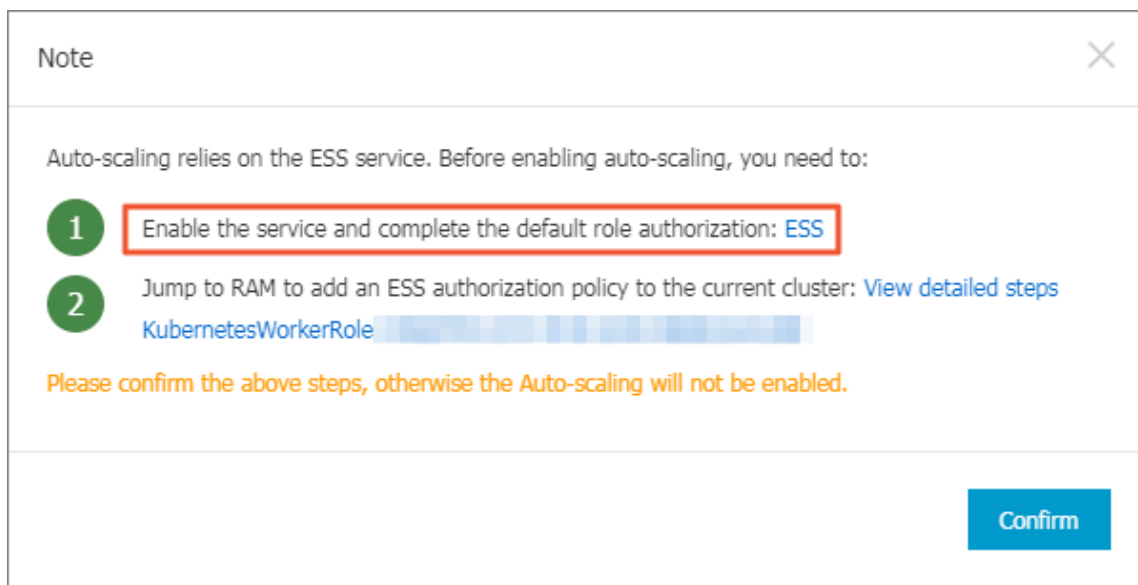
3. Find the target cluster. Then, in the Action column, choose More > Auto Scaling.

Cluster Name/ID	Cluster Type	Region (All) ▾	Network Type	Cluster Status	Time Created	Kubernetes Version	Action
mymanagedk8s1	ManagedKubernetes	China North 2 (Beijing)	VPC vpc-zzetk0nc8jr...	Running	09/07/2018,16:01:46	1.10.4	Manage View Logs Dashboard Scale Cluster More ▾
myserverlessk8s1	Serverless Kubernetes	China East 1 (Hangzhou)	VPC vpc-bp1bw6x01cy...	Running	09/04/2018,20:13:20	1.9.7	Manage View Logs Delete
mytest1	Kubernetes	China North 2 (Beijing)	VPC vpc-zze7tc50cju...	Running	09/03/2018,10:06:49	1.10.4	Manage View Logs Dashboard Scale Cluster More ▾
myk8s1	Kubernetes	China East 1 (Hangzhou)	VPC vpc-bp1faadn5o...	Running	08/26/2018,11:53:17	1.10.4	Manage Delete Add Existing Instance Upgrade Cluster Auto-scaling Addon Upgrade Upgrade monitoring service Deploy Istio

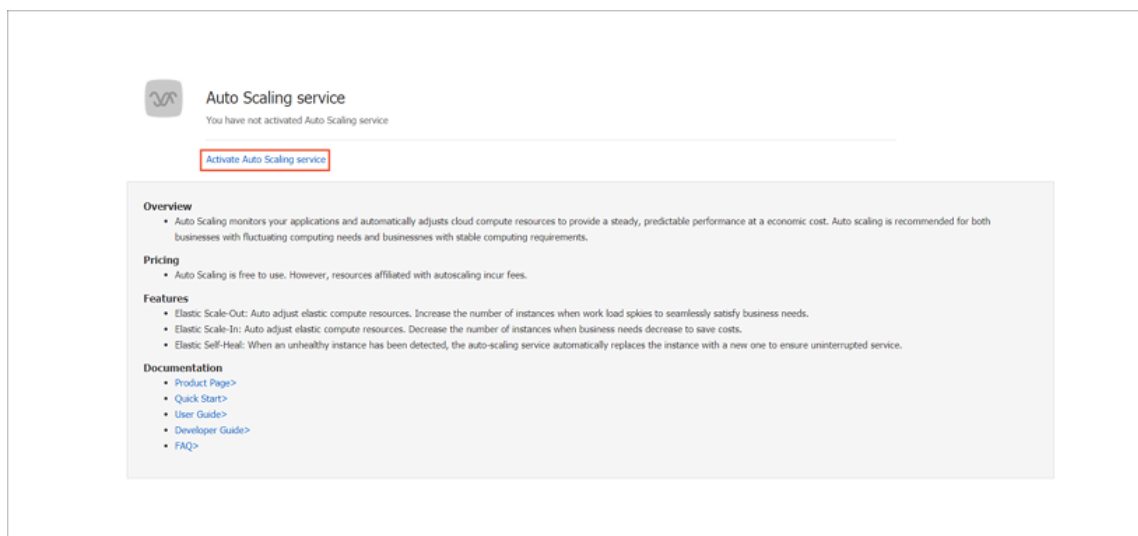
Grant required permissions for the Auto Scaling service and the cluster

- Activate the Auto Scaling service

1. Click ESS in the displayed dialog box.



2. Click Activate Auto Scaling service.



3. Read and confirm that you agree to the conditions by selecting the I agree with Auto Scaling Service Agreement of Service check box, and then click Enable Now.

Enable Service

Auto Scaling Service

Basic Configuration

Product

Auto Scaling Service

☒ I agree with [Auto Scaling Service Agreement of Service](#)

Enable Now

4. Click Console.

Confirm Order

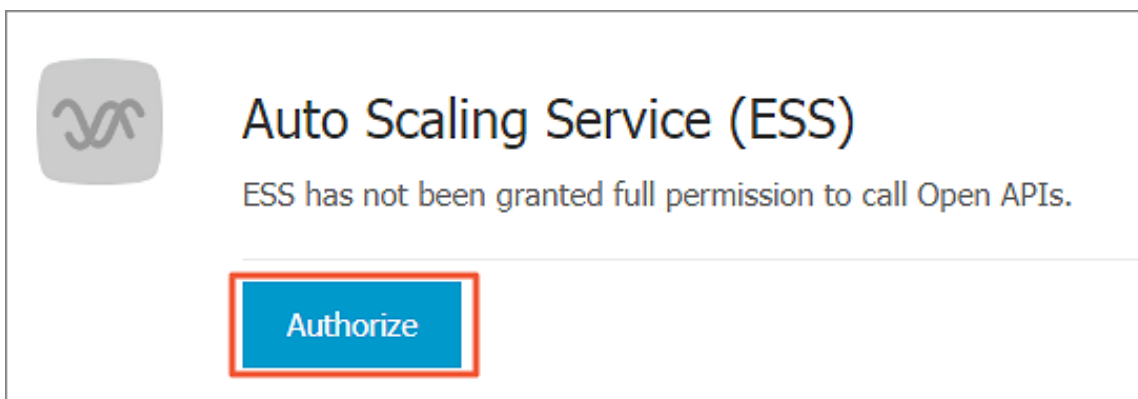
Order completed

The service you ordered is being opened. Please wait for a few minutes.

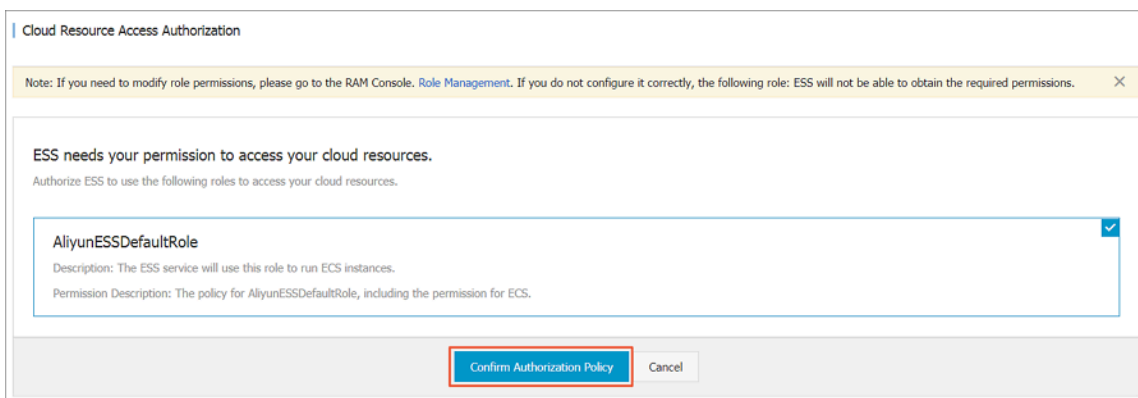
[Console](#)

[Purchase history](#)

5. Click Authorize.



6. Click Confirm Authorization Policy to grant ESS the permission to access your cloud resources.



Verify the result

If the page automatically redirects to the Auto Scaling console, the activation is successful.

• Add ESS authorization policies to the cluster

1. Click the Worker RAM role (`Kubernetes WorkerRole [xxx]`) in the following dialog box.



You need to use the primary account to log on to the console before perform this operation.

Note

Auto-scaling relies on the ESS service. Before enabling auto-scaling, you need to:

- 1 Enable the service and complete the default role authorization: [ESS](#)
- 2 Jump to RAM to add an ESS authorization policy to the current cluster: [View detailed steps](#)
KubernetesWorkerRole

Please confirm the above steps, otherwise the Auto-scaling will not be enabled.

Confirm

2. Click View Permissions on the right of the target authorization policy.

KubernetesWorkerRole-827d8853-3009-40c9-add1-3f378613c5d1

Edit Authorization Policy

Authorization Policy Name	Description	Type	Actions
k8sWorkerRolePolicy-827d8853-3009-40c9-add1-3f378613c5d1		Custom	View Permissions Revoke Authorization

3. In the upper-right corner of the page, click Modify Authorization Policy.

K8sWorkerRolePolicy-827d8853-3009-40c9-add1-3f378613c5d1 (Custom)

Authorization Policy...

Versions

References

Policy Details

Name k8sWorkerRolePolicy-827d8853-3009-40c9-add1-3f378613c5d1 Type Custom

Description

```

1 {
2   "Version": "1",
3   "Statement": [
4     {
5       "Action": [
6         "ecs:AttachDisk",
7         "ecs:DetachDisk",
8         "ecs:DescribeDisks",
9         "ecs:CreateDisk",
10        "ecs:CreateSnapshot",
11        "ecs>DeleteDisk",
12        "ecs>CreateNetworkInterface",
13        "ecs:DescribeNetworkInterfaces",
14        "ecs:AttachNetworkInterface",
15        "ecs:DetachNetworkInterface",
16        "ecs>DeleteNetworkInterface",

```

More information about the authorization policy formulation.

4. In the Action field of the Policy Document area, add the following policies:

```

"ess : Describe *",
"ess : CreateScalingRule ",
"ess : ModifyScalingGroup ",

```

```
" ess : RemoveInstances ",
" ess : ExecuteScalingRule ",
" ess : ModifyScalingRule ",
" ess : DeleteScalingRule ",
" ecs : DescribeInstancesType s ",
" ess : DetachInstances "
```



Note:

You must add a comma (,) to the end of the last line in the **Action** field before adding these policies.

5. Click Modify Authorization.

Set cluster auto scaling parameters

1. On the Automatic Scaling page, set the following parameters.


Configuration	Description
Cluster	Target cluster name.
Scale-in Threshold	The ratio of the amount of resources requested by the cluster workload to the amount of cluster resources. When the amount of resources requested by the cluster workload is less than or equal to the threshold, the system automatically scales in the cluster. The default value is 50%.
Defer Scale-in For	The number of minutes for which the system must wait to automatically scale in the cluster after the scale-in threshold is reached. The default value is 10 minutes.
Cool-Down Time	The period (in minutes) during which the system does not automatically scale in or scale out a cluster after the number of cluster nodes increases or decreases. The default is 10 minutes.


- Click **Create** on the right of the target type of resource (which can be CPU or GPU) that you want to autoscale.

On the **Scaling Group Configuration** page, set the following parameters to create a scaling group:

Configuration	Description
Region	The region to which the scaling group is deployed . You must ensure that the scaling group and the cluster where it is located share the same region. This region cannot be modified.
Zone	The zone where the scaling group is created.
VPC	The network where the scaling group is created. You must ensure that the scaling group and the cluster where it is located are in the same region.

Set worker nodes.

Configuration	Description
Instance Type	Set the specifications of instances in the scaling group .
System Disk	Set the system disks of the scaling group.
Attach Data Disk	Mount a data disk when you create a scaling group. By default, no data disk is mounted.
Instance Quantity	Set the number of instances in the scaling group. <div>  Note: <ul style="list-style-type: none"> The number does not include the existing instances in the cluster. By default, this parameter value is 0, and the cluster adds instances to the scaling group and the Kubernetes cluster where the scaling group is located when this parameter exceeds 0. </div>

Configuration	Description
Key Pair	<p>Set the key pair used to log on to the node added through autoscaling. You can create a new key pair in the Elastic Compute Service (ECS) console.</p> <div>  Note: Only key pair logon is supported. </div>
RDS Whitelist	Set the Relational Database Service (RDS) instances that can be accessed by the node added through autoscaling.

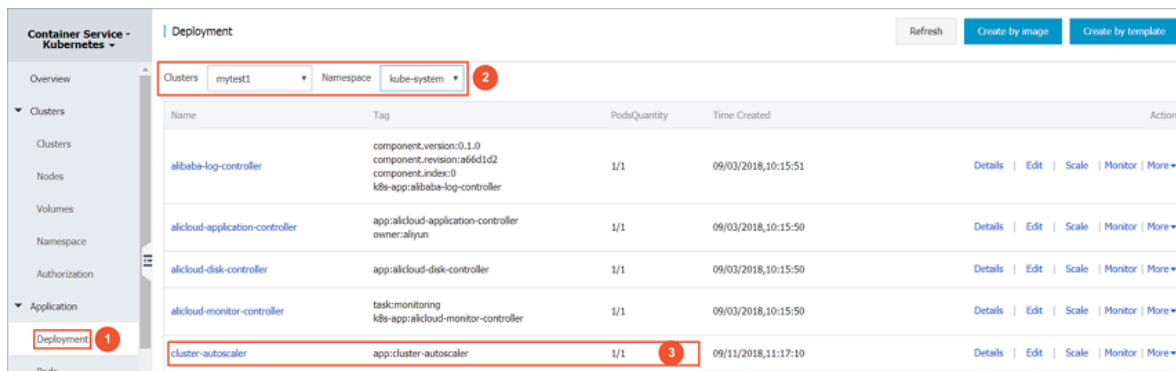
3. Click OK.

Verify the result

- You can directly verify that a scaling group under CPU is displayed on the Auto Scaling page.

Auto-scaling Back to Cluster List				
Overview				
Cluster ID: c68b3c5a4273c4fba8151f06276bcc224		Cluster Name: mytest1		
Shrinkage Threshold: 50%		Shrinkage Trigger Delay: 10Min.		
Scaling Group	Status	Instance Type	Total Instance Number	Min Instance Number
CPU asg-2ze9npi7u9efzddlmrzc	Active	ecs.sn1ne.xlarge	0	0
GPU				

- To verify the created autoscaling component, follow these steps:
 1. In the left-side navigation pane, choose Application > Deployment.
 2. Select the target cluster and the kube-system namespace to view the created component named cluster-autoscaler.



Troubleshooting

- If the cluster autoscaler cannot add nodes to a pod that requested more resources, you can perform the following checks:
 - Make sure that the amount of resources provided by the ECS instances that you set for the autoscaling group is greater than the amount of resources requested by the pod.
 - Make sure that you have granted the required permissions by following the preceding steps. You must grant the required permissions for each target cluster .
 - Make sure that the target Kubernetes cluster are connected to the Internet. The cluster autoscaler calls an API action from Alibaba Cloud, therefore you must ensure that the cluster nodes can be accessed through the Internet.

- If the cluster autoscaler cannot delete nodes from the autoscaling group, you can perform the following checks:
 - Make sure that the resource request threshold of pods on all nodes is not greater than that of used to scale in the cluster.
 - Make sure that no node runs the pods that belong to the `kube-system` namespace.
 - Make sure that no node runs the pod for which any constrained scheduling policies are set. This is because a constrained scheduling policy can limit a pod to a fixed node.
 - Make sure that no pod contains a `PodDisruptionBudget` object that has reached the minimum value allowed. For more information, see [How do Disruption Budgets work](#). For more information, see [How do Disruption Budgets work](#).

For more information, see [Cluster autoscaler](#).

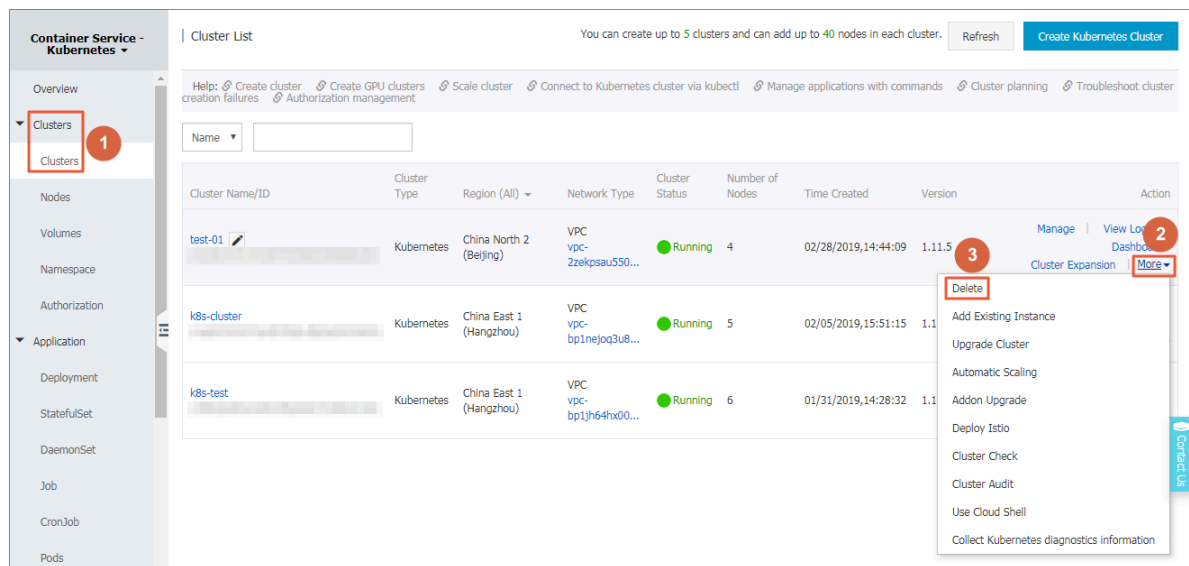
1.4.17 Delete a Kubernetes cluster

This topic describes how to delete a Kubernetes cluster in the Alibaba Cloud Container Service for Kubernetes console.

Procedure

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

3. On the right of the target cluster, choose More > Delete .



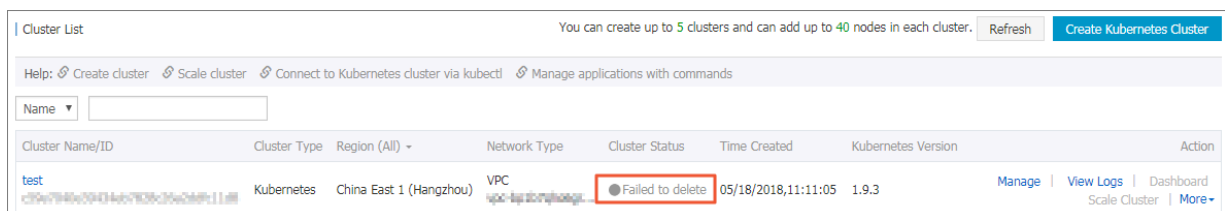
What's next

Troubleshoot a cluster deletion failure

If you manually add cloud resources into resources created by Resource Orchestration Service (ROS), ROS does not have the permission to delete the manually added resources. For example, if you manually add a VSwitch in the Virtual Private Cloud (VPC) created by ROS, ROS cannot delete this VPC when you delete the Kubernetes cluster. As a result, the cluster deletion will fail.

To solve this issue, Container Service allows you to forcibly delete the cluster. Specifically, you can delete the cluster record and the ROS stack if the cluster fails to be deleted. However, you must release the created resources manually.

If a cluster fails to be deleted, the cluster status is displayed as follows.

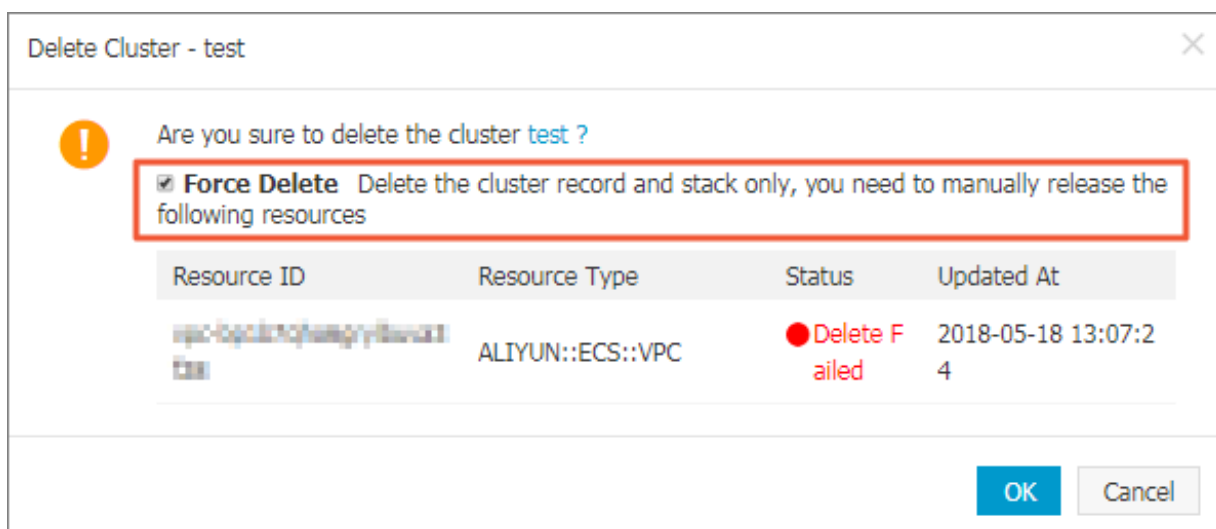


On the right of the cluster that failed to be deleted, choose More > Delete . In the displayed dialog box, select the Force Delete check box, and then click OK.



Note:

You must manually release the resources that failed to be deleted. For information, see [Failed to delete Kubernetes clusters: ROS stack cannot be deleted](#).



1.4.18 Upgrade the Heapster components to the metrics-server component

This topic describes how to upgrade the Heapster components to the metrics-server component without upgrading the target Kubernetes cluster.

Prerequisites

- A Kubernetes cluster is created. For more information, see [Create a Kubernetes cluster](#).
- The Kubernetes cluster version is earlier than v1.12.

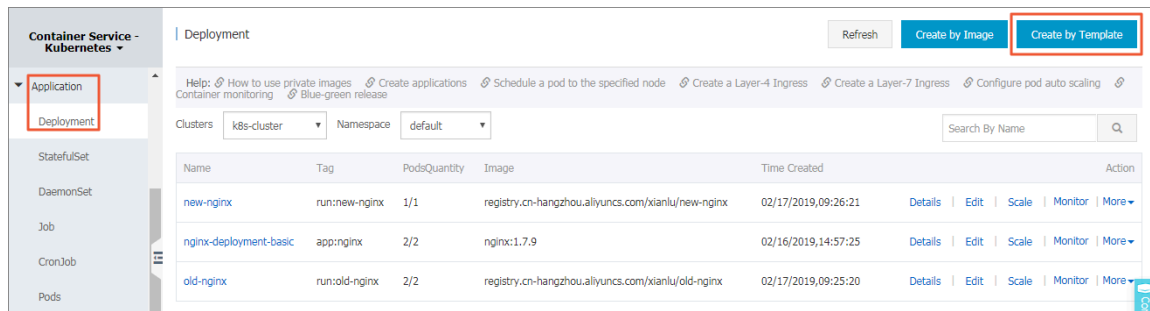
Procedure

You must reset the data collection component, reset the monitoring data link, and then modify component compatibility settings.

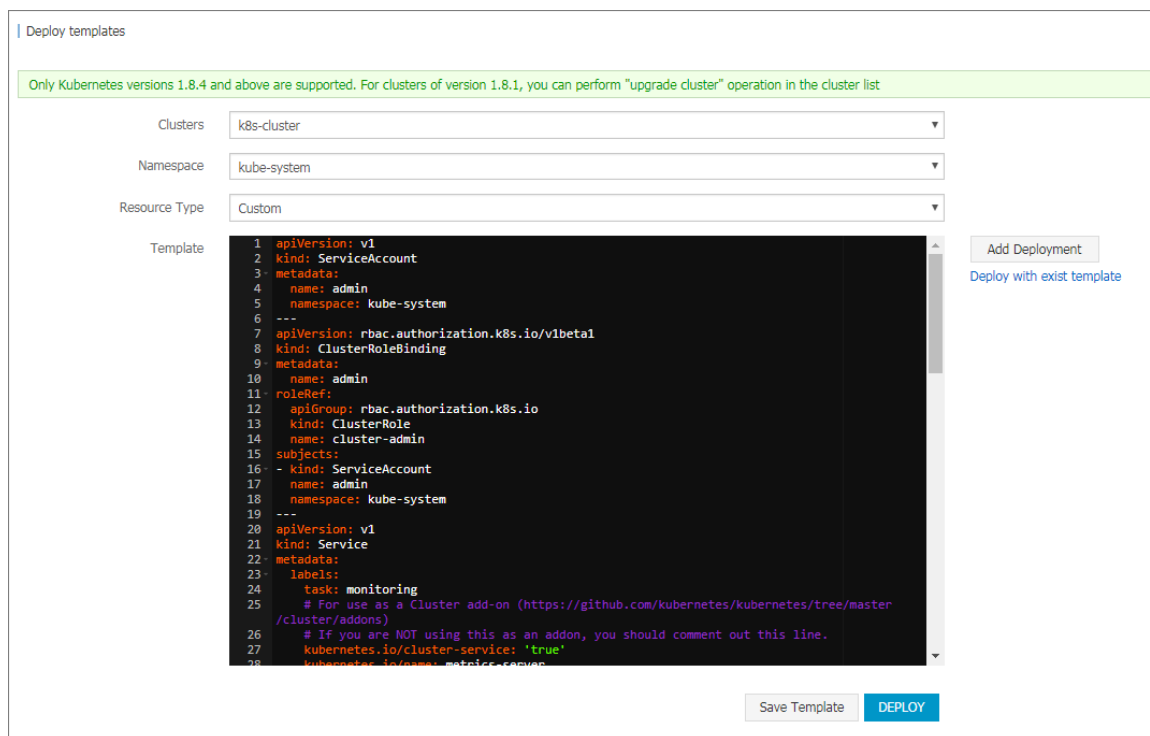
- Reset the data collection component.

To change the data collection component from Heapster to metrics-server, follow these steps:

- Log on to the [Container Service console](#).
- In the left-side navigation pane under Container Service-Kubernetes, choose **Application > Deployment**.



- In the upper-right corner, click **Create by Template**.



- Select the target cluster from the Clusters drop-down list, and then select kube-system from the Namespace drop-down list.
- Select Custom from the Resource Type, copy the following code and paste it to the Template area, then click **DEPLOY**.

```
apiVersion : v1
kind : ServiceAccount
metadata :
  name : admin
```

```

    namespace : kube - system
---
apiVersion : rbac . authorizat ion . k8s . io / v1beta1
kind : ClusterRol eBinding
metadata :
  name : admin
roleRef :
  apiGroup : rbac . authorizat ion . k8s . io
  kind : ClusterRol e
  name : cluster - admin
subjects :
- kind : ServiceAcc ount
  name : admin
  namespace : kube - system
---
apiVersion : v1
kind : Service
metadata :
  labels :
    task : monitoring
    # For use as a Cluster add - on ( https :// github
    . com / kubernetes / kubernetes / tree / master / cluster /
    addons )
    # If you are NOT using this as an addon ,
    you should comment out this line .
    kubernetes . io / cluster - service : ' true '
    kubernetes . io / name : metrics - server
  name : heapster
  namespace : kube - system
spec :
  ports :
  - port : 80
    targetPort : 8082
  selector :
    k8s - app : metrics - server
---
apiVersion : v1
kind : Service
metadata :
  name : metrics - server
  namespace : kube - system
  labels :
    kubernetes . io / name : metrics - server
spec :
  selector :
    k8s - app : metrics - server
  ports :
  - port : 443
    protocol : TCP
    targetPort : 443
---
apiVersion : apiregistr ation . k8s . io / v1beta1
kind : APIService
metadata :
  name : v1beta1 . metrics . k8s . io
spec :
  service :
    name : metrics - server
    namespace : kube - system
  group : metrics . k8s . io
  version : v1beta1
  insecureSk ipTLSVerif y : true
  groupPrior ityMinimum : 100
  versionPri ority : 100

```

```

---
apiVersion : extensions / v1beta1
kind : Deployment
metadata :
  name : metrics - server
  namespace : kube - system
  labels :
    k8s - app : metrics - server
spec :
  selector :
    matchLabels :
      k8s - app : metrics - server
  template :
    metadata :
      name : metrics - server
      labels :
        k8s - app : metrics - server
    spec :
      serviceAccountName : admin
      containers :
        - name : metrics - server
          image : registry .## REGION ##. aliyuncs . com / acs /
metrics - server : v0 . 2 . 1 - 9dd9511 - aliyun
          imagePullPolicy : Always
          command :
            - / metrics - server
            - '-- source = kubernetes : https :// kubernetes . default
,
            - '-- sink = socket : tcp :// monitor . csk .## REGION ##.
aliyuncs . com : 8093 ? clusterId =## CLUSTER_ID ##& public =
true '

```

**Note:**

You need to replace `## REGION ##` with the region (for example, China East 1:cn-hangzhou) to which your target cluster is located, and replace `## CLUSTER_ID ##` with your target cluster ID.

- Reset the monitoring data link.

- In the left-side navigation pane, choose Clusters > Nodes.
- Select the target Cluster.
- Click the instance ID of one Master node. In this step, the master-01 node is used.

Node List

Refresh

Label Management

Cluster Expansion

Add Existing Instance

Help: [Postpay instance to Prepay](#) [Node exception](#) [Node monitoring and alarms](#) [Collect Kubernetes diagnostics information](#)

Clusters

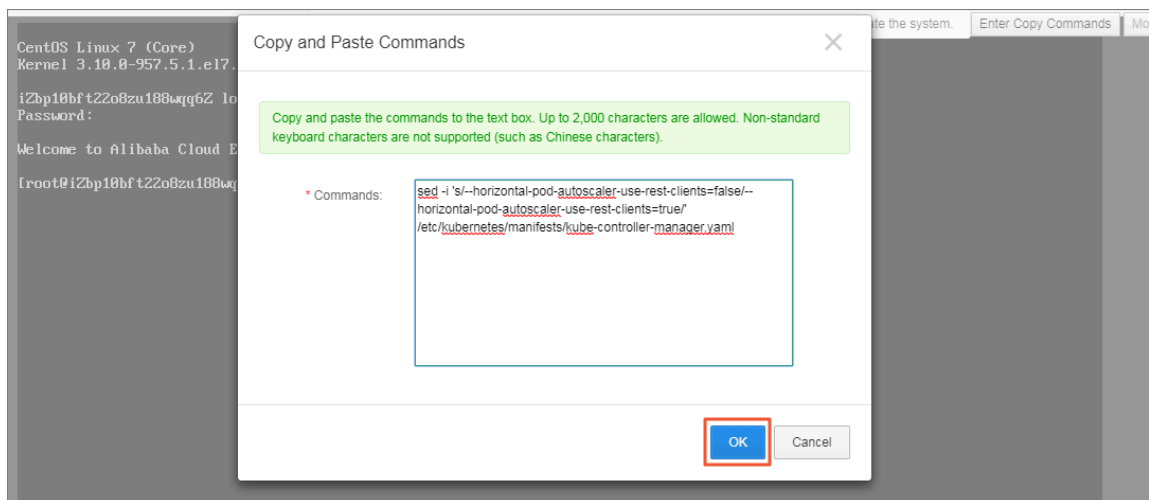
k8s-cluster

Filter by Labels

<input type="checkbox"/>	Instance ID/Name/IP Address	Role/Status	Configuration	Pods (Allocated)	CPU Request/Limit	CPU Usage	Memory Request/Limit	Memory Usage	Kubelet Version	Update Time	Action
<input type="checkbox"/>	master-01- 	Master <div>Running</div>	Pay-As-You-Go ecs.n1.large	14	30.00% 0.00%	<div></div> 5.15 %	8.90% 14.73%	<div></div> 53.66 %	v1.11.5	02/05/2019,15:58:07	Monitor More

- Click Connect. On the displayed page, enter the remote connection password and click OK. After you log on to the ECS instance, run the following command:

```
sed -i 's/--horizontal-pod-autoscaler-use-rest-clients=false/--horizontal-pod-autoscaler-use-rest-clients=true/' /etc/kubernetes/manifests/kube-controller-manager.yaml
```

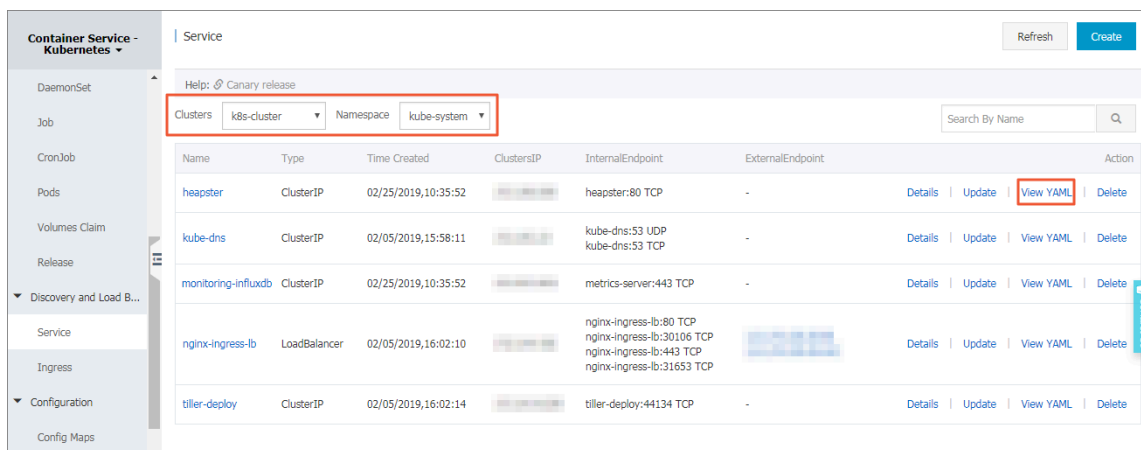


- Repeat step c to step d on the master-02 and master-03 nodes.

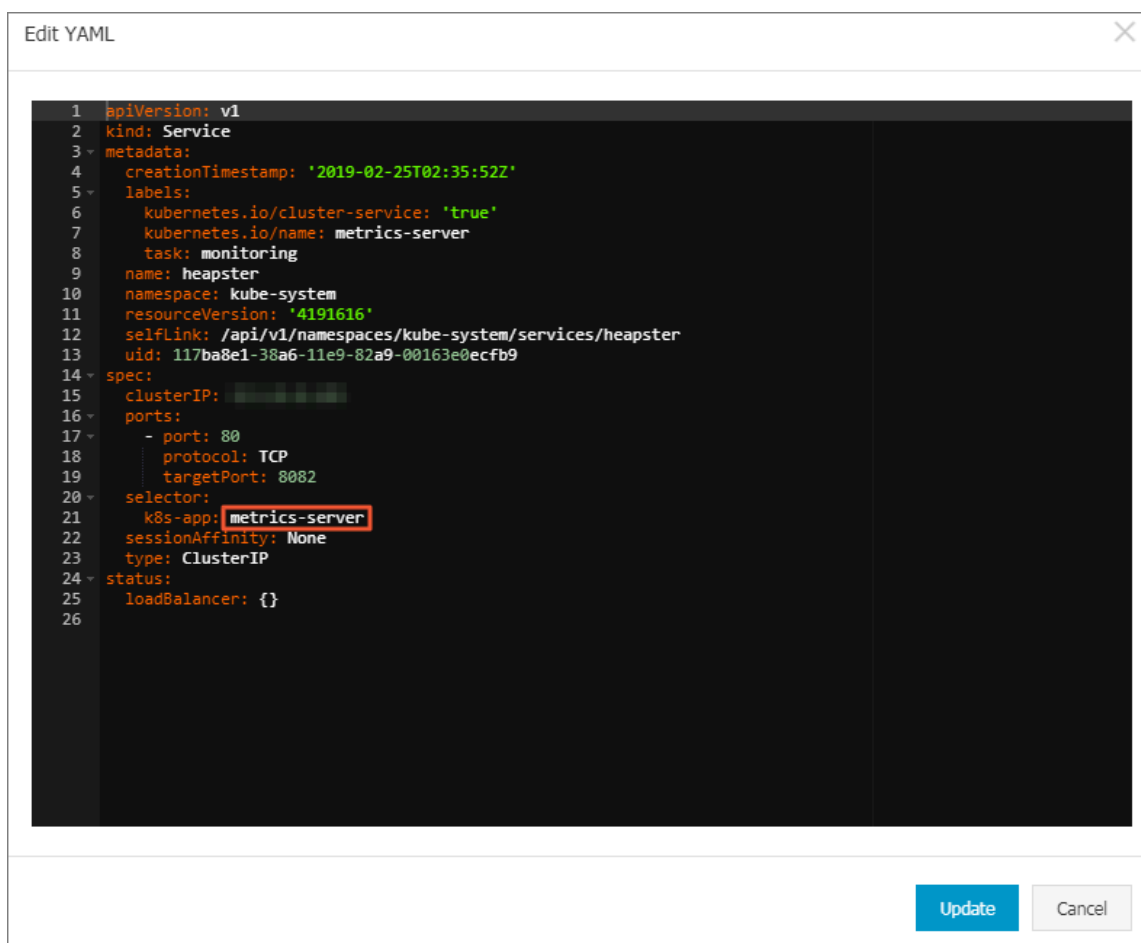
After you complete the preceding operations, kubelet automatically starts and updates the kube-controller-manager component.

- Reset the component compatibility settings.

- In the left-side navigation pane, choose Discovery and Load Balancing > Service.
- Select the target Cluster and the kube-system namespace. Then click View YAML on the right of heapster.



- In the displayed dialog box, reset the k8s-app parameter of the selector field to metrics-server. Then click Update.



- In the left-side navigation pane, choose Application > Deployment.
- Select the target Cluster and the kube-system Namespace.

Container Service - Kubernetes Overview Clusters Nodes Volumes Namespace Authorization Application Deployment StatefulSet DaemonSet Job CronJob Pods Volumes Claim Release Discovery and Load Balancing Service Ingress Configuration Config Maps	Deployment						Refresh	Create by Image	Create by Template
	Help: How to use private images Create applications Schedule a pod to the specified node Create a Layer-4 Ingress Create a Layer-7 Ingress Configure pod auto scaling Container monitoring Blue-green release								
	Clusters: kube-cluster		Namespace: kube-system		Search By Name				
	Name	Tag	PodsQuantity	Image	Time Created	Action			
	alibaba-log-controller	component.version:0.1.3 component.revision:527ff4d component.index:0 k8s-app:alibaba-log-controller	1/1	registry-vpc.cn-beijing.aliyuncs.com/acs/log-controller:0.1.3.0-527ff4d-aliyun	02/28/2019,14:55:37	Details Edit Scale Monitor More			
	alicloud-application-controller	app:alicloud-application-controller owner:aliyun	1/1	registry-vpc.cn-beijing.aliyuncs.com/acs/aliyun-app-lifecycle-manager:v0.1.0.1-f832bed-aliyun	02/28/2019,14:55:37	Details Edit Scale Monitor More			
	alicloud-disk-controller	app:alicloud-disk-controller	1/1	registry-vpc.cn-beijing.aliyuncs.com/acs/alicloud-disk-controller:v1.11.2.2-a390cfb-aliyun	02/28/2019,14:55:37	Details Edit Scale Monitor More			
	alicloud-monitor-controller	task:monitoring k8s-app:alicloud-monitor-controller	1/1	registry-vpc.cn-beijing.aliyuncs.com/acs/alicloud-monitor-controller:v1.0.0	02/28/2019,14:55:37	Details Edit Scale Monitor More			
	aliyun-acr-credential-helper	app:aliyun-acr-credential-helper	1/1	registry-vpc.cn-beijing.aliyuncs.com/acs/aliyun-acr-credential-helper:v18.10.29.0-1a28f02-aliyun	02/28/2019,14:55:37	Details Edit Scale Monitor More			
	coredns	k8s-app:kube-dns	2/2	registry-vpc.cn-beijing.aliyuncs.com/acs/coredns:1.1.3	02/28/2019,14:50:57	Details Edit Scale Monitor More			

6. Delete the components related to Heapster. In this example, you need to delete the heapster and monitoring-influxdb components.

- On the right of the heapster component, choose More > Delete. In the displayed dialog box, click OK.
- On the right of the monitoring-influxdb component, choose More > Delete. In the displayed dialog box, select the Delete associated services monitoring-influxdb check box, and then click OK.

Deployment						Refresh	Create by Image	Create by Template
Help: How to use private images Create applications Schedule a pod to the specified node Create a Layer-4 Ingress Create a Layer-7 Ingress Configure pod auto scaling Container monitoring Blue-green release								
Clusters: kube-cluster		Namespace: kube-system		Search By Name				
Name	Tag	PodsQuantity	Image	Time Created	Action			
alibaba-log-controller	component.version:0.1.3 component.revision:527ff4d component.index:0 k8s-app:alibaba-log-controller	1/1	registry-vpc.cn-beijing.aliyuncs.com/acs/log-controller:0.1.3.0-527ff4d-aliyun	02/28/2019,14:55:37	Details Edit Scale Monitor More			
alicloud-application-controller	app:alicloud-application-controller owner:aliyun	1/1	registry-vpc.cn-beijing.aliyuncs.com/acs/aliyun-app-lifecycle-manager:v0.1.0.1-f832bed-aliyun	02/28/2019,14:55:37	Details Edit Scale Monitor More			
alicloud-disk-controller	app:alicloud-disk-controller	1/1	registry-vpc.cn-beijing.aliyuncs.com/acs/alicloud-disk-controller:v1.11.2.2-a390cfb-aliyun	02/28/2019,14:55:37	Details Edit Scale Monitor More			
alicloud-monitor-controller	task:monitoring k8s-app:alicloud-monitor-controller	1/1	registry-vpc.cn-beijing.aliyuncs.com/acs/alicloud-monitor-controller:v1.0.0	02/28/2019,14:55:37	Details Edit Scale Monitor More			
aliyun-acr-credential-helper	app:aliyun-acr-credential-helper	1/1	registry-vpc.cn-beijing.aliyuncs.com/acs/aliyun-acr-credential-helper:v18.10.29.0-1a28f02-aliyun	02/28/2019,14:55:37	Details Edit Scale Monitor More			
coredns	k8s-app:kube-dns	2/2	registry-vpc.cn-beijing.aliyuncs.com/acs/coredns:1.1.3	02/28/2019,14:50:57	Details Edit Scale Monitor More			
heapster	task:monitoring k8s-app:heapster	1/1	registry-vpc.cn-beijing.aliyuncs.com/acs/heapster-amd64:v1.5.1.1	02/28/2019,14:55:36	Details Edit Scale Monitor More			
monitoring-influxdb	task:monitoring k8s-app:influxdb	1/1	registry-vpc.cn-beijing.aliyuncs.com/acs/heapster-influxdb-amd64:v1.1.1	02/28/2019,14:55:36	Details Edit Scale Monitor More			
nginx-ingress-controller	app:ingress-nginx	2/2	registry-vpc.cn-beijing.aliyuncs.com/acs/aliyun-ingress-controller:v0.20.0.1-4597ce2-aliyun	02/28/2019,14:55:37	Details Edit Scale Monitor More			

7. Check the new data link status.

Data link initialization takes about three minutes.

In the left-side navigation pane, choose **Application > Pods**. You can verify that the CPU and memory columns show normal values. This means the data link has been reset.



Note:

If both the displayed CPU and memory columns of each component are 0, it indicates an exception.

Container Service - Kubernetes		Pod								Refresh
		Clusters	k8s-cluster	Namespace	kube-system	Search By Name				
		Name	Status	max attempts	Pod IP	Node	Time Created	CPU	Memory	
		alibaba-log-controller-888499bc4-9qdtx	Running	1			02/05/2019,16:02:12	0.002	35.059 Mi	Details More
		alicloud-application-controller-b76c68764-xc2vk	Running	1			02/05/2019,16:02:10	0.001	36.574 Mi	Details More
		alicloud-disk-controller-84d64c55cd-6kn7	Running	1			02/05/2019,16:02:11	0	52.039 Mi	Details More
		alicloud-monitor-controller-84f86d5cd5-rfmcd	Running	1			02/05/2019,16:02:11	0.004	43.082 Mi	Details More
		aliyun-acr-credential-helper-56764b6dc-p99rj	Running	0			02/05/2019,16:02:11	0	36.672 Mi	Details More
		cloud-controller-manager-2f49v	Running	1			02/05/2019,16:02:10	0	41.383 Mi	Details More
		cloud-controller-manager-5hppl	Running	0			02/05/2019,16:02:42	0	8.621 Mi	Details More

1.4.19 Create a Kubernetes cluster that supports Windows

This topic describes how to use the Container Service console to create a Kubernetes cluster that supports Windows.

Prerequisites

The following services are activated: Container Service, Resource Orchestration Service (ROS), Resource Access Management (RAM), and Auto Scaling service.

Log on to the [Container Service console](#), [ROS console](#), [RAM console](#), and [Auto Scaling console](#) to activate the corresponding services.



Note:

The deployment of Kubernetes clusters that support Windows in Container Service is dependent on the application deployment capabilities of Alibaba Cloud ROS.

Therefore, you must activate ROS before you create a Kubernetes cluster that supports Windows.

Limits

- An SLB instance created with the cluster supports only the Pay-As-You-Go billing method.
- Kubernetes clusters that support Windows support only the VPC network type.
- Each account has a set quota of resources that can be used to create clusters. The default numbers are as follows:
 - Each account can create up to 100 security groups.
 - Each account can create up to 60 SLB instances of the Pay-As-You-Go billing method.
 - Each account can create up to 20 EIPs.



Note:

If any of the preceding quotas are exceeded when you create a cluster, the cluster fails to be created. In the case that you need a larger quota for any of these resources, you can open a ticket.

Procedure

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Container Service-Kubernetes, choose Clusters > Clusters. Then click Create Kubernetes Cluster in the upper-right corner.
3. On the displayed page, find Windows Cluster, and then click Create.
4. Enter the cluster name.



Note:

The cluster name must be 1 to 63 characters in length, and can contain letters, numbers, Chinese characters, and hyphens (-).

* Cluster Name

k8s-windows-cluster

The cluster name should be 1-63 characters long, and can contain numbers, Chinese characters, English letters and hyphens.

5. Select the region and zone where you want to locate the cluster.

Region	China North 2 (Beijing)	China North 3 (Zhangjiakou)	China North 5 (Huhehaote)	China East 1 (Hangzhou)	China East 2 (Shanghai)
	China South 1 (Shenzhen)	Hong Kong	Asia Pacific SE 1 (Singapore)	Asia Pacific SE 3 (Kuala Lumpur)	Asia Pacific SE 5 (Jakarta)
	Asia Pacific SOU 1 (Mumbai)	US East 1 (Virginia)	US West 1 (Silicon Valley)	EU Central 1 (Frankfurt)	
Zone	China East 1 Zone G				

6. Select a VPC.



Note:

Kubernetes clusters support only the VPC network type.

VPC: You can click Auto Create to create a VPC together with the Kubernetes cluster, or click Use Existing to use an existing VPC. If you click Use Existing, you can select a VPC and a VSwitch from the two displayed drop-down lists.

- If you click Auto Create, the system automatically creates a NAT gateway for a VPC when you create a cluster.
- If you click Use Existing and the selected VPC is associated with an existing NAT gateway, the system then uses the preexisting NAT gateway. If the selected VPC is not associated with any existing NAT gateway, the system automatically creates a NAT gateway for the selected VPC. Furthermore, if you do not want the system to automatically create a NAT gateway for the selected VPC, clear the Configure SNAT for VPC check box.



Note:

If you set the system not to create a NAT gateway automatically, then you need to manually set a NAT gateway, or set an SNAT entry to ensure that your selected

VPC is accessible to the Internet. Otherwise, instances in the VPC cannot access the Internet, which results in a cluster creation failure.

VPC

Auto Create
Use Existing

vpc-k8s-for-cs-c709888d6aa0c44af8de3817c78706e0...

(vsw-bp149pxcw4v9tdwzclr83) ZoneG

7. Set the node type by selecting a node billing method.



Note:

Pay-As-You-Go and Subscription are supported.

Node Type

Pay-As-You-Go
Subscription

View the differences between the two billing methods [Billing method comparison](#)

8. Set the Worker instance.



Note:

- Each cluster must contain at least 2 Worker nodes.
- Each cluster can contain up to 48 Worker nodes. To create more Worker nodes, open a ticket.
- By default, a system disk is attached to Worker nodes. For the system disk, available system disk types are SSD disks and Ultra disks.

- You can manually attach a data disk to Worker nodes. For data disks, available data types are SSD disks and Ultra disks.

Instance Configuration			
Instance Type	x86-Architecture Heterogeneous Computing ECS Bare Metal Instance Super Computing Cluster		
	4 Core(s) 8 G (ecs.c5.xlarge)		
Quantity	3 unit(s)		
System Disk	Ultra Disk	120 GiB	
<input checked="" type="checkbox"/> Attach Data Disk	Ultra Disk	100 GiB	

9. Select the Kubernetes version that supports the Windows operating system. By default, the Kubernetes version with this support is selected.

**Note:**

The system does not display the Operating Stem option if you select the Kubernetes version that does not support the Windows operating system.

Docker Version	18.09.2	
Kubernetes Version	1.12.6-aliyun.1	1.11.5

10. Select the Windows operating system.

Operating System	Linux	Windows (Beta)
------------------	--------------------	-----------------------------

11.Set the logon password.

- **Logon Password:** Set the node logon password.
- **Confirm Password:** Confirm your node logon password.

* Logon Password

The password should be 8-30 characters long and contain three types of characters (uppercase/lowercase letters, numbers and special characters).

* Confirm Password

⊗ Please fill in a valid password

12.Set the Pod Network CIDR and Service CIDR.



Note:

- If you choose to use an existing VPC, we recommend that you set these two parameters.
- The two Classless Inter-Domain Routing (CIDR) blocks cannot overlap with each other, or with your selected VPC, or with the CIDR blocks used by any other existing Kubernetes clusters in the selected VPC. Furthermore, they cannot be modified after the cluster is created. For more information, see [Plan Kubernetes CIDR blocks under VPC](#).

Pod Network CIDR

172.20.0.0/16

Please fill in a valid private CIDR, namely the following CIDR and its subnets: 10.0.0.0/8, 172.16-31.0.0/12-16, 192.168.0.0/16
Cannot be duplicated with the VPC CIDR and CIDR used by Kubernetes cluster in VPC, cannot be modified after creation
For more information about CIDR block planning for a cluster, see [Plan Kubernetes CIDR blocks under a VPC](#).

Service CIDR

172.21.0.0/20

Optional range: 10.0.0.0/16-24, 172.16-31.0.0/16-24, 192.168.0.0/16-24
Cannot be duplicated with the VPC CIDR and CIDR used by Kubernetes cluster in VPC, cannot be modified after creation

13.Configure an SNAT gateway for the VPC.



Note:

We recommend that you configure this setting. Otherwise, instances in the VPC cannot access the Internet, which will result in a cluster creation failure.

- If you set the system to automatically create a VPC, you must configure an SNAT gateway for the VPC.
- If you set the system to use an existing VPC, you can perform one of the following two operations:
 - Set the system to automatically configure an SNAT gateway.
 - Manually configure a NAT gateway or SNAT gateway for the VPC.

Configure SNAT

☒ Configure SNAT for VPC

If the VPC you choose does not have access to Internet, NAT gateway and EIP will be used to configure SNAT for the VPC. During this period, NAT gateway, EIP, and other resources may be created.

14.Enable the Use Public SLB to Expose API Server function.

The Kubernetes API supports such RESTful API actions as retrieving, creating, querying, updating, and deleting resources, such as pods and services.

- If you enable this function by selecting the check box, an Internet SLB instance is created and the Master node port (namely, port 6443) is opened. In this case, you can use kubeconfig to connect to and operate the cluster through the Internet.



Note:

The API server uses the Master node port.

- If you do not enable this function, no Internet SLB instance is created. In this case, you can only use kubeconfig to connect to and operate the cluster within the VPC.

Public Access

☒ Expose API Server with EIP

15.Enable Log Service.

You can select an existing project or create a new project.

If you select the Using Log Service check box, a Log Service plugin is automatically installed in the cluster. Then, when you create an application in the cluster, you

can immediately use Log Service with only a few configurations required. For more information, see [Use Log Service to collect Kubernetes cluster logs](#).

Log Service

☒ Using Log Service

Select Project

Create Project

A SLS Project named k8s-log-{ClusterID} will be created automatically

16.Set the RDS instance whitelist.

Add the IP address of the node to the RDS instance whitelist.



Note:

If you choose to use an existing VPC, we recommend that you set this parameter.

RDS Whitelist

Select RDS Instances

17.Set advanced configurations.

Set the number of pods for a node. This parameter specifies the maximum number of pods that can be run by a single node. We recommend that you retain the default setting.

Pod Number for

128

Node

18.Click Create, and then click Create in the displayed dialog box.



Note:

A Kubernetes cluster that contains multiple nodes typically takes ten minutes to be created.

Verify the result

After the cluster is created, you can view the cluster in the cluster list of the Container Service console.

Container Service - Kubernetes

Overview

Clusters

Nodes

Volumes

Namespaces

Authorization

Application

Cluster List

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

Refresh

Create Kubernetes Cluster

Help: [Create cluster](#) [Create GPU clusters](#) [Scale cluster](#) [Connect to Kubernetes cluster via kubectl](#) [Manage applications with commands](#) [Cluster planning](#) [Troubleshoot](#)

cluster creation failures [Authorization management](#)

Name

Tags

Cluster Name/ID	Tags	Cluster Type	Region (All)	Network Type	Cluster Status	Number of Nodes	Time Created	Version	Action
111		Serverless Kubernetes	China North 2 (Beijing)	VPC	Running		04/20/2019,18:27:30	1.11.5	Manage View Logs Delete Use Cloud Shell
k8s-windows-cluster		Windows Kubernetes	China East 1 (Hangzhou)	VPC	Running	3	04/20/2019,15:46:00	1.12.6-alinyun.1	Manage View Logs Scale Out More

- To view the cluster logs, click View Logs on the right of the cluster. To view more details, click Stack Events.

Container Service - Kubernetes	Cluster Logs: k8s-windows-cluster									
	Back to Cluster List Refresh									
Overview	Detailed resource deployment logs Stack Events									
Clusters										
Nodes										
Volumes										
Namespaces										
Authorization										
Application										
Deployments										
Stateful Sets										
Daemon Sets										
Jobs										
Time	Information									
04/22/2019,10:51:01	Start to DescribeK8sUserCertConfig									
04/22/2019,10:49:30	Start to DescribeK8sUserCertConfig									
04/22/2019,10:32:43	Start to DescribeK8sUserCertConfig									
04/20/2019,16:03:53	start to update cluster status CREATE_COMPLETE									
04/20/2019,16:03:53	Successfully to create managed kubernetes cluster									
04/20/2019,15:58:49	Install addons successfully									
04/20/2019,15:58:17	Start to install addons									
04/20/2019,15:57:15	Stack CREATE completed successfully:									
04/20/2019,15:47:14	Successfully to CreateStack with response &ros.CreateStackResponse[Id:"...", Name:"k8s-for-..."]									

- To view the basic information, the connection information, and other information, return to the cluster list page and click Manage in the action column of the cluster.

Cluster:k8s-windows-cluster

Refresh

Use Cloud Shell

Basic Information

Cluster ID: c-	VPC	<div>Running</div>	Region: China East 1 (Hangzhou)
----------------	-----	--------------------	---------------------------------

Cluster Information

API Server Internet endpoint	https://-
API Server Intranet endpoint	https://-
Pod Network CIDR	-
Service CIDR	-
Service Access Domain	-

Cluster Resource

ROS	-
VPC	-
Worker RAM Role	-

Connect to Kubernetes cluster via kubectl (Use Cloud Shell)

- Download the latest kubectl client from the [Kubernetes Edition page](#) .
- Install and set up the kubectl client. For more information, see [Installing and Setting Up kubectl](#)
- Configure the cluster credentials:

Cluster information

- **API Server Internet endpoint:** the IP address and the port through which the Kubernetes API server provides services to the Internet. It enables you to manage the cluster by using `kubectl` or other tools on your terminal.
- **API Server Intranet endpoint:** the IP address and the port through which the Kubernetes API server provides services within the cluster. This IP address is the SLB instance IP address. The three Master nodes on its backend provide services.
- **Pod Network CIDR:** the CIDR block where the pods of a Kubernetes cluster are located.
- **Service CIDR:** the CIDR block where the services of a Kubernetes cluster are located.

For example, you can log on to the Master node by using SSH (for more information, see [Connect to a Kubernetes cluster by using kubectl](#)), and run the `kubectl get node` command to view the cluster nodes.

```
shell@Alicloud:~$ kubectl get node
```

NAME	STATUS	ROLES	AGE	VERSION
cn-hangzhou. [REDACTED]	Ready	<none>	26m	v1.12.6-aliyun.1
cn-hangzhou. [REDACTED]	Ready	<none>	26m	v1.12.6-aliyun.1
cn-hangzhou. [REDACTED]	Ready	<none>	26m	v1.12.6-aliyun.1

```
shell@Alicloud:~$
```

1.5 Node management

1.5.1 Add an existing ECS instance

You can add existing Elastic Compute Service (ECS) instances to a created Kubernetes cluster. Currently, Kubernetes clusters only support adding worker nodes.

Prerequisites

- If you have not created a cluster before, create a cluster first. For how to create a cluster, see [Create a Kubernetes cluster](#).
- Add the ECS instance to the security group of the Kubernetes cluster first.

Context

- By default, each cluster can contain up to 40 nodes. To add more nodes, open a ticket.
- The ECS instance to be added must be in the same Virtual Private Cloud (VPC) region as the cluster.

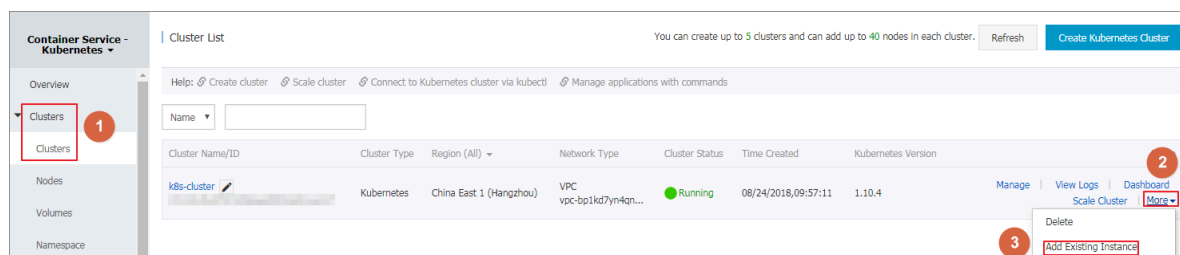
- When adding an existing instance, make sure that your instance has an Elastic IP (EIP) for the VPC network type, or the corresponding VPC is already configured with the NAT gateway. In short, make sure the corresponding node can access public network normally. Otherwise, the ECS instance fails to be added.
- The ECS instance to be added must be under the same account as the cluster.
- Only the ECS instance whose operating system is CentOS can be added.

Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Clusters in the left-side navigation pane.
3. Select the target cluster and click More > Add Existing Instance.

The Add Existing ECS Instance page appears. All the available ECS instances under the current account are displayed on this page. Select to add existing ECS instances automatically or manually.

If Automatically Add is selected, select the ECS instances to add them to the cluster automatically. If Manually Add is selected, you must obtain the command and then log on to the corresponding ECS instance to add the ECS instance to this cluster. You can only add one ECS instance at a time.



4. Select Automatically Add to add multiple ECS instances at a time.

- a) In the list of existing cloud servers, select the target ECS instance, and then click Next Step.

Add Existing ECS Instance - test [Back to Cluster List](#)

Progress: Select Existing ECS Instance(s) > Enter Instance Information > Added Successfully

Adding Method: **Automatically Add** (1) ☐ Manually Add

Adding an existing ECS instance to the Container Service will change the instance system disk. The disk ID will change and the previous system disk will be released.

Note:

1. Your previous system disk's user snapshots will be retained. Automatic snapshots will be retained or deleted according to the setting of the "Delete automatic snapshots when releasing disk" option. You can go to the disk list and click "Modify Attributes" to view and modify attribute values.
2. To retain enough snapshot quota for the periodic automatic snapshot policy, you can delete unwanted user and automatic snapshots.
3. Back up data before performing this operation to avoid data loss.

Select Existing ECS Instance(s):

Instance ID Enter the instance ID for exact search.

Instance ID	Instance Name	IP Address	Zone	Network Type	Instance Type
<input checked="" type="checkbox"/> [redacted]	test	47.87.208.17 (Elastic) 192.168.245.100 (Private)	cn-hangzhou-g	VPC	ecs.sn2ne.large

(2) **Next Step** (3)

- b) Enter the instance information, set the logon password, and then click Next Step.

Add Existing ECS Instance - test [Back to Cluster List](#)

Progress: Select Existing ECS Instance(s) > Enter Instance Information > Added Successfully

Cluster ID/Name: [redacted] / test
Information of the cluster to which to add the ECS instance(s).

Instance Information:

Instance ID	Instance Name
[redacted]	test

Next Step

- c) Click Confirm in the displayed dialog box. The selected ECS instances are automatically added to this cluster.

Add Existing ECS Instance - test [Back to Cluster List](#)

Progress: Select Existing ECS Instance(s) > Enter Instance Information > Added Successfully

Cluster ID/Name: [redacted] / test
Information of the cluster to which to add the ECS instance(s).

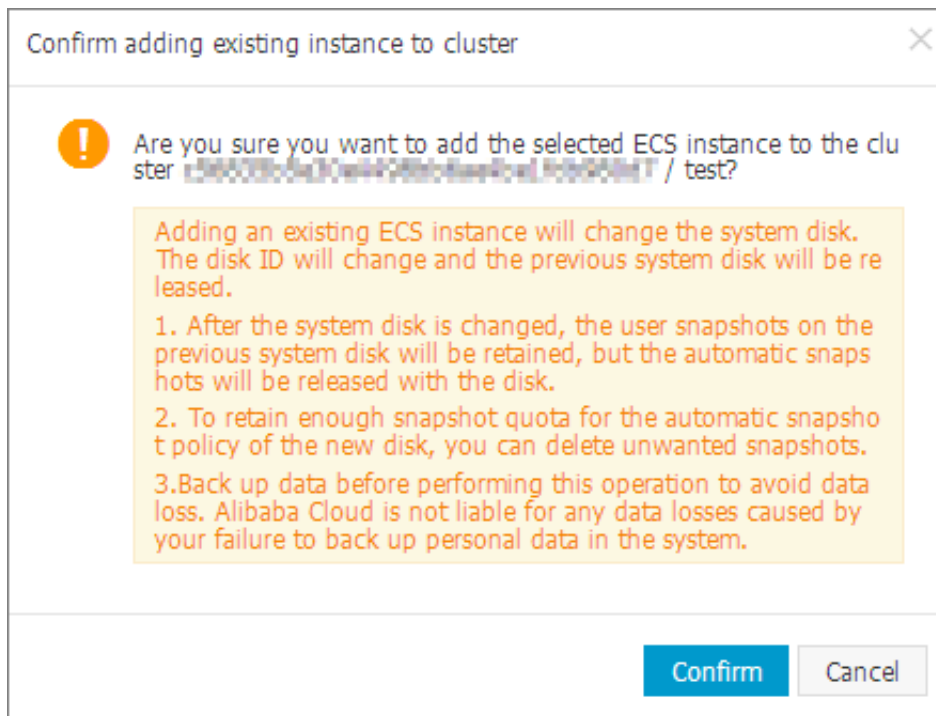
Instance Information:

Instance ID	Instance Name
[redacted]	test

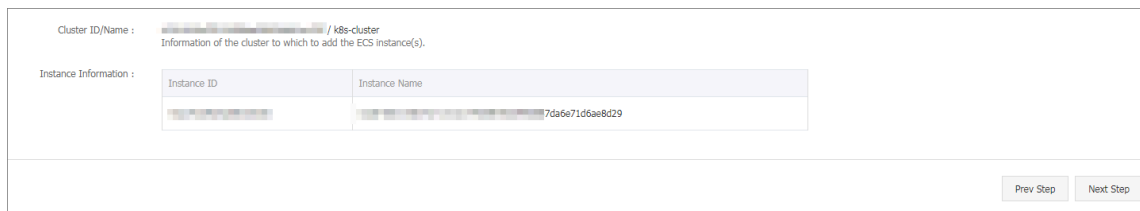
Next Step

5. Optional: You can also select Manually Add to manually add an existing ECS instance to the cluster.

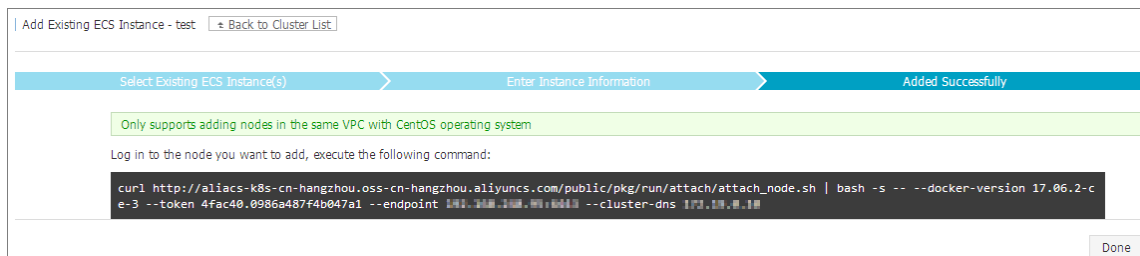
a) Select the ECS instance to be added and then click Next Step. You can add only one ECS instance at a time.



b) Confirm the information and then click Next Step.



c) Copy the command.

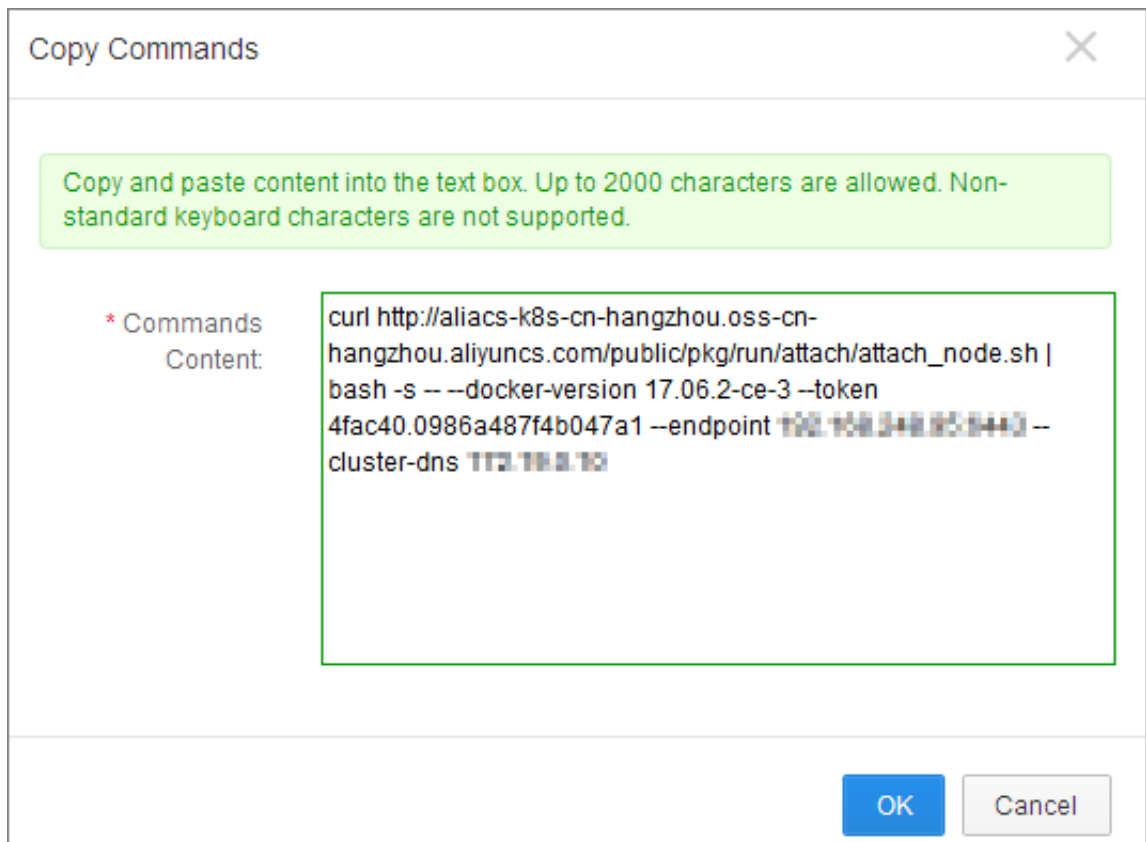


d) Click Done.

e) Log on to the [ECS console](#) and click Instances in the left-side navigation pane.

Select the region in which the cluster resides and the ECS instance to be added.

- f) Click **Connect** at the right of the ECS instance to be added. The **Enter VNC Password** dialog box appears. Enter the VNC password and then click **OK**. Enter the copied command and then click **OK** to run the script.



- g) After the script is successfully run, the ECS instance is added to the cluster. You can click the cluster ID on the **Cluster List** page to view the node list of the cluster and check if the ECS instance is successfully added to the cluster.

1.5.2 View node list

You can view the node list of the Kubernetes cluster by using commands, in the Container Service console, or in the Kubernetes dashboard.

View node list by using commands



Note:

Before using commands to view the node list of the Kubernetes cluster, [#unique_57](#) first.

After connecting to the Kubernetes cluster by using `kubectl`, run the following command to view the nodes in the cluster:

```
kubectl get nodes
```

Sample output:

```
$ kubectl get nodes
NAME                STATUS    AGE       VERSION
iz2ze2n6ep          Ready    19m       v1.6.1-2+
ed9e3d33a0          Ready    7m         v1.6.1-2+
iz2zeafr76          Ready    7m         v1.6.1-2+
ed9e3d33a0          Ready    7m         v1.6.1-2+
iz2zeafr76          Ready    14m        v1.6.1-2+
ed9e3d33a0          Ready    11m        v1.6.1-2+
iz2zef4dnn          Ready    8m         v1.6.1-2+
ed9e3d33a0          Ready    7m         v1.6.1-2+
iz2zeitvvo          Ready    7m         v1.6.1-2+
ed9e3d33a0          Ready    7m         v1.6.1-2+
```

View node list in Container Service console

1. Log on to the [Container Service console](#).
2. Click **Kubernetes > Clusters > > Nodes** in the left-side navigation pane.
3. Select the cluster from the Cluster drop-down list and then view the node list of this cluster.

View node list in Kubernetes dashboard

1. Log on to the [Container Service console](#).
2. Click **Kubernetes > Clusters** in the left-side navigation pane.
3. Click **Dashboard** at the right of the cluster to enter the Kubernetes dashboard.
4. In the Kubernetes dashboard, click **Nodes** in the left-side navigation pane to view the node list of this cluster.

1.5.3 Node monitoring

Kubernetes clusters integrate with the Alibaba Cloud monitoring service seamlessly. You can view the monitoring information of Kubernetes nodes and get to know the node monitoring metrics of the Elastic Compute Service (ECS) instances under Kubernetes clusters.

Procedure

1. Log on to the [Container Service console](#).

2. Under Kubernetes, click Clusters > Nodes to enter the Node List page.
3. Select the target cluster and node under the cluster.
4. Click Monitor at the right of the node to view the monitoring information of this node.
5. You are redirected to the CloudMonitor console. View the basic monitoring information of the corresponding ECS instance, including the CPU usage, network inbound bandwidth, network outbound bandwidth, disk BPS, and disk IOPS.

What's next

To view the monitoring metrics at the operating system level, install the CloudMonitor component. For more information, see [Host monitoring overview](#).

Kubernetes clusters can now monitor resources by using application groups. For more information, see [#unique_60](#).

1.5.4 Manage node labels

You can manage node labels in the Container Service console, including adding node labels in batches, filtering nodes by using a label, and deleting a node label quickly.

For how to use node labels to schedule pods to specified nodes, see [#unique_62](#).

Prerequisite

You have successfully created a Kubernetes cluster. For more information, see [#unique_63](#).

Add node labels in batches

1. Log on to the [Container Service console](#).
2. Click Kubernetes Clusters > Nodes in the left-side navigation pane.
3. Select the cluster from the Clusters drop-down list and then click Label Management in the upper-right corner.
4. Select one or more nodes by selecting the corresponding check boxes and then click Add Tag.

5. Enter the name and value of the label in the displayed dialog box and then click OK.

Nodes with the same label are displayed on the Label Management page.

Filter nodes by using a label

1. Log on to the [Container Service console](#).
2. Click Kubernetes Clusters > Nodes in the left-side navigation pane.
3. Select the cluster from the Clusters drop-down list and then click Label Management in the upper-right corner.
4. Click the label at the right of a node to filter nodes by using the label. In this example, click `group : worker`.

Nodes with the label `group : worker` are filtered.

Delete a node label

1. Log on to the [Container Service console](#).
2. Click Kubernetes Clusters > Nodes in the left-side navigation pane.
3. Select the cluster from the Clusters drop-down list and then click Label Management in the upper-right corner.
4. Click the delete (x) button of a node label, for example, `group : worker`.

Click Confirm in the displayed dialog box. The node label is deleted.

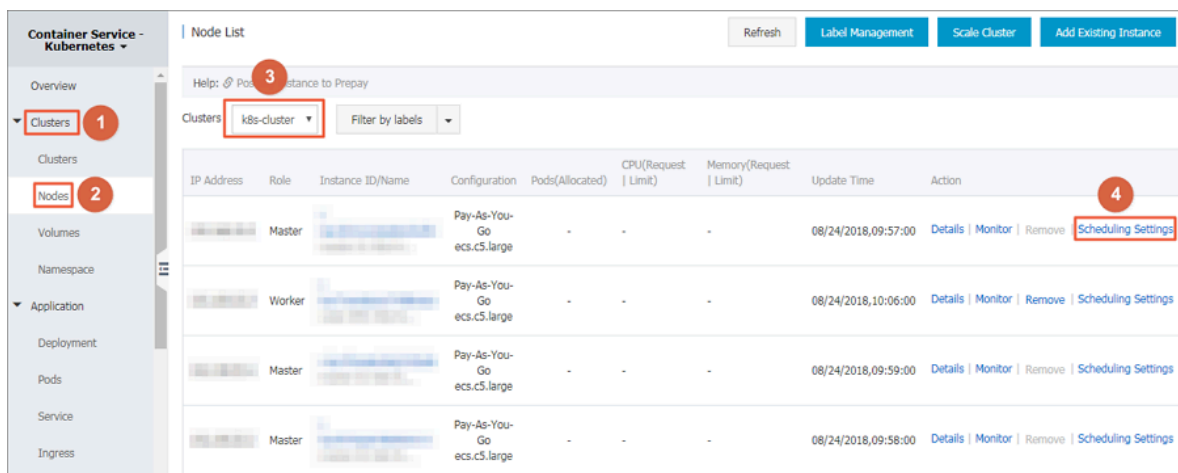
1.5.5 Set node scheduling

You can set node scheduling through the web interface so that you can allocate loads to each node properly.

Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Clusters > Nodes to enter the Node List page.

3. Select a cluster, select a node under the cluster, and click Schedule Settings on the right.



The screenshot shows the 'Node List' page in the Container Service - Kubernetes console. The sidebar on the left contains a navigation menu with 'Clusters' (1) and 'Nodes' (2) highlighted. The main content area shows a table of nodes. The 'Clusters' dropdown is set to 'k8s-cluster' (3). The table lists four nodes, each with a 'Scheduling Settings' link in the 'Action' column (4).

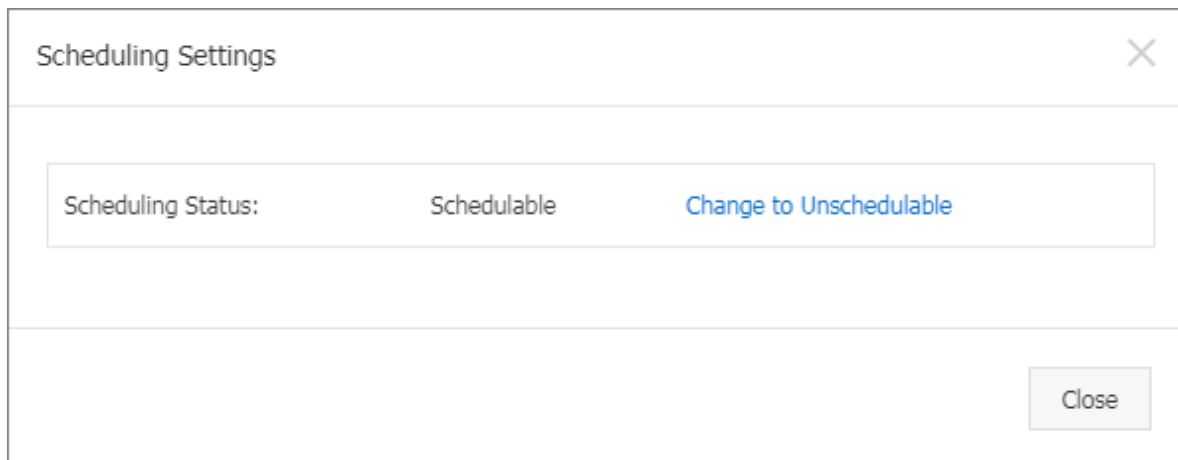
IP Address	Role	Instance ID/Name	Configuration	Pods(Allocated)	CPU(Request Limit)	Memory(Request Limit)	Update Time	Action
192.168.1.1	Master	ecs-c5.large	Pay-As-You-Go	-	-	-	08/24/2018,09:57:00	Details Monitor Remove Scheduling Settings
192.168.1.2	Worker	ecs-c5.large	Pay-As-You-Go	-	-	-	08/24/2018,10:06:00	Details Monitor Remove Scheduling Settings
192.168.1.3	Master	ecs-c5.large	Pay-As-You-Go	-	-	-	08/24/2018,09:59:00	Details Monitor Remove Scheduling Settings
192.168.1.4	Master	ecs-c5.large	Pay-As-You-Go	-	-	-	08/24/2018,09:58:00	Details Monitor Remove Scheduling Settings

4. Set node scheduling in the displayed dialog box. In this example, click Change to Unschedulable to set the node to unschedulable.

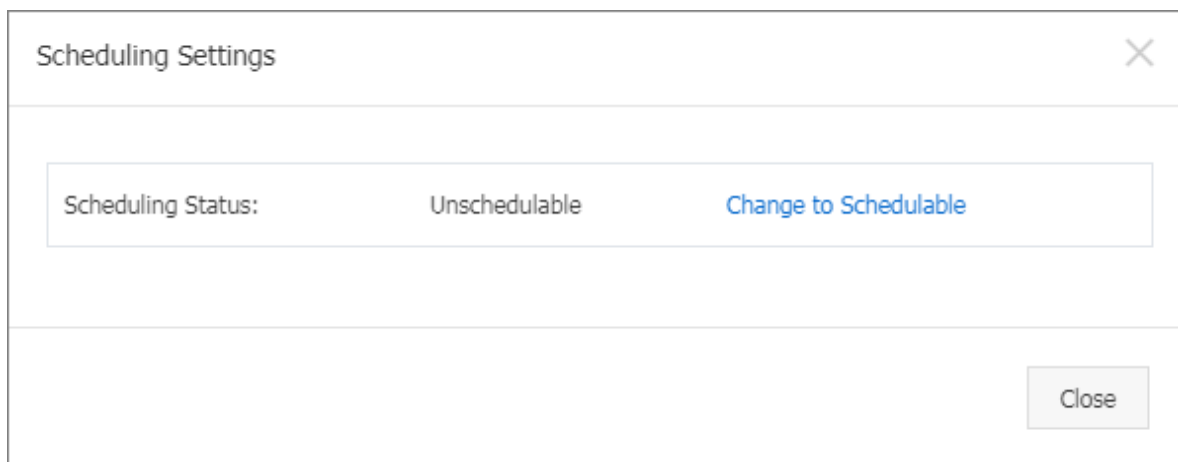


Note:

The scheduling status of the current node is displayed in the Scheduling Settings dialog box, which is schedulable by default. You can change the status.



After the status is set, the scheduling status of the node changes in the dialog box.



What's next

When you deploy your application later, you can find that pods are not scheduled to the node.

1.5.6 Remove a node

Before you restart or release an ECS instance in a Kubernetes cluster, you need to remove the ECS node from the cluster. This topic describes how to remove a node from a Kubernetes cluster.

Prerequisites

- You have created a Kubernetes cluster. For more information, see [Create a Kubernetes cluster](#).

- You have connected to the Kubernetes cluster by using `kubectl`, see [Connect to a Kubernetes cluster by using kubectl](#).

Context

- Removing a node causes pod migration. This may affect the services provided by the pods running on the node. Therefore, we recommend that you remove a node only when fewer services are in demand.
- Removing a node may cause unintended risks. We recommend that you back up your data in advance and exercise caution when performing this action.
- Only Worker nodes can be removed.

Procedure

1. Run the following command to migrate the pods on the target node to other nodes:



Note:

You must ensure that other nodes in the Kubernetes cluster have sufficient resources to run the pods that you want to migrate.

```
kubectl drain node-name
```



Note:

The `node-name` parameter must be in the format of `your-region-name.node-id`.

- `your-region-name` indicates the name of the region where your cluster resides.
- `node-id` indicates the ID of the ECS instance in which the node to be removed resides. For example, `cn - hangzhou . i - xxx`.

2. Set the node to be removed as the non-schedulable node.

Method 1: Use a command

- Run the following command to set the node to be removed as the non-schedulable node:

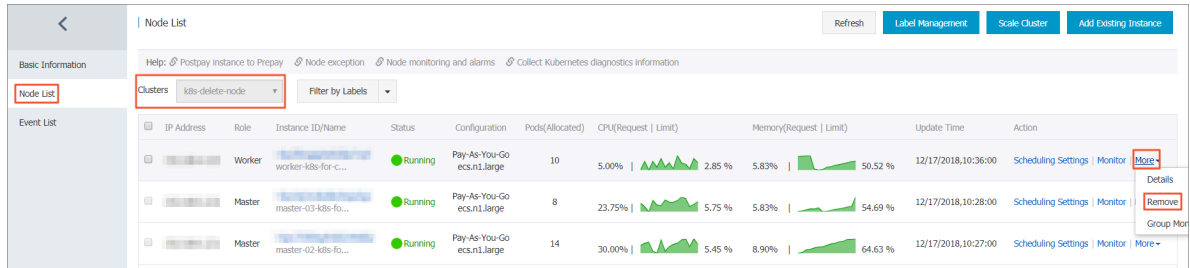
```
kubectl cordon node-name
```

Method 2: Use the Container Service console

For more information, see [Set node scheduling](#).

3. In the left-side navigation pane under Kubernetes, choose Clusters > Nodes.

4. Under the target cluster, select the target node, and choose More > Remove in the Action column.



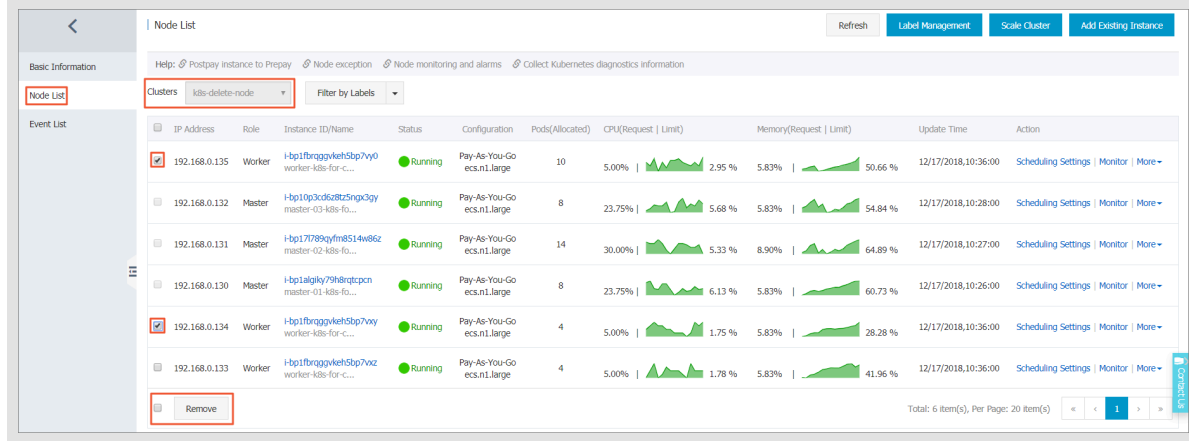
The screenshot shows the 'Node List' page. At the top, there are tabs for 'Basic Information', 'Node List', and 'Event List'. The 'Node List' tab is selected. Below the tabs, there is a 'Clusters' dropdown menu set to 'k8s-delete-node'. A table lists nodes with columns: IP Address, Role, Instance ID/Name, Status, Configuration, Pods(Allocated), CPU(Request | Limit), Memory(Request | Limit), Update Time, and Action. The 'Action' column for the first node (Worker) has a 'More' button highlighted with a red box.

IP Address	Role	Instance ID/Name	Status	Configuration	Pods(Allocated)	CPU(Request Limit)	Memory(Request Limit)	Update Time	Action
192.168.0.135	Worker	h-p1fbraggvkeh5bp7yy0	Running	Pay-As-You-Go ecs.n1.large	10	5.00% 2.95 %	5.83% 50.52 %	12/17/2018,10:36:00	Scheduling Settings Monitor More
192.168.0.132	Master	h-p1tp3dc5c8t25ng0gy	Running	Pay-As-You-Go ecs.n1.large	8	23.75% 5.68 %	5.83% 54.84 %	12/17/2018,10:28:00	Scheduling Settings Monitor More
192.168.0.131	Master	h-p17f89qym8514w86z	Running	Pay-As-You-Go ecs.n1.large	14	30.00% 5.33 %	8.90% 64.89 %	12/17/2018,10:27:00	Scheduling Settings Monitor More
192.168.0.130	Master	h-p1aigky79h8qtcpn	Running	Pay-As-You-Go ecs.n1.large	8	23.75% 6.13 %	5.83% 60.73 %	12/17/2018,10:26:00	Scheduling Settings Monitor More
192.168.0.134	Worker	h-p1fbraggvkeh5bp7xy	Running	Pay-As-You-Go ecs.n1.large	4	5.00% 1.75 %	5.83% 28.28 %	12/17/2018,10:36:00	Scheduling Settings Monitor More
192.168.0.133	Worker	h-p1fbraggvkeh5bp7wz	Running	Pay-As-You-Go ecs.n1.large	4	5.00% 1.78 %	5.83% 41.96 %	12/17/2018,10:36:00	Scheduling Settings Monitor More



Note:

If you want to remove multiple nodes at a time, you can select the target cluster on the Node List page, select all the nodes to be removed, and then click Remove.

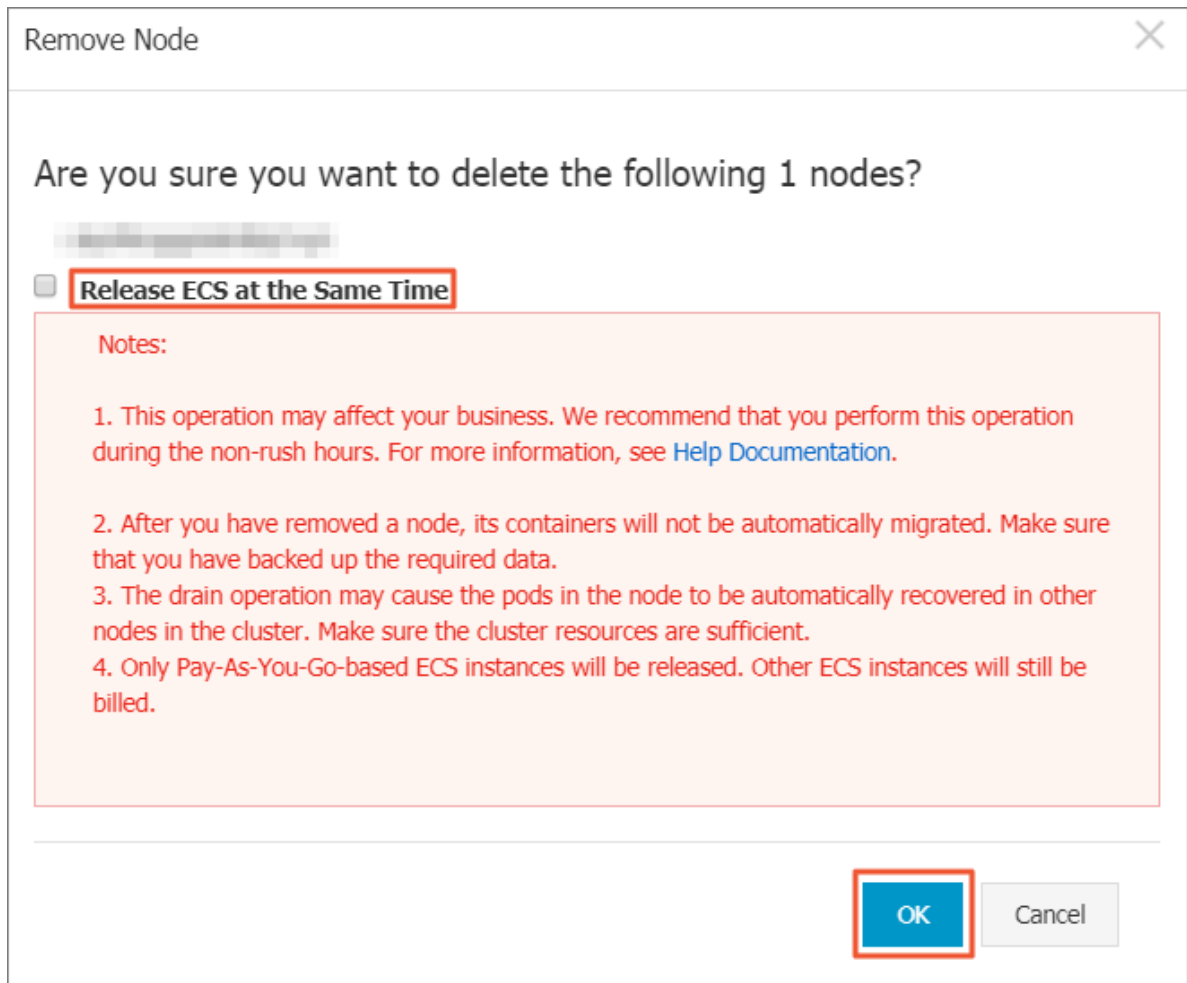


The screenshot shows the 'Node List' page with a table of nodes. The 'Action' column for the first node (Worker) has a 'More' button highlighted with a red box. Below the table, there is a 'Remove' button highlighted with a red box. The table has columns: IP Address, Role, Instance ID/Name, Status, Configuration, Pods(Allocated), CPU(Request | Limit), Memory(Request | Limit), Update Time, and Action. The 'Action' column for the first node (Worker) has a 'More' button highlighted with a red box.

IP Address	Role	Instance ID/Name	Status	Configuration	Pods(Allocated)	CPU(Request Limit)	Memory(Request Limit)	Update Time	Action
192.168.0.135	Worker	h-p1fbraggvkeh5bp7yy0	Running	Pay-As-You-Go ecs.n1.large	10	5.00% 2.95 %	5.83% 50.52 %	12/17/2018,10:36:00	Scheduling Settings Monitor More
192.168.0.132	Master	h-p1tp3dc5c8t25ng0gy	Running	Pay-As-You-Go ecs.n1.large	8	23.75% 5.68 %	5.83% 54.84 %	12/17/2018,10:28:00	Scheduling Settings Monitor More
192.168.0.131	Master	h-p17f89qym8514w86z	Running	Pay-As-You-Go ecs.n1.large	14	30.00% 5.33 %	8.90% 64.89 %	12/17/2018,10:27:00	Scheduling Settings Monitor More
192.168.0.130	Master	h-p1aigky79h8qtcpn	Running	Pay-As-You-Go ecs.n1.large	8	23.75% 6.13 %	5.83% 60.73 %	12/17/2018,10:26:00	Scheduling Settings Monitor More
192.168.0.134	Worker	h-p1fbraggvkeh5bp7xy	Running	Pay-As-You-Go ecs.n1.large	4	5.00% 1.75 %	5.83% 28.28 %	12/17/2018,10:36:00	Scheduling Settings Monitor More
192.168.0.133	Worker	h-p1fbraggvkeh5bp7wz	Running	Pay-As-You-Go ecs.n1.large	4	5.00% 1.78 %	5.83% 41.96 %	12/17/2018,10:36:00	Scheduling Settings Monitor More

Total: 6 item(s), Per Page: 20 item(s)

5. Optional: Select the Release ECS at the Same Time check box to permanently release the ECS instance where the node resides.



Remove Node

Are you sure you want to delete the following 1 nodes?

☐ **Release ECS at the Same Time**

Notes:

1. This operation may affect your business. We recommend that you perform this operation during the non-rush hours. For more information, see [Help Documentation](#).
2. After you have removed a node, its containers will not be automatically migrated. Make sure that you have backed up the required data.
3. The drain operation may cause the pods in the node to be automatically recovered in other nodes in the cluster. Make sure the cluster resources are sufficient.
4. Only Pay-As-You-Go-based ECS instances will be released. Other ECS instances will still be billed.

OK Cancel



Note:

- Only Pay-As-You-Go ECS instances can be released.
- A Subscription ECS instance will be released automatically when it expires.
- If you do not select the Release ECS at the Same Time check box, the ECS instance in which the node resides will continue to be charged.

6. Click OK.

1.5.7 Use Alibaba Cloud Kubernetes GPU node labels for scheduling

When you implement GPU computing through a Kubernetes cluster, you can schedule an application to the node installed with GPU devices as needed by using GPU node labels.

Prerequisites

- You have created a Kubernetes cluster that has GPU nodes. For more information, see [Configure a Kubernetes GPU cluster to support GPU scheduling](#).
- You have connected to the Master node, which makes it easier to view node labels and other information. For more information, see [Connect to a Kubernetes cluster by using kubectl](#).

Context

When deploying NVIDIA GPU nodes, Kubernetes that runs on Alibaba Cloud discovers the GPU attribute and exposes it as the node label information. Node labels provide the following benefits:

1. Node labels help you filter GPU nodes.
2. Node labels can be used as the scheduling conditions for application deployment.

Procedure

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Kubernetes, choose Clusters > Nodes.



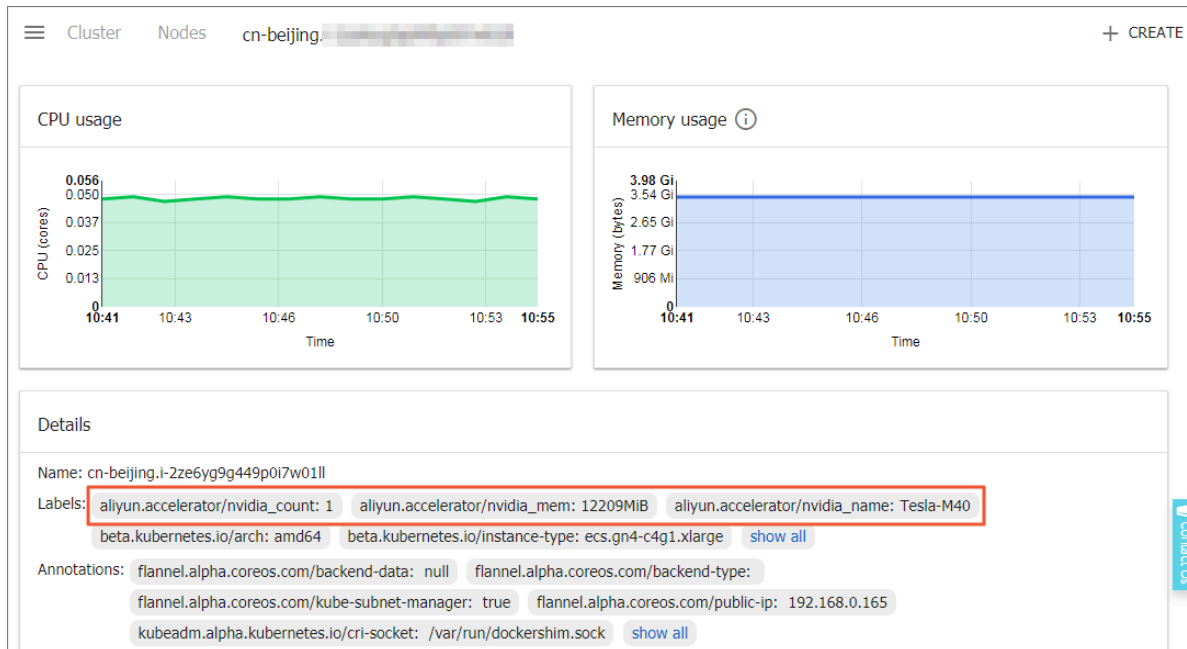
Note:

In this example, the cluster has three Worker nodes of which two Worker nodes are mounted with GPU devices. You need to view the node IP addresses for verification.

IP Address	Role	Instance ID/Name	Configuration	Pods(Allocated)	CPU(Request Limit)	Memory(Request Limit)	Update Time	Action
165	Worker	Pay-As-You-Go ecs.gn4-c4g1.xlarge	Pay-As-You-Go ecs.gn4-c4g1.xlarge	5	30.00% 1.23 %	2.52% 11.20 %	12/05/2018,13:40:00	Scheduling Settings Monitor More
166	Worker	Pay-As-You-Go ecs.gn4-c4g1.xlarge	Pay-As-You-Go ecs.gn4-c4g1.xlarge	11	30.00% 1.88 %	2.52% 16.33 %	12/05/2018,13:40:00	Scheduling Settings Monitor More
164	Master	Pay-As-You-Go ecs.n1.large	Pay-As-You-Go ecs.n1.large	8	23.75% 4.68 %	5.83% 38.71 %	12/05/2018,13:30:00	Scheduling Settings Monitor More
163	Master	Pay-As-You-Go ecs.n1.large	Pay-As-You-Go ecs.n1.large	8	23.75% 4.30 %	5.83% 39.02 %	12/05/2018,13:29:00	Scheduling Settings Monitor More
162	Master	Pay-As-You-Go ecs.n1.large	Pay-As-You-Go ecs.n1.large	14	30.00% 6.25 %	8.90% 47.82 %	12/05/2018,13:28:00	Scheduling Settings Monitor More
167	Worker	Pay-As-You-Go ecs.n4.large	Pay-As-You-Go ecs.n4.large	4	10.00% 1.90 %	12.03% 52.84 %	12/05/2018,13:55:00	Scheduling Settings Monitor More

Total: 6 Item(s) , Per Page : 20 Item(s)

3. Select a GPU node, and choose More > Details in the action column. Then, you can view the GPU node label on the Kubernetes dashboard.



You can also log on to a Master node and run the following command to view the GPU node label:

```
# kubectl get nodes
NAME                                STATUS    ROLES    AGE
cn-beijing.i-2ze2dy2h9w.97v65uuaft Ready     master
cn-beijing.i-2ze8o1a45q.dv5q8a7luz Ready     < none >
cn-beijing.i-2ze8o1a45q.dv5q8a7lv0 Ready     < none >
cn-beijing.i-2ze9xylyn1.1vop7g5bwe Ready     master
cn-beijing.i-2zed5sw8sn.jniq6mf5e5 Ready     master
cn-beijing.i-2zej9s0zij.ykp9pwf7lu Ready     < none >
```

with the node displayed in the console to determine the GPU node.

Select a GPU node and run the following command to view the GPU node label:

```
# kubectl describe node cn-beijing.i-2ze8o1a45q
Name: cn-beijing.i-2ze8o1a45q.dv5q8a7luz
Roles: < none >
Labels: aliyun.accelerator/nvidia_count=1
        # This field is important.
        aliyun.accelerator/nvidia_mem=12209MiB
        aliyun.accelerator/nvidia_name=Tesla-M40
        beta.kubernetes.io/arch=amd64
```

```

        beta . kubernetes . io / instance - type = ecs
    . gn4 - c4g1 . xlarge
        beta . kubernetes . io / os = linux
        failure - domain . beta . kubernetes . io /
region = cn - beijing
        failure - domain . beta . kubernetes . io /
zone = cn - beijing - a
        kubernetes . io / hostname = cn - beijing . i -
2ze8o1a45q dv5q8a7luz
.....

```

In this example, the GPU node contains the following three node labels:

Key	Value
aliyun . accelerator / nvidia_count	Number of GPU cores
aliyun . accelerator / nvidia_mem	GPU memory in MiB
aliyun . accelerator / nvidia_name	Name of the GPU computing card of the NVIDIA device

The GPU cloud servers of the same type share the same GPU computing card name. Therefore, you can use this label to filter nodes.

```

# kubectl get node -l aliyun . accelerator / nvidia_name
= Tesla - M40
NAME                                STATUS    ROLES    AGE
VERSION
cn - beijing . i - 2ze8o1a45q dv5q8a7luz    Ready    < none >
2d                                  v1 . 11 . 2
cn - beijing . i - 2ze8o1a45q dv5q8a7lv0     Ready    < none >
2d                                  v1 . 11 . 2

```

4. Return to the Container Service console home page, choose Application >

Deployment in the left-side navigation pane, and click Create by Template in the upper-right corner.

a) Create a TensorFlow application and schedule this application to the GPU node.

Clusters: k8s-test

Namespace: default

Resource Type: Custom

Template:

```

1 ---
2 # Define the tensorflow deployment
3 apiVersion: apps/v1
4 kind: Deployment
5 metadata:
6   name: tf-notebook
7   labels:
8     app: tf-notebook
9 spec:
10  replicas: 1
11  selector: # define how the deployment finds the pods it manages
12    matchLabels:
13      app: tf-notebook
14  template: # define the pods specifications
15    metadata:
16      labels:
17        app: tf-notebook
18    spec:
19      nodeSelector:
20        aliyun.accelerator/nvidia_name: Tesla-M40
21    containers:
22      - name: tf-notebook
23        image: tensorflow/tensorflow:1.4.1-gpu-py3
24        resources:
25          limits:
26            nvidia.com/gpu: 1
27        ports:
28          - containerPort: 8888
29            hostPort: 8888
30        env:
31          - name: PASSWORD
32            value: mypasswdv
  
```

Add Deployment

Deploy with exist template

Save Template DEPLOY

In this example, the YAML template is orchestrated as follows:

```

---
# Define the tensorflow deployment
apiVersion: apps / v1
kind: Deployment
metadata:
  name: tf - notebook
  labels:
    app: tf - notebook
spec:
  replicas: 1
  selector: # define how the deployment finds the
pods it manages
    matchLabels:
      app: tf - notebook
  template: # Define the pod specifications .
    metadata:
      labels:
        app: tf - notebook
    spec:
      nodeSelector:
        # This field is important .
        aliyun . accelerato r / nvidia_nam e : Tesla - M40
      containers:
  
```

```

- name : tf - notebook
  image : tensorflow / tensorflow : 1 . 4 . 1 - gpu - py3
  resources :
    limits :
      nvidia . com / gpu : 1
      # This field is important .
  ports :
- containerPort : 8888
  hostPort : 8888
  env :
- name : PASSWORD
  value : mypassw0rd v

```

b) You can also avoid deploying an application to a GPU node. The following deploys an Nginx pod and schedules it by using the node affinity feature. For more information about node affinity, see [Create a deployment application by using an image](#).

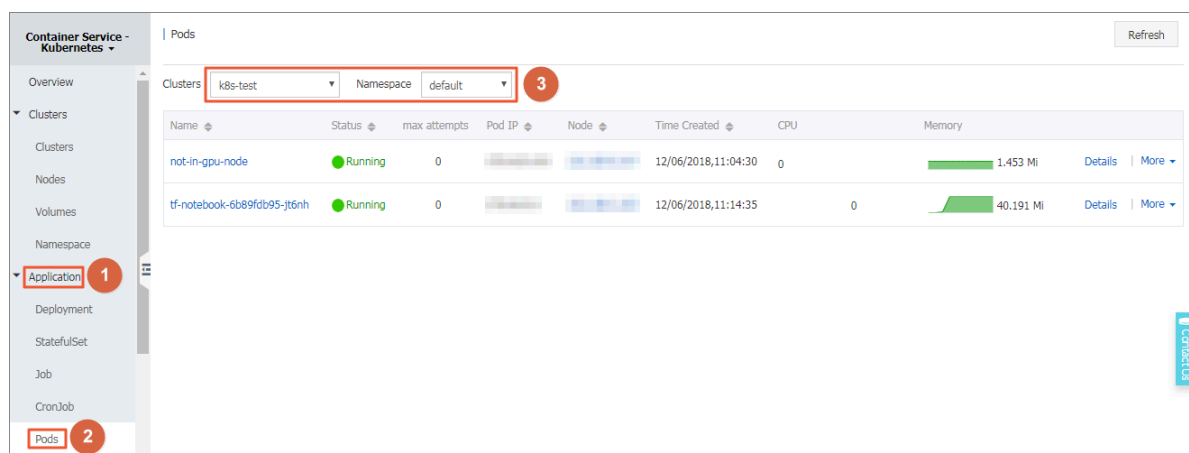
The example YAML template is orchestrated as follows:

```

apiVersion : v1
kind : Pod
metadata :
  name : not - in - gpu - node
spec :
  affinity :
    nodeAffinity :
      requiredDuringSchedulingIgnoredDuringExecution :
        nodeSelectorTerms :
- matchExpressions :
- key : aliyun . accelerator / nvidia_name
  operator : DoesNotExist
  containers :
- name : not - in - gpu - node
  image : nginx

```

5. In the left-side navigation pane, choose **Application > Pods**, and select the target cluster and namespace.



Result

In the pod list, you can see that the two example pods have been scheduled to the target nodes, indicating you have implemented flexible scheduling by using GPU node labels.

1.5.8 View resource request and limit on nodes

The Container Service Console allows you to view resource usage of each node in a Kubernetes cluster.

Prerequisites

You have created a Kubernetes cluster. For more information, see [Create a Kubernetes cluster](#).

Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Clusters > Nodes.

You can view the resource usage for the CPU and memory of each node, namely, the request and limit, which are calculated as follows:

- CPU request = sum (CPU request value from all pods on the current node) /total CPU of the current node.
- CPU limit= sum (actual CPU usage of all pods on the current node)/total CPU of the current node.
- Memory request = sum (memory request value from all pods on the current node) /total memory of the current node.
- Memory limit= sum (actual memory usage of all pods on the current node)/total memory of the current node.



Note:

- You can allocate loads to a node based on the resource usage on the node. For more information, see [Set node scheduling](#).

- When both the request and limit on a node is 100%, no new pod is scheduled to the node.

IP Address	Role	Instance ID/Name	Configuration	Pods(Allocated)	CPU(Request Limit)	Memory(Request Limit)	Update Time	Action
	Master		Pay-As-You-Go ecs.n4.xlarge	12	26.25% 3.15 %	4.35% 53.87 %	09/18/2018,14:02:00	Scheduling Settings Monitor More
	Master		Pay-As-You-Go ecs.n4.xlarge	7	21.25% 4.05 %	2.56% 53.01 %	09/18/2018,14:04:00	Scheduling Settings Monitor More
	Master		Pay-As-You-Go ecs.n4.xlarge	7	21.25% 3.88 %	2.56% 54.15 %	09/18/2018,14:01:00	Scheduling Settings Monitor More
	Worker		Pay-As-You-Go ecs.n4.xlarge	18	52.75% 2.13 %	2.81% 30.70 %	09/18/2018,14:13:00	Scheduling Settings Monitor More

1.5.9 Mount a disk to a Kubernetes cluster node

This topic describes how to mount a disk to a Kubernetes cluster node. Mounting a disk allows you to expand the Docker data directory and maintain a sufficient disk capacity when the number of containers or images that run on a node increases.

Prerequisites

Your Kubernetes cluster version must be v1.10.4 or later.

You can mount a disk to an existing Kubernetes cluster node by using either of the following methods:

- If no disk is mounted to the existing node, see [Mount a disk to the Docker data directory](#).
- If you have created a disk for the existing node, but you have failed to mount the disk to the node, you can follow these steps.



Note:

- We recommend that you create a snapshot of the target node or back up node data to avoid data loss.
- Additionally, you must ensure that you can schedule your cluster applications to other nodes.
- We recommend that you perform this operation during off-peak service hours to avoid disruptions to your business.

- Draining a node reschedules pods on the node to other nodes. Therefore, you must ensure that your Kubernetes cluster contains sufficient nodes. We recommend that you add cluster nodes in advance as needed.

Before performing the operation, you need to determine whether a disk is already mounted to the target cluster node. To do so, run the `df` command on the target Worker node, and then check whether `/var/lib/docker` has been mounted to `/dev/vdb1`. If the disk mounting operation failed, you can mount the disk by following these steps.

```
[root@aliyun-ecs-20190614 ~]# df
Filesystem      1K-blocks    Used Available Use% Mounted on
/dev/vda1       41151808 2273772  36764604   6% /
devtmpfs        3995592    0    3995592   0% /dev
tmpfs           4005096    0    4005096   0% /dev/shm
tmpfs           4005096    508    4004588   1% /run
tmpfs           4005096    0    4005096   0% /sys/fs/cgroup
/dev/vdb1       101441464 61668  96120584   1% /var/lib/docker
tmpfs           801020    0    801020    0% /run/user/0
```

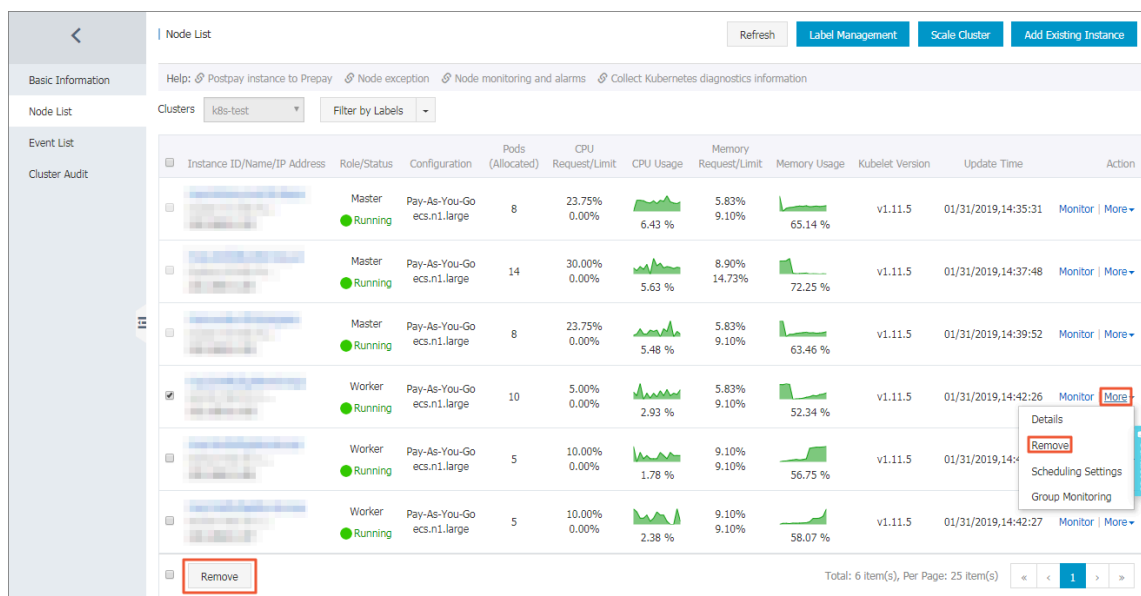
1. Set the target node as unschedulable. For more information, see [Mark node as unschedulable](#).
2. Drain the target node. For more information, see [Safely drain a node](#).

3. Remove the target node. This topic uses the Container Service console as an example.

a. Log on to the [Container Service console](#).

b. In the left-side navigation pane, click Node.

c. Select the target node, and click Remove or choose More > Remove.



The screenshot displays the 'Node List' page in the Container Service console. The left sidebar contains navigation options: Basic Information, Node List, Event List, and Cluster Audit. The main area shows a table of nodes for the 'k8s-test' cluster. The table columns include Instance ID/Name/IP Address, Role/Status, Configuration, Pods (Allocated), CPU Request/Limit, CPU Usage, Memory Request/Limit, Memory Usage, Kubelet Version, Update Time, and Action. A 'Remove' button is highlighted in the bottom left corner of the table. A 'More' button is also highlighted in the 'Action' column for one of the nodes, with a dropdown menu showing options like 'Details', 'Remove', 'Scheduling Settings', and 'Group Monitoring'.

Instance ID/Name/IP Address	Role/Status	Configuration	Pods (Allocated)	CPU Request/Limit	CPU Usage	Memory Request/Limit	Memory Usage	Kubelet Version	Update Time	Action
[Instance ID]	Master Running	Pay-As-You-Go ecs.n1.large	8	23.75% 0.00%	6.43 %	5.83% 9.10%	65.14 %	v1.11.5	01/31/2019,14:35:31	Monitor More
[Instance ID]	Master Running	Pay-As-You-Go ecs.n1.large	14	30.00% 0.00%	5.63 %	8.90% 14.73%	72.25 %	v1.11.5	01/31/2019,14:37:48	Monitor More
[Instance ID]	Master Running	Pay-As-You-Go ecs.n1.large	8	23.75% 0.00%	5.48 %	5.83% 9.10%	63.46 %	v1.11.5	01/31/2019,14:39:52	Monitor More
[Instance ID]	Worker Running	Pay-As-You-Go ecs.n1.large	10	5.00% 0.00%	2.93 %	5.83% 9.10%	52.34 %	v1.11.5	01/31/2019,14:42:26	Monitor More
[Instance ID]	Worker Running	Pay-As-You-Go ecs.n1.large	5	10.00% 0.00%	1.78 %	9.10% 9.10%	56.75 %	v1.11.5	01/31/2019,14:42:27	Monitor More
[Instance ID]	Worker Running	Pay-As-You-Go ecs.n1.large	5	10.00% 0.00%	2.38 %	9.10% 9.10%	58.07 %	v1.11.5	01/31/2019,14:42:27	Monitor More

Total: 6 item(s), Per Page: 25 item(s)

d. In the displayed Remove Node dialog box, click OK.

Remove Node

Are you sure you want to delete the following 1 nodes?

☐ Release ECS at the Same Time

Notes:

1. This operation may affect your business. We recommend that you perform this operation during the non-rush hours. For more information, see [Help Documentation](#).
2. After you have removed a node, its containers will not be automatically migrated. Make sure that you have backed up the required data.
3. The drain operation may cause the pods in the node to be automatically recovered in other nodes in the cluster. Make sure the cluster resources are sufficient.
4. Only Pay-As-You-Go-based ECS instances will be released. Other ECS instances will still be billed.

OK Cancel

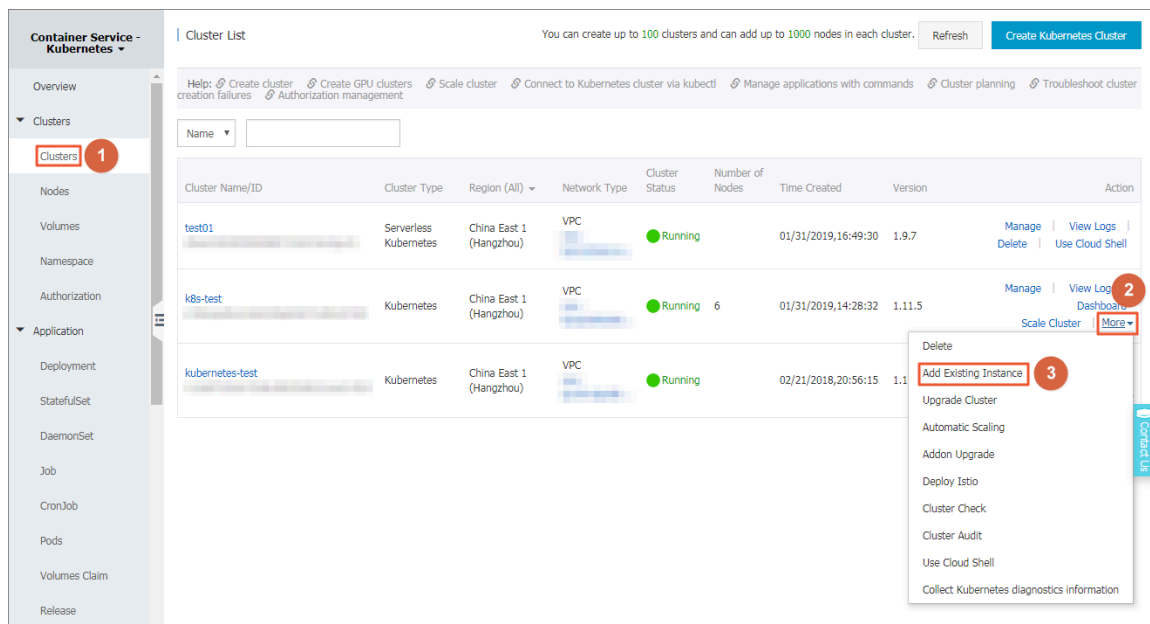


Note:

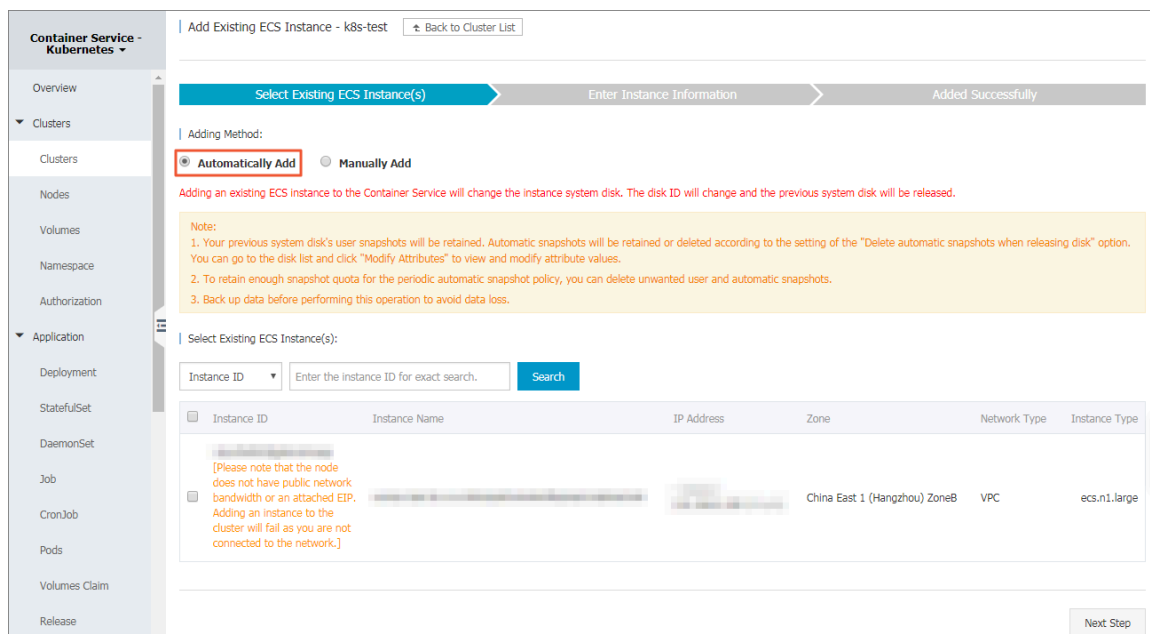
We recommend that you do not select the Release ECS at the same time check box. Otherwise, the ECS instance used by the target node will be released.

4. Add the removed node to the cluster.

- a. In the left-side navigation pane, click Clusters.
- b. On the right of the target cluster, choose More > Add Existing Instance.



- c. Select Automatically Add or Manually Add. In this example, the instance is added automatically.



- d. Select the existing instance and then click Next Step.
- e. Turn on the Format Data Disk switch.

Select Existing ECS Instance(s) Enter Instance Information

Cluster ID/Name :

Information of the cluster to which to add the ECS instance(s).

Format Data Disk : ☒

If an ECS instance already has data disks attached, the first data disk will be automatically formatted and mounted to `/var/lib/docker`.
The original data in the data disk will be lost after formatting. Make sure that you have backed up all important data.
If the ECS instance does not have any data disks attached, no new data disk will be mounted.

Login : ☒ Password

f. Complete other required settings.

After the node has been added to the cluster, you can log on to the node to run the `df` command to check whether a disk has been mounted to the target node.

The following figure shows the disk has been amounted to the target node.

```
[root@xxxxxxxxxxxxxxxxxxxx~]# df
Filesystem      1K-blocks    Used Available Use% Mounted on
/dev/vda1        41151808 2273772  36764604   6% /
devtmpfs         3995592    0      3995592   0% /dev
tmpfs            4005096    0      4005096   0% /dev/shm
tmpfs            4005096    508      4004588   1% /run
tmpfs            4005096    0      4005096   0% /sys/fs/cgroup
/dev/vdb1        101441464 61668  96120584   1% /var/lib/docker
tmpfs            801020     0      801020    0% /run/user/0
```

1.5.10 Mount a disk to the Docker data directory

This topic describes how to mount a disk to the Docker data directory. If the number of containers or images that run on an ECS instance increases constantly, the ECS instance disk capacity may be insufficient. In this case, you can expand the Docker data directory by mounting a disk to the ECS instance.

Docker data directory

Docker data is stored in disks through a union file system (UnionFS). The default container data and image data of Docker is stored in the `/var/lib/docker` directory. You can run the `du` command to view the disk space size occupied by this directory.

```
# du -h --max-depth=0 /var/lib/docker
```

```
7 . 9G      / var / lib / docker
```

Scenarios

Generally, a Docker image occupies a large amount of disk space. If you want to use multiple Docker images or a large number of containers, you must mount a disk to the Docker data directory to ensure sufficient disk capacity is available.

Mount a disk

To mount a disk to the Docker data directory, follow these steps:

1. Create a disk and mount it to the target ECS instance for which you want to expand the disk capacity.
 - a. Log on to the [ECS console](#) to create a disk.
 - b. In the left-side navigation pane, click Instances.
 - c. Click the target ECS instance ID.
 - d. In the left-side navigation pane, click Disks.
 - e. In the upper-right corner, click Mount.
 - f. In the displayed dialog box, select the created disk from the target disk drop-down list, and then click OK.
 - g. Click Mount to mount the new disk to the target ECS instance, and record the new disk mounting point which is in the format of `/ dev / xvd *` or `/ dev / vd *`.

2. Log on to the target ECS instance to format the new disk.

- a. Run the `ls -l /dev/xvd*` or `ls -l /dev/vd*` command to verify whether a disk that has the recorded mounting point has been mounted to the ECS instance.
- b. Run the `fdisk` command to partition the new disk, and then run the `mkfs -ext4` command to format the new disk.

```

root@c836831d69e4040e797eff4d3c4dcd983-node2:~# ll /dev/xvd*
brw-rw---- 1 root disk 202,  0 May 26 15:44 /dev/xvda
brw-rw---- 1 root disk 202,  1 May 26 15:44 /dev/xvda1
brw-rw---- 1 root disk 202, 16 May 27 13:03 /dev/xvdb
root@c836831d69e4040e797eff4d3c4dcd983-node2:~# fdisk -S 56 /dev/xvdb
Device contains neither a valid DOS partition table, nor Sun, SGI or OSF disklabel
Building a new DOS disklabel with disk identifier 0x446953ae.
Changes will remain in memory only, until you decide to write them.
After that, of course, the previous content won't be recoverable.

Warning: invalid flag 0x0000 of partition table 4 will be corrected by w(rite)

Command (m for help): n
Partition type:
   p   primary (0 primary, 0 extended, 4 free)
   e   extended
Select (default p): p
Partition number (1-4, default 1): 1
First sector (2048-62914559, default 2048):
Using default value 2048
Last sector, +sectors or +size[K,M,G] (2048-62914559, default 62914559):
Using default value 62914559

Command (m for help): wq
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.
root@c836831d69e4040e797eff4d3c4dcd983-node2:~# ll /dev/xvd*
brw-rw---- 1 root disk 202,  0 May 26 15:44 /dev/xvda
brw-rw---- 1 root disk 202,  1 May 26 15:44 /dev/xvda1
brw-rw---- 1 root disk 202, 16 May 27 13:08 /dev/xvdb
brw-rw---- 1 root disk 202, 17 May 27 13:08 /dev/xvdb1
root@c836831d69e4040e797eff4d3c4dcd983-node2:~# mkfs.ext4 /dev/xvdb1
mke2fs 1.42.9 (4-Feb-2014)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
1966080 inodes, 7864064 blocks
393203 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=4294967296
240 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000

Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done

```

3. Migrate the Docker data to the new disk.

If you do not want to suspend the applications that run on the target ECS instance, you must migrate the applications. For how to migrate applications on a Swarm cluster, see [Schedule an application to specified nodes](#). For how to migrate applications on a Kubernetes cluster, see [Safely drain a node while respecting application SLOs](#).

- a. To ensure that data can be migrated, run the `service docker stop` command to stop Docker daemon, and run the `service kubelet stop` command to stop kubelet.
- b. Migrate the Docker directory data to a backup directory. For example, `mv /var / lib / docker / var / lib / docker_dat a .`
- c. Mount the new disk to the `/ var / lib / docker` and `/ var / lib / kubelet` directories. For example,

```
echo "/ dev / xvdb1 / var / lib / container / ext4
defaults 0 0 " >> / etc / fstab
echo "/ var / lib / container / kubelet / var / lib / kubelet
none defaults , bind 0 0 " >> / etc / fstab
echo "/ var / lib / container / docker / var / lib / docker
none defaults , bind 0 0 " >> / etc / fstab

mkdir / var / lib / docker

mount - a
```

- d. Migrate the backed up Docker data to the new disk. For example, `mv / var / lib / docker_dat a /* / var / lib / docker /.`

4. Start the Docker daemon and kubelet, and check the data location.
 - a. Run the `service docker start` command to start the Docker daemon, and run the `service kubelet start` command to start kubelet.
 - b. Run the `df` command to verify whether `/var/lib/docker` has been mounted to the new disk. If you need to start the Kubernetes cluster, skip this step.

```
root@c836831d69e4040e797eff4d3c4dcd983-node2:/var/lib# df
Filesystem      1K-blocks    Used Available Use% Mounted on
udev            497280         4    497276   1% /dev
tmpfs           101628        712    100916   1% /run
/dev/xvda1      41151808 1928420    37109960   5% /
none              4           0         4   0% /sys/fs/cgroup
none            5120          0        5120   0% /run/lock
none           508136        288    507848   1% /run/shm
none           102400         0    102400   0% /run/user
/dev/xvdb1      30831612 667168    28575248   3% /var/lib/docker
```

- c. Run the `docker ps` command to check whether containers are lost. Restart containers as needed. For example, you can restart a container that has not been set the `restart : always` label.

```
root@c836831d69e4040e797eff4d3c4dcd983-node2:/var/lib# docker ps
CONTAINER ID        IMAGE                                     COMMAND                  CREATED             STATUS
4f564091bffa        registry.aliyuncs.com/acs/logspout:0.1-41e0e21  "/bin/logspout"         21 hours ago        Up 3 minutes
gspout_2
a5aba5fbadae        registry.aliyuncs.com/acs/ilogtail:0.9.9        "/bin/sh -c 'sh /usr/"  21 hours ago        Up 3 minutes
gtail_2
5e3d8fe154bb        registry.aliyuncs.com/acs/monitoring-agent:0.7-1cf85e6  "acs-mon-run.sh --hel"  21 hours ago        Up 3 minutes
_acs-monitoring-agent_1
fb72c2388b0e        registry.aliyuncs.com/acs/volume-driver:0.7-252cb09  "acs-agent volume_exe"  21 hours ago        Up 3 minutes
er_volumedriver_2
604fcb4ad720        registry.aliyuncs.com/acs/routing:0.7-c8c15f0        "/opt/run.sh"           21 hours ago        Up 3 minutes
uting_1
8fe1d6ed15b5        registry.aliyuncs.com/acs/agent:0.7-6967e86         "acs-agent join --nod"  21 hours ago        Up 3 minutes
999da3883264        registry.aliyuncs.com/acs/tunnel-agent:0.21         "/acs/agent -config=c"  21 hours ago        Up 3 minutes
```

5. If a container has been migrated to other nodes, you can schedule it back to the target node to which you mounted the new disk.

For more information, see [Container Service](#).

1.6 Namespace management

1.6.1 Create a namespace

This topic describes how to create a namespace.

Prerequisites

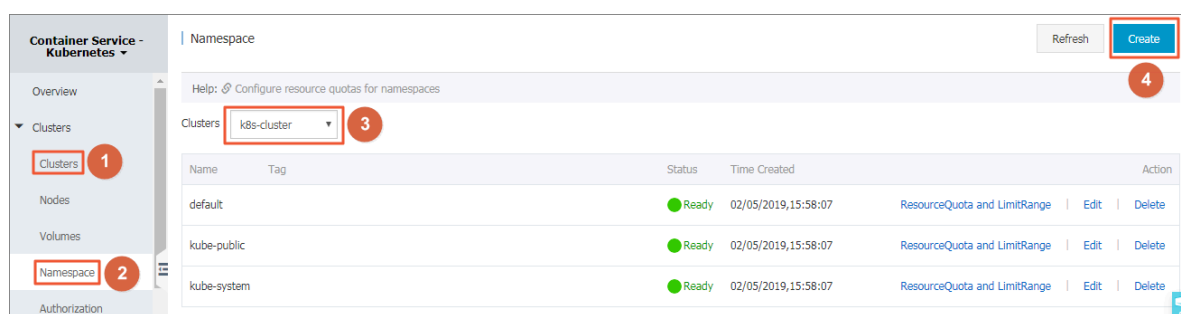
You have created a Kubernetes cluster. For more information, see [Create a Kubernetes cluster](#).

Context

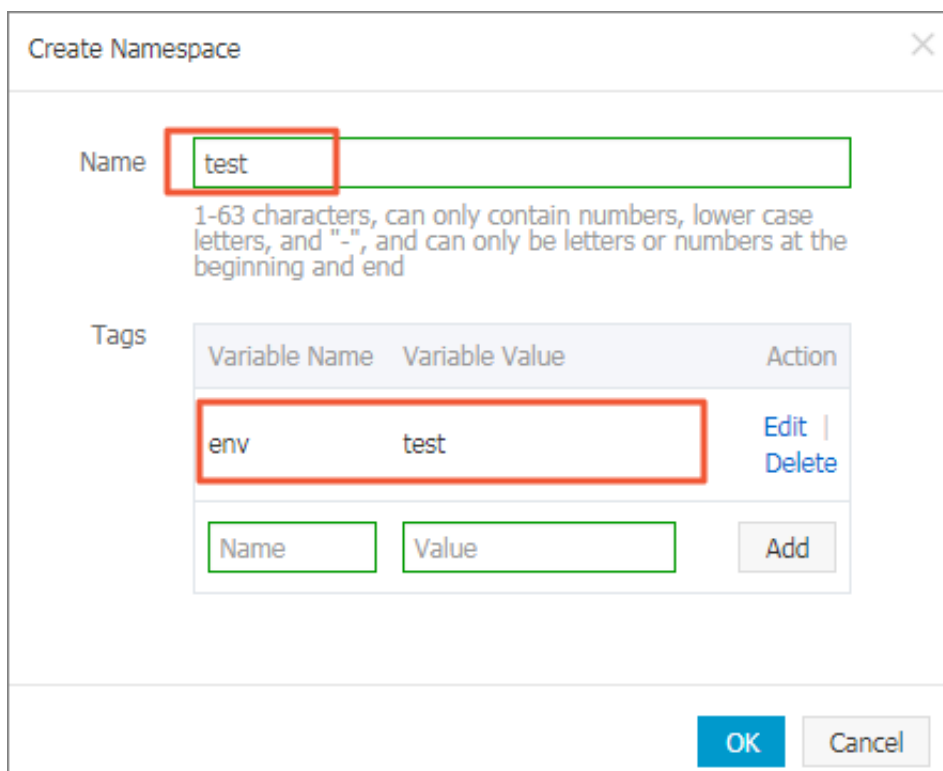
In a Kubernetes cluster, you can use namespaces to create multiple virtual spaces. When a large number of users share a cluster, multiple namespaces can be used to effectively divide different work spaces and assign cluster resources to different tasks. Furthermore, you can use [resource quotas](#) to assign resources to each namespace.

Procedure

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Kubernetes, choose Clusters > Namespace.
3. Select the target cluster, and then click Create in the upper-right corner.



4. In the displayed dialog box, set a namespace.



The dialog box titled "Create Namespace" contains the following elements:

- Name:** A text input field containing "test". Below it, a note states: "1-63 characters, can only contain numbers, lower case letters, and '-', and can only be letters or numbers at the beginning and end".
- Tags:** A table with columns "Variable Name", "Variable Value", and "Action".

Variable Name	Variable Value	Action
env	test	Edit Delete

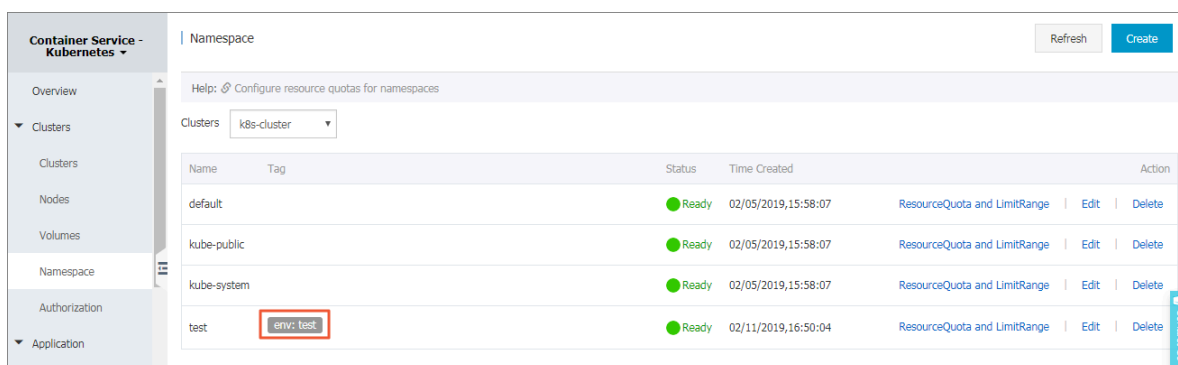
 Below the table are input fields for "Name" and "Value", and an "Add" button.
- Buttons:** "OK" and "Cancel" buttons at the bottom right.

- **Name:** Enter a name for the namespace name. The name must be 1 to 63 characters in length and can contain numbers, letters, and hyphens (-). It must start and end with a letter or number. In this example, test is used as the name.
- **Tags:** Add one or multiple tags to the namespace to identify the characteristics of the namespace. For example, you can set a tag to identify that this namespace is used for the test environment.

You can enter a variable name and a variable value, and then click Add on the right to add a tag to the namespace.

5. Click OK.

6. The namespace named test is displayed in the namespace list.



The screenshot shows the "Namespace" page in the "Container Service - Kubernetes" console. The table lists the following namespaces:

Name	Tag	Status	Time Created	Action
default		Ready	02/05/2019,15:58:07	ResourceQuota and LimitRange Edit Delete
kube-public		Ready	02/05/2019,15:58:07	ResourceQuota and LimitRange Edit Delete
kube-system		Ready	02/05/2019,15:58:07	ResourceQuota and LimitRange Edit Delete
test	env: test	Ready	02/11/2019,16:50:04	ResourceQuota and LimitRange Edit Delete

1.6.2 Set resource quotas and limits for a namespace

This topic describes how to set resource quotas and limits for a namespace through the Container Services console.

Prerequisites

- You have created a Kubernetes cluster. For more information, see [Create a Kubernetes cluster](#).
- You have created a namespace. In this topic, a namespace named test is used. For more information, see [Create a namespace](#).
- You have connected to the Master node of the cluster. For more information, see [Connect to a Kubernetes cluster by using kubectl](#).

Context

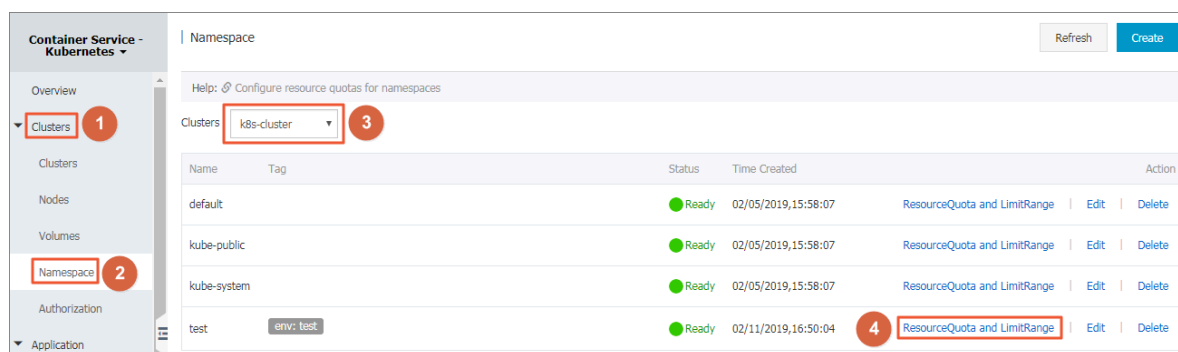
By default, a running pod uses the CPU and memory resources of nodes without limit. That is, any pod can use the computing resources of the cluster without restraint. Therefore, pods of a namespace may deplete the cluster resources.

Namespaces can be used as virtual clusters to serve multiple users. Therefore, setting resource quotas for a namespace is regarded as a best practice.

For a namespace, you can set the quotas of resources, such as CPU, memory, and number of pods. For more information, see [Resource quotas](#).

Procedure

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Kubernetes, choose Clusters > Namespace. Select the target cluster and click ResourceQuota and LimitRange on the right of the test namespace.



3. In the displayed dialog box, set resource quotas and default resource limits.



Note:

After setting CPU/memory quotas for a namespace, you must specify CPU/memory resource limits or set the default resource limits for the namespace when creating a pod. For more information, see [Resource quotas](#).

a) Set resource quotas for the namespace.

ResourceQuota and LimitRange

Tip: After setting the CPU/memory quota (ResourceQuota) for the namespace, you must specify the CPU/memory resource limit when configuring Pods, or configure the default resource limit (LimitRange) for the namespace. For details, please refer to: [Resource Quotas](#)

ResourceQuota
LimitRange

^ Compute Resource Quota

☒ CPU Limit
Total
2
Core(s)

☒ Memory Limit
Total
4Gi
?

^ Storage Resource Quota

☒ storage
Total
1024Gi
?

☒ persistentvolumeclaims
Total
50

^ Object Count Quota

☒ configmaps
Total
100

☒ pods
Total
50

☒ services
Total
20

☒ services.loadbalancers
Total
5

☒ secrets
Total
10

OK
Cancel

b) To control the amount of resources consumed by containers, set resource limits and resource requests for containers in this namespace. For more information, see <https://kubernetes.io/memory-default-namespace/>.

ResourceQuota and LimitRange

Tip: After setting the CPU/memory quota (ResourceQuota) for the namespace, you must specify the CPU/memory resource limit when configuring Pods, or configure the default resource limit (LimitRange) for the namespace. For details, please refer to: [Resource Quotas](#)

ResourceQuota
LimitRange

	CPU		Memory ?
Limit	0.5	Core(s)	512Mi
Request	0.1	Core(s)	256Mi

OK
Cancel

4. Connect to the Master node and then run the following commands to view the resources of the test namespace:

```
# kubectl get limitrange , ResourceQuota -n test
NAME      AGE
limitrange / limits    8m

NAME      AGE
resourcequota / quota  8m

# kubectl describe limitrange / limits resourcequota /
quota -n test
Name : limits
Namespace : test
Type   Resource   Min   Max   Default   Request   Default   Limit
Max   Limit / Request   Ratio
-----
Container  cpu - - 100m  500m -
Container  memory - - 256Mi  512Mi -

Name : quota
Namespace : test
Resource   Used   Hard
-----
configmaps  0    100
limits . cpu  0    2
limits . memory  0    4Gi
persistent volumeclaim  0    50
pods  0    50
requests . storage  0    1Ti
secrets  1    10
services  0    20
```

```
services . loadbalancers 0 5
```

1.6.3 Edit a namespace

This topic describes how to edit a namespace.

Prerequisites

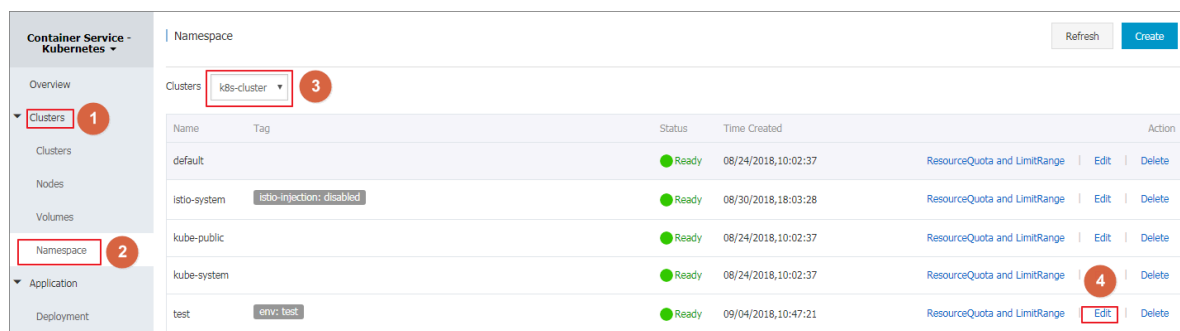
- You have created a Kubernetes cluster. For more information, see [Create a Kubernetes cluster](#).
- You have created a namespace. In this topic, a namespace named test is used. For more information, see [Create a namespace](#).

Context

Editing a namespace means to add, modify, or delete the details of a namespace tag.

Procedure

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Kubernetes, choose Clusters > Namespace.
3. Select the target cluster and then click Edit on the right of the target namespace tag.



- In the displayed dialog box, click Edit to modify the namespace tag. For example, change the tag to `env : test - V2` and click Save.

Edit Namespace

Name:

1-63 characters, can only contain numbers, lower case letters, and "-", and can only be letters or numbers at the beginning and end

Tags:

Variable Name	Variable Value	Action
env	test-V2	Save Delete

- Click OK. The edited namespace tag is then displayed in the namespace list.

Name	Tag	Status	Time Created	Action
default		Ready	08/24/2018,10:02:37	ResourceQuota and LimitRange Edit Delete
istio-system	istio-injection: disabled	Ready	08/30/2018,18:03:28	ResourceQuota and LimitRange Edit Delete
kube-public		Ready	08/24/2018,10:02:37	ResourceQuota and LimitRange Edit Delete
kube-system		Ready	08/24/2018,10:02:37	ResourceQuota and LimitRange Edit Delete
test	env: test-V2	Ready	09/04/2018,10:47:21	ResourceQuota and LimitRange Edit Delete

1.6.4 Delete a namespace

This topic describes how to delete a namespace you no longer require.

Prerequisites

- You have created a Kubernetes cluster. For more information, see [Create a Kubernetes cluster](#).
- You have created a namespace. In this topic, a namespace named test is used. For more information, see [Create a namespace](#).

Context

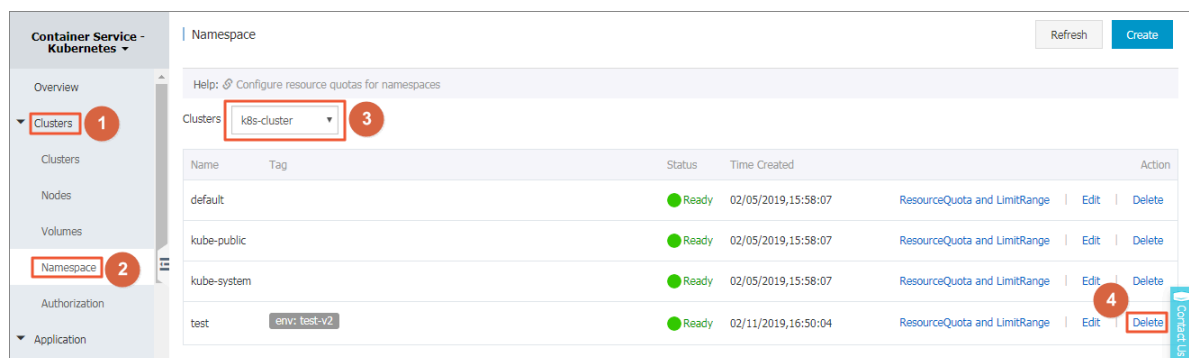


Note:

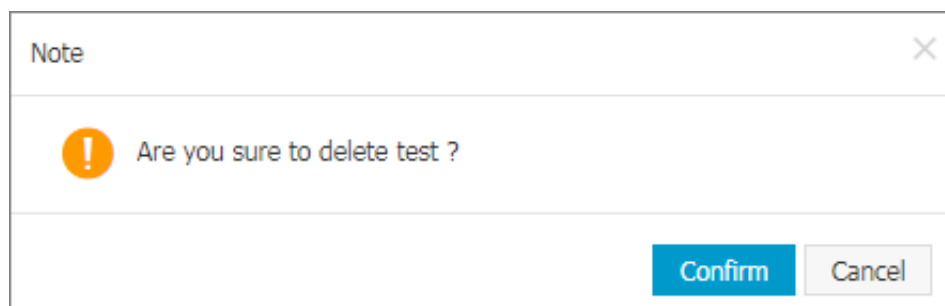
Deleting a namespace also deletes all of its resource objects. Exercise caution when performing this action.

Procedure

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Kubernetes, choose Clusters > Namespace.
3. Select the target cluster and then click Delete on the right of the cluster.



4. In the displayed dialog box, click Confirm.



5. The namespace is then deleted from the namespace list, and its resource objects are also deleted.

1.7 Service catalog management

1.7.1 Overview

Applications running on the cloud platform need some basic services such as databases, application servers, and other generic basic softwares. For example, a WordPress application, as a Web application, needs a database service (such as MariaDB) in the backend. Traditionally, you can create the MariaDB service on which

the application depends in the WordPress application orchestration, and integrate the MariaDB service with the Web application. To develop applications on the cloud in this way, developers must spend time and energy deploying and configuring the dependent infrastructure softwares, which increases the costs of hosting and migrating applications.

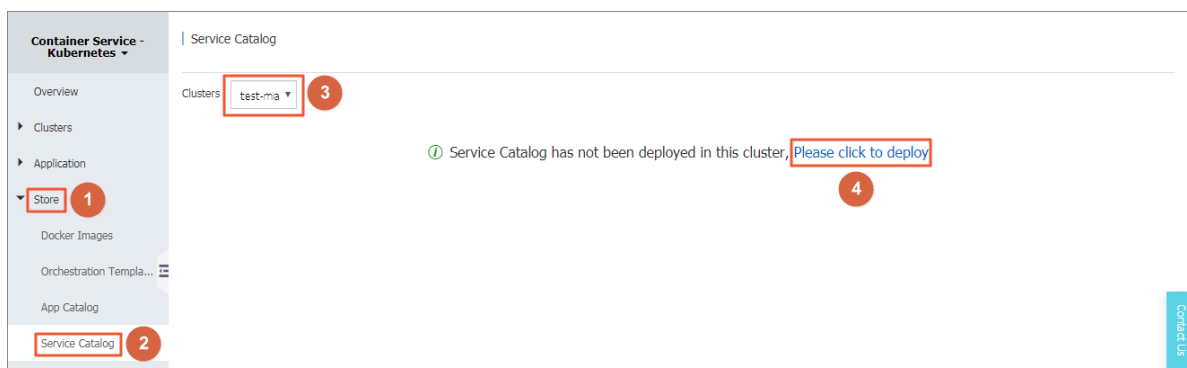
Alibaba Cloud Container Service supports and integrates with the service catalog function. The service catalog function aims to access and manage the service brokers , which allows applications running in Kubernetes clusters to use the managed services offered by service brokers. A series of infrastructure softwares are supported by the service catalog function, which allows the developers to use these softwares as services and focus on the applications, the core of the development, without concerning about the availability and scalability of the softwares or managing the softwares.

The service catalog uses the Open service broker API of Kubernetes to communicate with service brokers, acting as an intermediary for the Kubernetes API server to negotiate the initial provisioning and obtain the credentials necessary for the applications to use the managed services. For more information about the implementation principle of the service catalog, see [Service catalog](#).

1.7.2 Enable service catalog function

Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Store > Service Catalog in the left-side navigation pane. Select the cluster from the Cluster drop-down list in the upper-right corner.
3. If you have not deployed the service catalog, click to install the service catalog as instructed on the page.

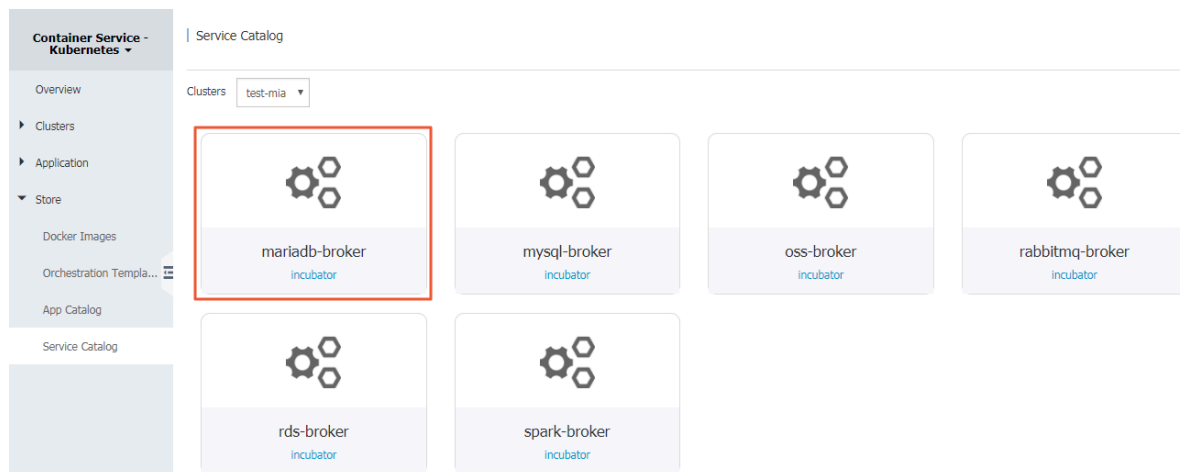


4. After the installation, the service broker, which is installed by default, is displayed on the Service Catalog page. You can click the mariadb-broker to view the details.

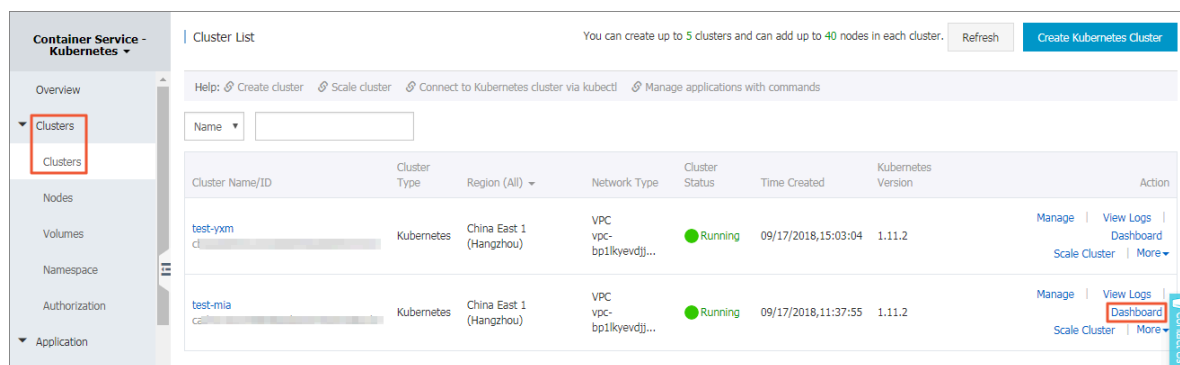


Note:

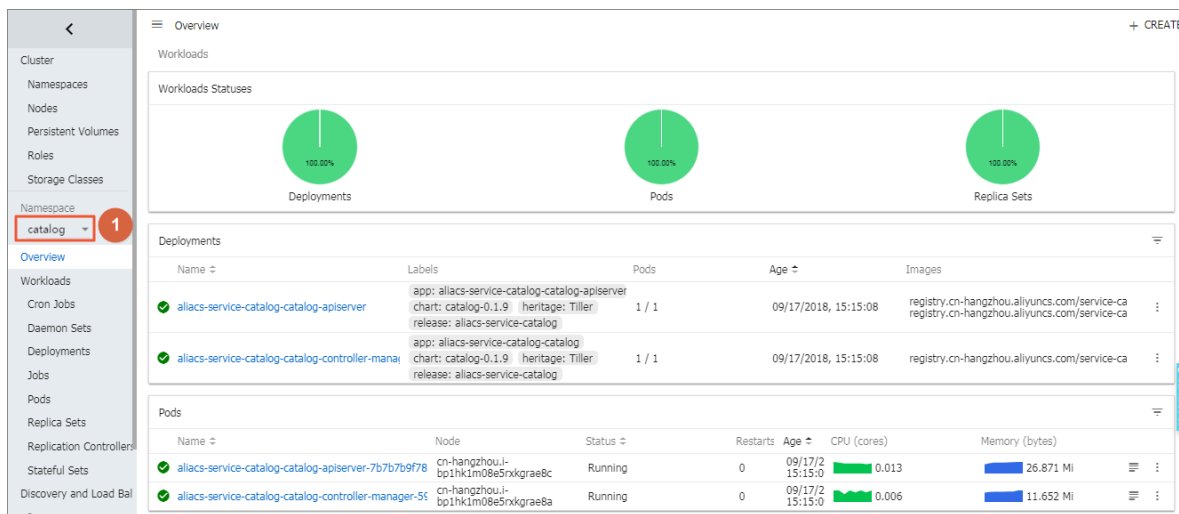
The service catalog is implemented as an extension API server and a controller. After Alibaba Cloud Container Service installs the service catalog function, the namespace catalog is created.



5. Click Clusters in the left-side navigation pane. Click Dashboard at the right of a cluster.



6. In the Kubernetes dashboard, select `catalog` as the Namespace in the left-side navigation pane. You can see the resource objects related to catalog apiserver and controller are installed under this namespace.



What's next

Then, you have successfully enabled the service catalog function. You can create a managed service by using the service broker in the service catalog, and apply the managed service to your applications.

1.8 Application management

1.8.1 Create a deployment application by using an image

This topic describes how to use an image to create a deployment application. In this topic, an Nginx application that is accessible to the Internet is created.

Prerequisites

A Kubernetes cluster is created. For more information, see [Create a Kubernetes cluster](#).

Procedure

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Kubernetes, choose Application > Deployment, and then click Create by Image in the upper-right corner.

3. Set Name, Cluster, Namespace, Replicas, Type, Tag, and Annotation. The replicas parameter indicates the number of pods contained in the application. Then click Next.

**Note:**

In this example, you need to select the Deployment type.

If you do not set Namespace, the system automatically uses the default namespace.

4. Configure a container.

**Note:**

You can configure multiple containers for the pod of the application.

- a) Set general container parameters.

- **Image Name:** Click Select image to select the image in the displayed dialog box and then click OK. In this example, select the Nginx image.

You can also enter a private registry in the format of `domainname / namespace / imagename : tag` to specify an image.

- **Image Version:** Click Select image version to select a version. If you do not select an image version, the system uses the latest version by default.
- **Always pull image:** Container Service caches the image to improve deployment efficiency. During deployment, if the tag of the newly specified image is the same as that of the cached image, Container Service reuses the cached image, instead of re-pulling the same image. Therefore, if you do not

modify the image tag when changing your code and image, the early image in the local cache is used in the application deployment. If you select this check box, Container Service ignores the cached image and re-pulls an image when deploying the application to make sure the latest image and code are always used.

- **Image pull secret:** Create a Secret for the image. A secret is required to pull a image from a private image repository. For more information, see [Use an image Secret](#).
- **Resource Limit:** Specify the upper limit for the resources (CPU and memory) that can be used by this application to avoid occupying excessive resources. CPU is measured in millicores, that is, one thousandth of one core. Memory is measured in bytes, which can be Gi, Mi, or Ki.
- **Resource Request:** Specify how many resources (CPU and memory) are reserved for the application. These resources can be set to be exclusive to the container by using this parameter. If you do not set this parameter, other

services or processes will compete for resources. Then the application may become unavailable due to resource shortage.

- **Init Container:** Select this check box to create an Init Container that contains useful tools. For more information, see [Init containers](#).

The screenshot shows the 'Container' configuration page with the following details:

- Tab:** Container (selected)
- Section:** General
- Image Name:** nginx
- Image Version:** latest
- Always pull image:** ☐ (checked)
- Image pull secret:** (empty)
- Resource Limit:** CPU: 500m, Core: (empty), Memory: 128Mi
- Resource Request:** CPU: 500m, Core: (empty), Memory: 128Mi
- Init Container:** ☐ (unchecked)

b) Optional: Set environment variables.

You can use key-value pairs to set environment variables for the pods.

Environment variables are used to add environment labels or pass configurations for the pods. For more information, see [Pod variable](#).

c) Optional: Set health checks.

You can set liveness probes and readiness probes. Liveness probes are used to detect when to restart the container. Readiness probes determine if the

container is ready to receive traffic. For more information about health checks, see [Configure liveness and readiness probes](#).

Health Check

Liveness

☒ Enable

HTTP

TCP

Command

▼

Protocol

HTTP

▼

path

Port

Http Header

name

value

Initial Delay

3

Period

10

Timeout

1

Success Threshold

1

Failure Threshold

3

Readiness

☒ Enable

HTTP

TCP

Command

▼

Protocol

HTTP

▼

path

Port

Http Header

name

value

Initial Delay

3

Period

10

Timeout

1

Request method	Description
HTTP request	<p>With this health check method, you can send an HTTP GET request to the container. The following parameters are supported:</p> <ul style="list-style-type: none"> • Protocol: HTTP/HTTPS. • Path: the path to access the HTTP server. • Port: the number or name of the access port exposed by the container. The port number must be in the range of 1 to 65535. • HTTP Header: custom headers in the HTTP request. HTTP allows repeated headers. You can use a key-value pair to set an HTTP Header. • Initial Delay (in seconds): the initialDelaySeconds parameter, indicating the number of seconds for which the first probe must wait after the container is started. The default value is 3. • Period (in seconds): the periodseconds parameter, indicating the interval at which probes are performed. The default value is 10. The minimum value is 1. • Timeout (in seconds): the timeoutSeconds parameter, indicating the number of time that the probe has timed out. The default value is 1 and the minimum value is 1. • Success Threshold: The minimum number of consecutive successful probes needed for determining a probe success after a failed probe. The default value is 1 and the minimum value is 1. It must be set to 1 for a liveness probe. • Failure Threshold: The minimum number of consecutive failed probes needed for determining a probe failure after a successful probe. The default value is 3. The minimum value is 1.

Request method	Description
TCP connection	<p>If you use this health check method, a TCP socket is sent to the container. The kubelet then attempts to open the socket of the container on a specified port. If a connection can be established, the container is considered healthy. If not, it is considered unhealthy. The following parameters are supported:</p> <ul style="list-style-type: none"> • Port: the number or name of the access port exposed by the container. The port number must be in the range of 1 to 65535. • Initial Delay (in seconds): the initialDelaySeconds parameter, indicating the seconds for the first liveness or readiness probe must wait for after the container is started. The default value is 15. • Period (in seconds): the periodseconds parameter, indicating the interval at which probes are performed. The default value is 10. The minimum value is 1. • Timeout (in seconds): the timeoutSeconds parameter, indicating the number of time that the probe has timed out. The default value is 1 and the minimum value is 1. • Success Threshold: The minimum number of consecutive successful probes needed for determining a probe success after a failed probe. The default value is 1 and the minimum value is 1. It must be set to 1 for a liveness probe. • Failure Threshold: The minimum number of consecutive failed probes needed for determining a probe failure after a successful probe. The default value is 3. The minimum value is 1.

Request method	Description
Command line	<p>With this health check method, you can detect the container health by executing a probe detection command in the container. The following parameters are supported:</p> <ul style="list-style-type: none"> • Command: a probe command used to detect the container health. • Initial Delay (in seconds): the <code>initialDelaySeconds</code> parameter, indicating the number of seconds for which the first liveness or readiness probe must wait after the container is started. The default value is 5. • Period (in seconds): the <code>periodSeconds</code> parameter, indicating the interval at which probes are performed. The default value is 10. The minimum value 1. • Timeout (in seconds): the <code>timeoutSeconds</code> parameter, indicating the number of time that the probe has timed out. The default value is 1 and the minimum value is 1. • Success Threshold: The minimum number of consecutive successful probes needed for determining a probe success after a failed probe . The default value is 1 and the minimum value is 1. It must be set to 1 for a liveness probe. • Failure Threshold: The minimum number of consecutive failed probes needed for determining a probe failure after a successful probe. The default value is 3. The minimum value is 1.

d) Set life cycle rules.

You can set the following parameters for the container life cycle: start, post start, and pre-stop. For more information, see [Attach handlers to container lifecycle events](#).

- **Start:** Set a pre-start command and parameter for the container.
- **Post Start:** Set a post-start command for the container.
- **Pre Stop:** Set a pre-stop command for the container.

Life cycle	Start:	Command	<code>["/bin/sh","-c","echo Hello > /user/share/message"]</code>
		Parameter	
	Post Start:	Command	
	Pre Stop:	Command	<code>["/user/sbin/nginx","-s","quit"]</code>

e) Optional: Set volumes.

You can configure local storage and cloud storage.

- **local storage:** Supported storage types include HostPath, ConfigMap, Secret, and EmptyDir. By setting a type of local storage, you can mount its mount source to the container path. For more information, see [Volumes](#).
- **cloud storage:** Supported types of cloud storage include disks, Network Attached Storage (NAS), and Object Storage Service (OSS).

This example sets a disk as the volume and mounts the disk to the `/tmp` container path. Then container data generated in this path is stored to the disk.

Data Volume:		
+ Add local storage		
Storage type	Mount source	Container Path
+ Add cloud storage		
Storage type	Mount source	Container Path
Disk	pvc-disk	/tmp

f) Optional: Set Log Service. You can set collection parameters and customize tags.



Note:

Make sure that you have deployed a Kubernetes cluster and installed the log plugin on the cluster.

Set the following log collection parameters:

- **Log Store:** Set a Logstore. After you specify the Logstore name, the Logstore is generated in Log Service to store collected logs.
- **Log path in the container:** Set this parameter to stdout or set a log path.
 - **stdout:** If you set the log path parameter to stdout, you can collect the standard output logs of the container.
 - **text log:** If you specify a container log path, you can collect the text logs of the path. Wildcards can be used in setting the log file name for a log path. In this example, text logs in the path of /var/log/nginx are collected.

You can also customize log tags. The customized log tags can be collected together with container output logs and can benefit log analysis actions such as collecting log statistics and filtering specific logs.

Log Service:

Note: please ensure that cluster has deployed log plug-ins.

Log Configuration

Configuration

Log Store	Log path in the container (can be set to stdout)	
catalina	stdout	-
access	/var/log/nginx	-

Custom Tag

Name Of Tag	Value Of Tag	
app	nginx	-

5. Click Next.

6. Configure advanced settings.

a) Set Access Control.

You can set the methods to expose the application pod and then click Create. In this example, a cluster IP service and an Ingress are set to create an Nginx application that is accessible for the Internet.



Note:

You can set access methods according to the communication requirements of your application.

- **Internal application** : an application that works only inside the cluster. You can create a cluster IP service or a node port service as needed for communication within the cluster.

- **External application** : an application that needs to be exposed to the Internet. You can set how the application is accessed by using either of the following two methods:
 - Create a Server Load Balancer service. This method uses Alibaba Cloud Server Load Balancer (SLB) to provide Internet accessibility for the application.
 - Create a cluster IP service or a node port service, and create an Ingress. This method provides Internet accessibility through the Ingress. For more information, see [Ingress](#).

Basic Information		Container	Advanced	Done
Access Control	Service(Service)	Create		
	Ingress(Ingress)	Create		
Scale	HPA	<input type="checkbox"/> Enable		
Scheduling	Node Affinity	Add		
	Pod Affinity	Add		
	Pod Anti Affinity	Add		
			Prev	Create

A. Click Create on the right of Service. Configure a service in the displayed dialog box, and then click Create.

Create Service

Name:

nginx-svc

Type:

Server Load Balancer

public

Port Mapping:

+ Add

Name	service port	Container Port	Protocol
nginx-svc	80	80	TCP

Annotation:

+ Add Annotations for load balancer

Tag:

+ Add

Create

Cancel

- **Name:** Enter the service name. The default is `applicationname - svc`.
- **Type:** Select one service type.
 - **Cluster IP:** Exposes the service by using the internal IP address of your cluster. If you select this service type, the service is accessible only within the cluster.
 - **Node port:** Exposes the service by using the IP address and the static port (NodePort) of each node. A node port service routes to a cluster IP service that is automatically created. You can access the node port service from outside the cluster by requesting `< NodeIP > : < NodePort >`.
 - **Server Load Balancer:** Alibaba Cloud Server Load Balancer service. With this type of service, you can set an Internet or intranet access method

for your application. SLB can route to a node port service and a cluster IP service.

- **Port Mapping:** Add a service port and a container port, and select the TCP or UDP protocol. If you select the node port Type, you must add a node port to avoid port conflict.
- **annotation:** Add an annotation to the service. You can set SLB parameters. For more information, see [Access services by using Server Load Balancer](#).
- **Tag:** Add a tag to the service to identify the service.

B. Click Create on the right of Ingress. In the displayed dialog box, configure an Ingress rule for the application pod, and then click Create. For more information, see [Ingress configurations](#).



Note:

When you create an application by using an image, you can create an Ingress rule for only one service. In this example, a virtual host name is used as the

test domain name. You need to add a record to the host. You must use a filing domain name when you create your application.

101 . 37 . 224 . 146 foo . bar . com # This is the
IP address of the Ingress .

Create

Name:

Rule:

+ Add

Domain

foo.bar.com

Select *

ainer.com or Custom

path

e.g./

Services

+ Add

Name

Port

nginx-svc

80

-

EnableTLS

Service weight:

Enable

Grayscale release:

+ Add

After the gray rule is set, the request meeting the rule will be routed to the new service. If you set a weight other than 100, the request to satisfy the gamma rule will continue to be routed to the new and old version services according to the weights.

Annotation:

+ Add

rewrite annotation

Tag:

+ Add

Create

Cancel

C. In the access control area, the created service and Ingress are displayed. You can perform further configurations by clicking Update or Delete.

Issue: 20190614

171

Create Application

Basic Information > Container > **Advanced** > Done

Services(Service) [Update](#) [Delete](#)

service port	Container Port	Protocol
80	80	TCP

Ingresses(Ingress) [Update](#) [Delete](#)

Domain	path	Name	service port
foo.bar.com		nginx-svc	80

Scale

HPA ☐ Enable

Update Method ☒ Enable

☒ RollingUpdate ☐ OnDelete

Max Unavailable: 25 %

Max Surge: 25 %

Node Affinity [Add](#)

Pod Affinity [Add](#)

Pod Anti Affinity [Add](#)

b) Optional: Set Horizontal Pod Autoscaling (HPA).

You enable HPA by selecting the Enable check box. Alibaba Cloud Container Service for Kubernetes provides pod auto scaling to deal with different application workloads. That is, you can change the number of pods according to the container CPU and memory usage.

Scale

HPA ☒ Enable

Metric: CPU Usage

Condition: Usage 70 %

Maximum Replicas: 10 Range : 2-100

Minimum Replicas: 1 Range : 1-100



Note:

To use this function, you must set required resources for the pod. Otherwise, pod auto scaling cannot take effect. For more information, see general container settings.

- **Metric:** resource type. CPU or memory is available. This parameter must be specified with a resource type that is the same as the required resource type.
- **Condition:** the percentage value of resource usage. The number of containers increases when the resource usage exceeds this value.
- **Maximum Replicas:** the maximum number of the containers that the deployment can include.
- **Minimum Replicas:** the minimum number of the containers that the deployment can include.

c) Optional: Set Scheduling.

You can set an update method, node affinity, pod affinity, and pod anti affinity. For more information, see [Affinity and anti-affinity](#).



Note:

Affinity scheduling depends on node tags and pod tags. You can use built-in or customized tags to schedule nodes or pods.

A. Set Update Method.

You can select the `RollingUpdate` or `Recreate` (OnDelete) method to replace old pods with new ones. For more information, see [Deployments](#).

B. Set Node Affinity by using node tags.

Create

Required:

+

Add Rule

Selector

+

Add

Tag Name	Operator	Tag Value
kubernetes.io/hostname	In	
kubernetes.io/hostname	In	

Preferred:

+

Add Rule

Weight

100

Selector

+

Add

Tag Name	Operator	Tag Value
kubernetes.io/hostname	NotIn	

OK

Cancel

Required rules and preferred rules are supported, and available operators include `In` , `NotIn` , `Exists` , `DoesNotExist` , `Gt` , and `Lt` .

- Required rules must be satisfied and correspond to `requiredDuringSchedulingIgnoredDuringExecution` . The required rules

have the same effect as `NodeSelect` or `NodeAffinity`. In this example, the pod can be scheduled to only a node with the specified tags.

You can add multiple required rules, but only one required rule needs to be satisfied for pod scheduling.

- Preferred rules can be unnecessarily satisfied and correspond to `preferredDuringSchedulingIgnoredDuringExecution`. With the scheduling setting in this example, the system tries not to schedule the pod to the nodes with the specified tag.

You can also set `Weight` for each preferred rule. If multiple nodes satisfies the preferred rules, the system schedules the pod to a node with the highest weight.

You can add multiple preferred rules, and all the rules must be satisfied for pod scheduling.

- C. Set Pod Affinity to deploy the application pod in a topology domain together with other pods. For example, to reduce network latency between the services

that communicate with each other, you can deploy their pods to a topology domain (for example, a host).

The screenshot shows a 'Create' dialog for configuring pod scheduling rules. It is divided into two sections: 'Required' and 'Preferred'.

Required Rule Configuration:

- Namespace:** default (highlighted with a red box and number 1)
- Topology Key:** kubernetes.io/hostname
- Selector:** Add (highlighted with a red box and number 2)
- Tag Configuration:**

Tag Name	Operator	Tag Value
app (highlighted with a green box and number 3)	In	nginx

Preferred Rule Configuration:

- Weight:** 100
- Namespace:** default
- Topology Key:** kubernetes.io/hostname
- Selector:** Add
- Tag Configuration:**

Tag Name	Operator	Tag Value
app	NotIn	wordpress

At the bottom right, there are 'OK' and 'Cancel' buttons.

You can schedule pods according to tags of pods running on nodes. Required rules and preferred rules are supported, and available operators include `In`, `NotIn`, `Exists`, `DoesNotExist`.

- Required rules must be satisfied and correspond to `requiredDuringSchedulingIgnoredDuringExecution`. All specified conditions of required rules must be met for pod affinity scheduling.
- Namespace: Set a namespace. This parameter is required because the scheduling policy is based on pod tags.

- **Topology Key:** Set a topology domain to which pods are scheduled. This parameter takes effect through node tags. For example, if you set `kubernetes . io / hostname` as the topology key, a node is used to identify a topology. If you set `beta . kubernetes . io / os` as the topology key, a node operating system is used to identify a topology.
- **Selector:** Click this button to add a required rule.
- **View Application List:** Click View Application List, a dialog box is displayed. In the dialog box, you can view applications in each namespace and export application tags to the dialog box in which you set pod affinity.
- **Required rule tag:** Set a tag name, its operator, and the tag value for existing applications. This example schedules the application to be created to a host on which applications tagged with `app : nginx` run.
- Preferred rules can be unnecessarily satisfied and correspond to `preferredDuringSchedulingIgnoredDuringExecution`. Specified conditions of required rules will be met as many as possible for pod affinity scheduling.

You can set Weight for each preferred rule. The weight value range is 1 to 100. If multiple nodes satisfies the preferred rules, the system schedules the pod to a node with the highest weight. Other parameters are the same with the required rule setting.

D. Set Pod Anti Affinity to deploy the application pods in a topology domain that excludes other pods. Scenarios that use pod anti affinity scheduling include:

- Distribute the pods of a service to different topology domains (for example , different hosts) to improve the service stability.
- Grant a pod the exclusive access to a node so as to guarantee that no other pods use the resources of the node.
- Distribute pods of the services that may affect each other to different hosts.



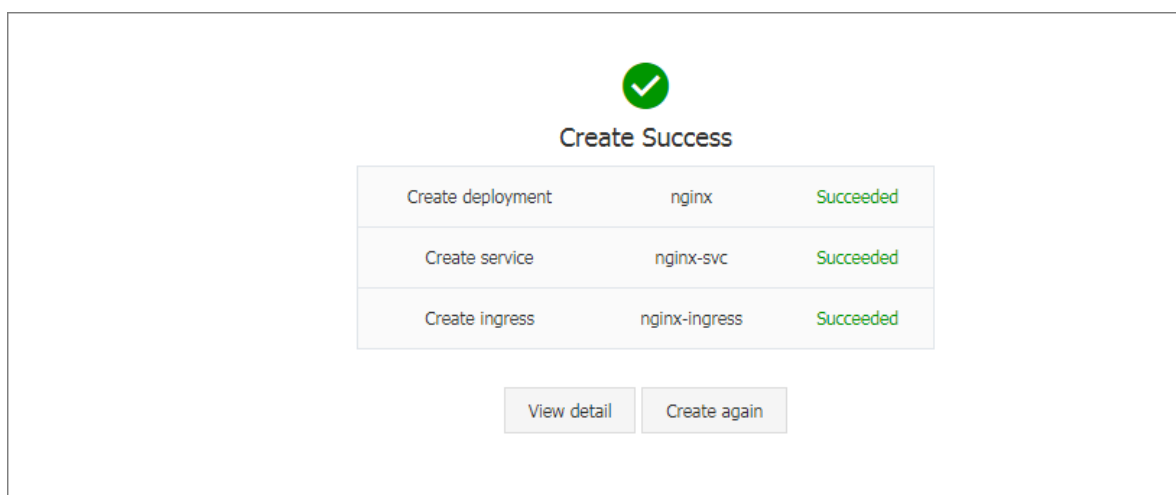
Note:

You can set pod anti affinity scheduling by using the same method as setting pod affinity scheduling. But the same scheduling rules have different

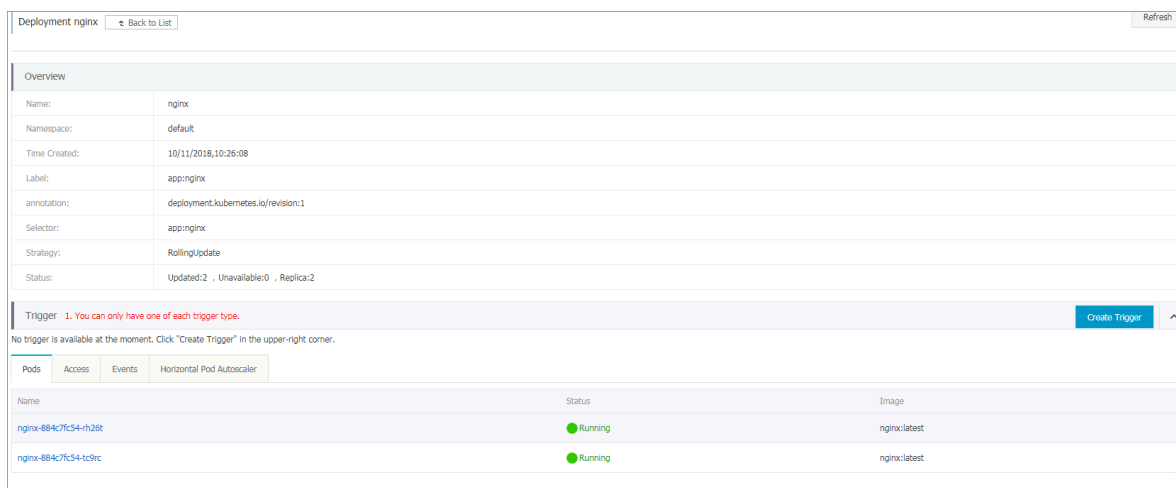
meanings for these two types of scheduling. You need to select appropriate scheduling rules as needed.

7. Click Create.

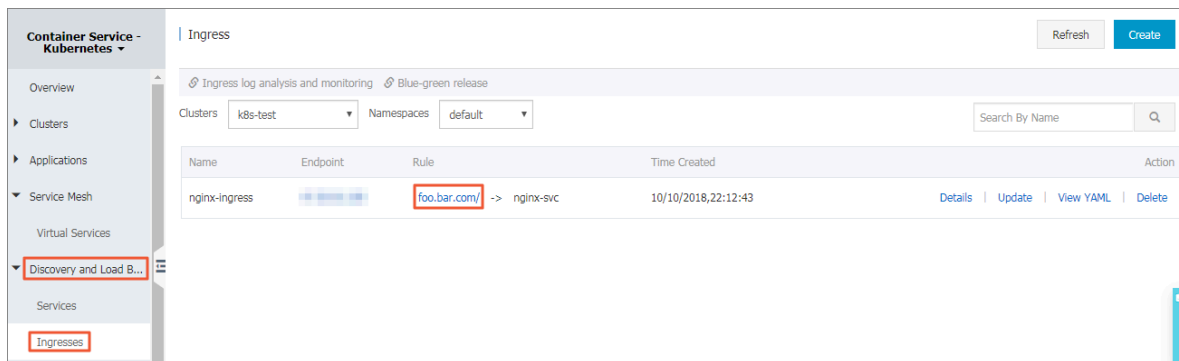
8. After you create the application, a new page is displayed by default to prompt that you have created the application and lists objects included in the application. You can click View detail to view the deployment details.



The nginx-deployment page is displayed by default.



9. Choose Discovery and Load Balancing > Ingress to verify that a rule is displayed in the Ingress list.



10. Access the test domain name in your browser to verify that you can visit the Nginx welcome page.



1.8.2 Create a StatefulSet application by using an image

Kubernetes clusters of Alibaba Cloud Container Service allows you to quickly create applications of the StatefulSet type through the web interface. In this example, create a StatefulSet Nginx application and show features of a StatefulSet application.

Prerequisites

- You have created a Kubernetes cluster. For more information, see [Create a Kubernetes cluster](#).
- You have successfully created a cloud disk storage volume claim. For more information, see [Create a persistent volume claim](#).
- You have successfully connected to the master node of the Kubernetes cluster. For more information, see [Connect to a Kubernetes cluster by using kubectl](#).

Context

StatefulSet features are as follows:

Scenarios	Description
Pod consistency	Contains order (such as startup and stop order) and network consistency. This consistency is related to pods and has nothing to do with the node to which the pods are to be scheduled.
Stable persistent storage	Create a PV for each pod through VolumeClaimTemplate. Deleting or reducing replicas does not delete relevant volumes.
Stable network marker	The hostname mode for a pod is: (statefulset name)-(sequence number).
Stable order	For StatefulSet of N replicas, each pod is assigned a unique order number within the range of 0 to N.

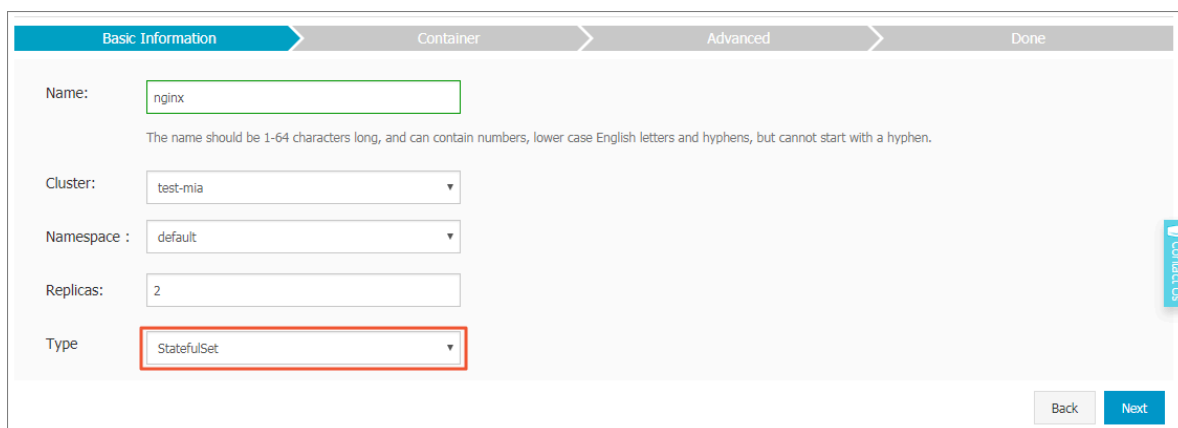
Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Application > Deployment in the left-side navigation pane, and then click Create by image in the upper-right corner.
3. Configure the basic parameters and then click Next.
 - Name: Enter the application name.
 - Cluster: Select a cluster to which the application is deployed.
 - Namespace: Select a namespace in which the application deployment is located. By default, the default namespace is used.
 - Replicas: Set the number of pods included in the application.
 - Type: Deployment type and StatefulSet type are available.



Note:

In this example, select the StatefulSet type.



The screenshot shows the 'Basic Information' tab of a configuration page. It includes the following fields:

- Name:** A text input field containing 'nginx'. Below it, a note states: 'The name should be 1-64 characters long, and can contain numbers, lower case English letters and hyphens, but cannot start with a hyphen.'
- Cluster:** A dropdown menu showing 'test-mia'.
- Namespace :** A dropdown menu showing 'default'.
- Replicas:** A text input field containing '2'.
- Type:** A dropdown menu showing 'StatefulSet', which is highlighted with a red rectangular border.

At the bottom right, there are 'Back' and 'Next' buttons. A 'Contact Us' link is visible on the right side of the form.

4. Configure containers.



Note:

You can configure multiple containers for the pod of the application.

a) Configure the general settings for the application.

- **Image Name:** Click Select image to select the image in the displayed dialog box and then click OK. In this example, select the nginx image.

You can also enter the private registry in the format of `domainname / namespace / imagename : tag` to specify an image.

- **Image Version:** Click Select image version to select a version. If the image version is not specified, the system uses the latest version by default.
- **Always pull image:** Container Service caches the image to improve deployment efficiency. During deployment, if the image tag is found consistent with that on the local cache, the image on the local cache is reused and is not pulled again. Therefore, if you do not modify the image tag when changing your codes and image for convenience of upper-layer business, the early image on the local cache is used in the application deployment. With this check box selected, Container Service ignores the cached image and re-pulls the image from the repository when deploying the application to make sure the latest image and codes are always used.
- **Resource Limit:** Specify the upper limit for the resources (CPU and memory) that can be used by this application to avoid occupying excessive resources.

CPU is measured in millicores, that is, one thousandth of one core. Memory is measured in bytes, which can be Gi, Mi, or Ki.

- **Resource Request:** Specify how many resources (CPU and memory) are reserved for the application, that is, these resources are exclusive to the container. Other services or processes will compete for resources when the resources are insufficient. By specifying the Resource Request, the application will not become unavailable because of insufficient resources.
- **Init Container:** Selecting this check box creates an Init Container which contains useful tools. For more information, see <https://kubernetes.io/docs/concepts/workloads/pods/init-containers/>.

The screenshot shows the 'Container' configuration page with three tabs: 'Basic Information', 'Container' (selected), and 'Advanced'. Under the 'Container' tab, there's a 'Container1' header with an 'Add Container' button. A 'General' sidebar is on the left. The main area contains the following fields:

- Image Name:** A text box with 'nginx' and a 'Select image' link.
- Image Version:** A text box with 'latest' and a 'Select image version' link.
- Always pull image:** A checkbox that is unchecked, with an 'Image pull secret' link.
- Resource Limit:** Two rows of fields. The first row has 'CPU' (500m), 'Core' (empty), and 'Memory' (128Mi). The second row has 'CPU' (500m), 'Core' (empty), and 'Memory' (128Mi). Units 'MiB' are shown at the end of each row.
- Resource Request:** Similar to Resource Limit, with 'CPU' (500m), 'Core' (empty), and 'Memory' (128Mi) fields.
- Init Container:** A checkbox that is unchecked.

b) Optional: Configure Environment .

You can configure environment variables for the pod by using key-value pairs. Environment variables are used to add environment labels or pass configurations for the pod. For more information, see [Pod variable](#).

c) Optional: Configure Health Check.

The health check function includes liveness probes and readiness probes. Liveness probes are used to detect when to restart the container. Readiness probes determine if the container is ready for receiving traffic. For more

information about health check, see <https://kubernetes.io/docs/tasks/configure-pod-container/configure-liveness-readiness-probes>.

Health Check

Liveness

☒ Enable

HTTP

TCP

Command

▼

Protocol

HTTP ▼

path

Port

Http Header

name

value

Initial Delay

3

Period

10

Timeout

1

Success Threshold

1

Failure Threshold

3

Readiness

☒ Enable

HTTP

TCP

Command

▼

Protocol

HTTP ▼

path

Port

Http Header

name

value

Initial Delay

3

Period

10

Timeout

1

Request method	Description
HTTP request	<p>An HTTP GET request is sent to the container. The following are supported parameters:</p> <ul style="list-style-type: none">• Protocol: HTTP/HTTPS• Path: Path to access the HTTP server• Port: Number or name of the port exposed by the container. The port number must be in the range of 1 to 65535.• HTTP Header: Custom headers in the HTTP request. HTTP allows repeated headers. Supports the correct configuration of key values.• Initial Delay (in seconds): Namely, the initialDelaySeconds. Seconds for the first probe has to wait after the container is started. The default is 3.• Period (in seconds): Namely, the periodseconds. Intervals at which the probe is performed. The default value is 10. The minimum value is 1.• Timeout (in seconds): Namely, the timeoutSeconds. The time of probe timeout. The default value is 1 and the minimum value is 1.• Success Threshold: The minimum number of consecutive successful probes that are considered as successful after a failed probe. The default is 1 and the minimum is 1. It must be 1 for a liveness probe.• Failure Threshold: The minimum number of consecutive failed probes that are considered as failed after a successful probe. The default value is 3. The minimum value is 1.

Request method	Description
TCP connection	<p>A TCP socket is send to the container. The kubelet attempts to open a socket to your container on the specified port. If a connection can be established, the container is considered healthy. If not, it is considered as a failure. The following are supported parameters:</p> <ul style="list-style-type: none">• Port: Number or name of the port exposed by the container. The port number must be in the range of 1 to 65535.• Initial Delay (in seconds): Namely, the initialDelaySeconds. Seconds for the first liveness or readiness probe has to wait after the container is started. The default is 15.• Period (in seconds): Namely, the periodseconds. Intervals at which the probe is performed. The default value is 10. The minimum value is 1.• Timeout (in seconds): Namely, the timeoutSeconds. The time of probe timeout. The default value is 1 and the minimum value is 1.• Success Threshold: The minimum number of consecutive successful probes that are considered as successful after a failed probe. The default is 1 and the minimum is 1. It must be 1 for a liveness probe.• Failure Threshold: The minimum number of consecutive failed probes that are considered as failed after a successful probe. The default value is 3. The minimum value is 1.

Request method	Description
Command line	<p>Detect the health of the container by executing probe detection commands in the container. The following are supported parameters:</p> <ul style="list-style-type: none"> • Command: A probe command used to detect the health of the container • Initial Delay (in seconds): Namely, the <code>initialDelaySeconds</code>. Seconds for the first liveness or readiness probe has to wait after the container is started. The default is 5. • Period (in seconds): Namely, the <code>periodSeconds</code>. Intervals at which the probe is performed. The default value is 10. The minimum value 1. • Timeout (in seconds): Namely, the <code>timeoutSeconds</code>. The time of probe timeout. The default value is 1 and the minimum value is 1. • Success Threshold: The minimum number of consecutive successful probes that are considered as successful after a failed probe. The default is 1 and the minimum is 1. It must be 1 for a liveness probe. • Failure Threshold: The minimum number of consecutive failed probes that are considered as failed after a successful probe. The default value is 3. The minimum value is 1.

d) Optional: Configure the lifecycle rule.

You can configure the following parameters for the container lifecycle: container config start, post start, and pre-stop. For more information, see <https://kubernetes.io/docs/tasks/configure-pod-container/attach-handler-lifecycle-event/>.

- **Container Config:** Select the stdin check box to enable standard input for the container. Select the tty check box to assign an virtual terminal to for the

container to send signals to the container. These two options are usually used together, which indicates to bind the terminal (tty) to the container standard

input (stdin). For example, an interactive program obtains standard input from you and then displays the obtained standard input in the terminal.

- **Start:** Configure a pre-start command and parameter for the container.
- **Post Start:** Configure a post-start command for the container.
- **Pre Stop:** Configure a pre-end command for the container.

Life cycle	Start:	Command	<code>["/bin/sh", "-c", "echo Hello > /user/share/message"]</code>
		Parameter	
	Post Start:	Command	
	Pre Stop:	Command	<code>["/user/sbin/nginx", "-s", "quit"]</code>

e) Configure data volumes.

Local storage and cloud storage can be configured.

- **Local storage:** Supports hostPath, configmap, secret, and temporary directory. The local data volumes mount the corresponding mount source to the container path. For more information, see [Volumes](#).
- **Cloud storage:** Supports three types of cloud storage: cloud disks, Network Attached Storage (NAS), and Object Storage Service (OSS).

In this example, configure a data volume claim named disk-ssd of cloud disk type and mount it to the `/data` path.

Data Volume:		
+ Add local storage		
Storage type	Mount source	Container Path
+ Add cloud storage		
Storage type	Mount source	Container Path
Disk	disk-ssd	/data

f) Optional: Configure Log Service. You can configure collection methods and customize tags for this service.



Note:

Make sure that a Kubernetes cluster is deployed and that the log plug-in is installed on the cluster.

Configure log collection methods as follows:

- Log Store: Configure a Logstore generated in Log Service which is used to store collected logs.
- Log path in the container: Supports stdout and text logs.
 - stdout: Collects standard output logs of containers.
 - text log: Collects logs in the specified path in the container. In this example, collect text logs in the path of /var/log/nginx. Wildcards are also supported.

You can also set custom tags. The customized tags are collected to the container output logs. A custom tag can help you tag container logs, providing convenience to log analysis such as log statistics and filter.

Log Service:

Note: please ensure that cluster has deployed log plug-ins.

Configuration

Log Store

Log path in the container (can be set to stdout)

catalina

stdout

–

access

/var/log/nginx

–

Custom Tag

Name Of Tag

Value Of Tag

app

nginx

–

5. Click Next after completing the configurations.
6. Configure advanced settings. In this example, configure only access settings.
 - a) Set Access Control.

You can set the methods to expose the application pod and then click Create. In this example, a cluster IP service and an Ingress are set to create an Nginx application that is accessible for the Internet.



Note:

You can set access methods according to the communication requirements of your application.

- **Internal application** : an application that works only inside the cluster. You can create a cluster IP service or a node port service as needed for communication within the cluster.

- **External application** : an application that needs to be exposed to the Internet. You can set how the application is accessed by using either of the following two methods:
 - Create a Server Load Balancer service. This method uses Alibaba Cloud Server Load Balancer (SLB) to provide Internet accessibility for the application.
 - Create a cluster IP service or a node port service, and create an Ingress. This method provides Internet accessibility through the Ingress. For more information, see [Ingress](#).

Basic Information		Container	Advanced	Done
Access Control	Service(Service)	Create		
	Ingress(Ingress)	Create		
Scale	HPA	<input type="checkbox"/> Enable		
Scheduling	Node Affinity	Add		
	Pod Affinity	Add		
	Pod Anti Affinity	Add		
			Prev	Create

A. Click **Create** on the right of **Service**. Configure a service in the displayed dialog box, and then click **Create**.

Create Service

Name:

Type:

Server Load Balancer

public

Port Mapping:

+ Add

Name	service port	Container Port	Protocol
nginx-svc	80	80	TCP

Annotation:

+ Add Annotations for load balancer

Tag:

+ Add

Create

Cancel

- **Name:** Enter the service name. The default is `applicationname - svc`.
- **Type:** Select one service type.
 - **Cluster IP:** Exposes the service by using the internal IP address of your cluster. If you select this service type, the service is accessible only within the cluster.
 - **Node port:** Exposes the service by using the IP address and the static port (NodePort) of each node. A node port service routes to a cluster IP service that is automatically created. You can access the node port service from outside the cluster by requesting `< NodeIP > : < NodePort >`.
 - **Server Load Balancer:** Alibaba Cloud Server Load Balancer service. With this type of service, you can set an Internet or intranet access method

for your application. SLB can route to a node port service and a cluster IP service.

- **Port Mapping:** Add a service port and a container port, and select the TCP or UDP protocol. If you select the node port Type, you must add a node port to avoid port conflict.
- **annotation:** Add an annotation to the service. You can set SLB parameters. For more information, see [Access services by using Server Load Balancer](#).
- **Tag:** Add a tag to the service to identify the service.

B. Click Create on the right of Ingress. In the displayed dialog box, configure an Ingress rule for the application pod, and then click Create. For more information, see [Ingress configurations](#).



Note:

When you create an application by using an image, you can create an Ingress rule for only one service. In this example, a virtual host name is used as the

test domain name. You need to add a record to the host. You must use a filing domain name when you create your application.

101 . 37 . 224 . 146 foo . bar . com # This is the
IP address of the Ingress .

Create

Name:

Rule:

+ Add

Domain

foo.bar.com

Select *

ainer.com or Custom

path

e.g./

Services

+ Add

Name

Port

nginx-svc

80

-

EnableTLS

Service weight:

Enable

Grayscale release:

+ Add

After the gray rule is set, the request meeting the rule will be routed to the new service. If you set a weight other than 100, the request to satisfy the gamma rule will continue to be routed to the new and old version services according to the weights.

Annotation:

+ Add

rewrite annotation

Tag:

+ Add

Create

Cancel

C. In the access control area, the created service and Ingress are displayed. You can perform further configurations by clicking Update or Delete.

194

Issue: 20190614

Create Application

Basic Information > Container > **Advanced** > Done

Services(Service) [Update](#) [Delete](#)

service port	Container Port	Protocol
80	80	TCP

Ingresses(Ingress) [Update](#) [Delete](#)

Domain	path	Name	service port
foo.bar.com		nginx-svc	80

Scale

HPA ☐ Enable

Update Method ☒ Enable

RollingUpdate ☐ OnDelete

Max Unavailable: 25 %

Max Surge: 25 %

Scheduling

Node Affinity [Add](#)

Pod Affinity [Add](#)

Pod Anti Affinity [Add](#)

b) Optional: Set Horizontal Pod Autoscaling (HPA).

You enable HPA by selecting the Enable check box. Alibaba Cloud Container Service for Kubernetes provides pod auto scaling to deal with different application workloads. That is, you can change the number of pods according to the container CPU and memory usage.

Scale

HPA ☒ Enable

Metric: CPU Usage

Condition: Usage 70 %

Maximum Replicas: 10 Range : 2-100

Minimum Replicas: 1 Range : 1-100



Note:

To use this function, you must set required resources for the pod. Otherwise, pod auto scaling cannot take effect. For more information, see general container settings.

- **Metric:** resource type. CPU or memory is available. This parameter must be specified with a resource type that is the same as the required resource type.
- **Condition:** the percentage value of resource usage. The number of containers increases when the resource usage exceeds this value.
- **Maximum Replicas:** the maximum number of the containers that the deployment can include.
- **Minimum Replicas:** the minimum number of the containers that the deployment can include.

c) Optional: Set Scheduling.

You can set an update method, node affinity, pod affinity, and pod anti affinity. For more information, see [Affinity and anti-affinity](#).



Note:

Affinity scheduling depends on node tags and pod tags. You can use built-in or customized tags to schedule nodes or pods.

A. Set Update Method.

You can select the `RollingUpdate` or `Recreate` (OnDelete) method to replace old pods with new ones. For more information, see [Deployments](#).

B. Set Node Affinity by using node tags.

Required rules and preferred rules are supported, and available operators include `In` , `NotIn` , `Exists` , `DoesNotExist` , `Gt` , and `Lt` .

- Required rules must be satisfied and correspond to `requiredDuringSchedulingIgnoredDuringExecution` . The required rules

have the same effect as `NodeSelect` or `NodeAffinity`. In this example, the pod can be scheduled to only a node with the specified tags.

You can add multiple required rules, but only one required rule needs to be satisfied for pod scheduling.

- Preferred rules can be unnecessarily satisfied and correspond to `preferredDuringSchedulingIgnoredDuringExecution`. With the scheduling setting in this example, the system tries not to schedule the pod to the nodes with the specified tag.

You can also set `Weight` for each preferred rule. If multiple nodes satisfies the preferred rules, the system schedules the pod to a node with the highest weight.

You can add multiple preferred rules, and all the rules must be satisfied for pod scheduling.

- C. Set Pod Affinity to deploy the application pod in a topology domain together with other pods. For example, to reduce network latency between the services

that communicate with each other, you can deploy their pods to a topology domain (for example, a host).

Create

Required:

+

Add Rule

1

Namespace

default

Topology Key

kubernetes.io/hostname

2

Selector

+

Add

View Application List

Tag Name	Operator	Tag Value	
3	app	In	nginx

Preferred:

+

Add Rule

Weight

100

Namespace

default

Topology Key

kubernetes.io/hostname

Selector

+

Add

View Application List

Tag Name	Operator	Tag Value
app	NotIn	wordpress

OK

Cancel

You can schedule pods according to tags of pods running on nodes. Required rules and preferred rules are supported, and available operators include `In`, `NotIn`, `Exists`, `DoesNotExist`.

- Required rules must be satisfied and correspond to `requiredDuringSchedulingIgnoredDuringExecution`. All specified conditions of required rules must be met for pod affinity scheduling.
- Namespace: Set a namespace. This parameter is required because the scheduling policy is based on pod tags.

Issue: 20190614

199

- **Topology Key:** Set a topology domain to which pods are scheduled. This parameter takes effect through node tags. For example, if you set `kubernetes . io / hostname` as the topology key, a node is used to identify a topology. If you set `beta . kubernetes . io / os` as the topology key, a node operating system is used to identify a topology.
- **Selector:** Click this button to add a required rule.
- **View Application List:** Click View Application List, a dialog box is displayed. In the dialog box, you can view applications in each namespace and export application tags to the dialog box in which you set pod affinity.
- **Required rule tag:** Set a tag name, its operator, and the tag value for existing applications. This example schedules the application to be created to a host on which applications tagged with `app : nginx` run.
- Preferred rules can be unnecessarily satisfied and correspond to `preferredDuringSchedulingIgnoredDuringExecution`. Specified conditions of required rules will be met as many as possible for pod affinity scheduling.

You can set Weight for each preferred rule. The weight value range is 1 to 100. If multiple nodes satisfies the preferred rules, the system schedules the pod to a node with the highest weight. Other parameters are the same with the required rule setting.

D. Set Pod Anti Affinity to deploy the application pods in a topology domain that excludes other pods. Scenarios that use pod anti affinity scheduling include:

- Distribute the pods of a service to different topology domains (for example , different hosts) to improve the service stability.
- Grant a pod the exclusive access to a node so as to guarantee that no other pods use the resources of the node.
- Distribute pods of the services that may affect each other to different hosts.



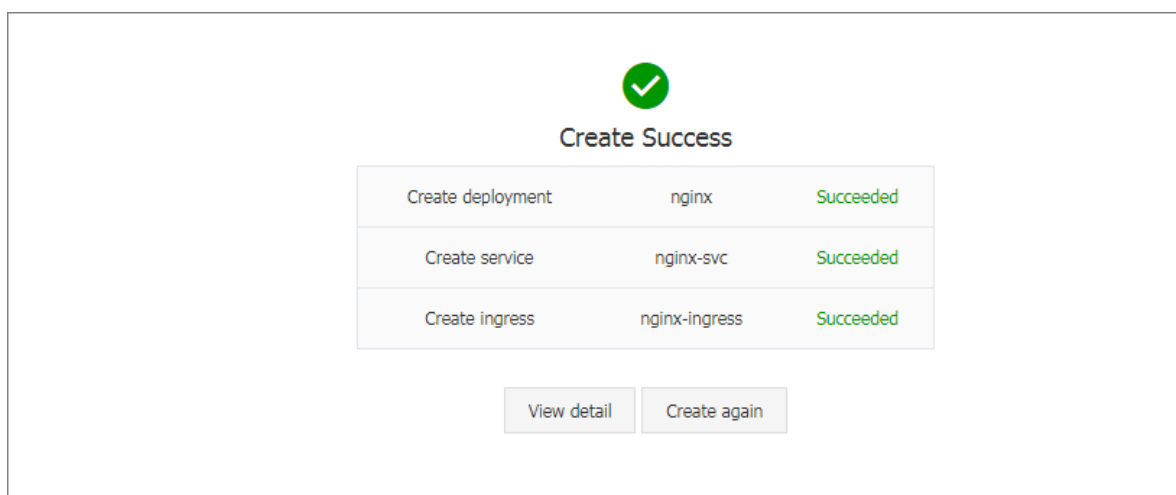
Note:

You can set pod anti affinity scheduling by using the same method as setting pod affinity scheduling. But the same scheduling rules have different

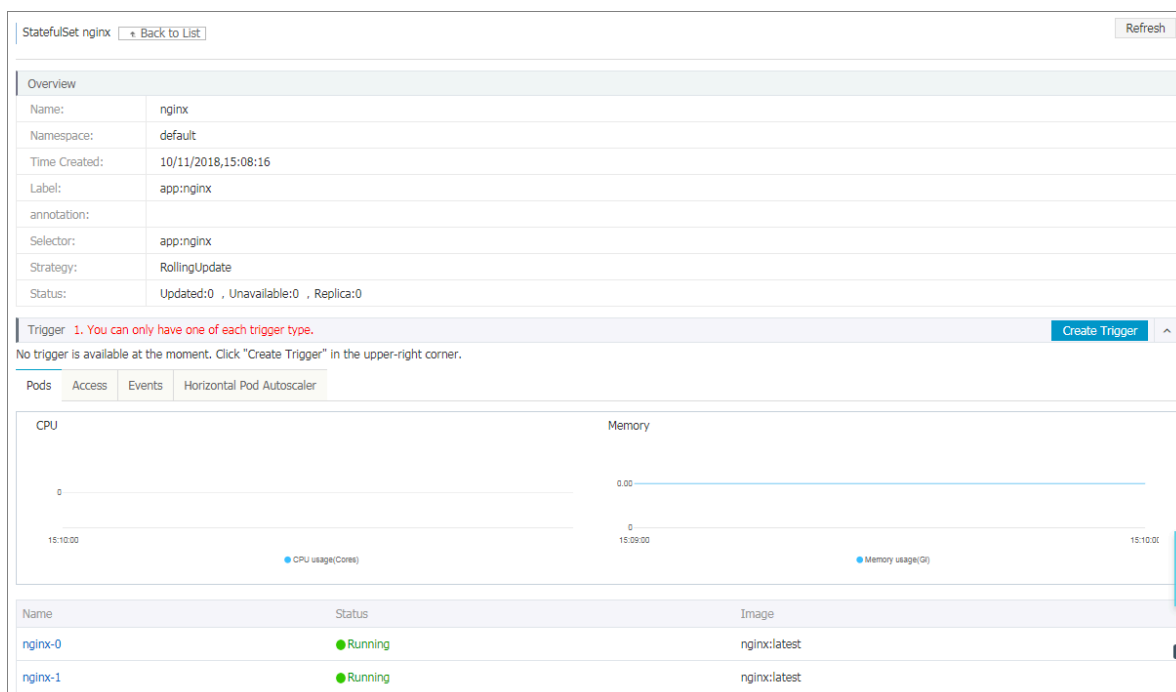
meanings for these two types of scheduling. You need to select appropriate scheduling rules as needed.

7. Click Create.

8. After you create the application, the create success page is displayed by default and objects contained in the application are listed. You can click View detail to view the deployment details.



The StatefulSet page is displayed by default.



9. Then click Back to list in the upper-left corner to view the created StatefulSet application in the StatefulSet list page.

StatefulSet					Refresh	Create by Image	Create by Template
Clusters	test-mia	Namespace	default				
Name	Tag	PodsQuantity	Image	Time Created	Action		
nginx		2/2	nginx	10/11/2018,15:55:57	Details	Edit	Scale More

- 10.Optional: To verify service scalability, click Scale at the right of a target nginx application.

- a) In the displayed dialog box, set the number of pod to 3. You can see that when you expand pods, the pods are in the increment order; when you contract pods, the pods are in the descending order. This shows the order stability of pods in StatefulSet.

Name	Status	Image
nginx-0	Running	nginx:latest
nginx-1	Running	nginx:latest
nginx-2	Running	nginx:latest

- b) Click Application > Volumes Claim in the left-side navigation pane, you can see that as the application expands, new cloud disk volumes are created with pods; if the application contracts, created PV/PVC will not be deleted.

Volumes Claims

RefreshCreate

Clusters

test-mia

Namespace

default

Name	Capacity	Access Mode	Status	Storage Class Name	Relate Volume	Time Created	Action
disk-ssd-nginx-0	20Gi	ReadWriteOnce	Bound	alicloud-disk-ssd	d-bp1cy8o56jfgodpst28f	10/11/2018,15:55:57	Delete
disk-ssd-nginx-1	20Gi	ReadWriteOnce	Bound	alicloud-disk-ssd	d-bp1gdkenf5ki7pt5pct9	10/11/2018,15:56:09	Delete
disk-ssd-nginx-2	20Gi	ReadWriteOnce	Bound	alicloud-disk-ssd	d-bp1f2xopk3sz02ug12ls	10/11/2018,15:57:02	Delete

What's next

Connect to the master node and run following commands to verify the persistent storage feature.

Create a temporary file on a cloud disk:

```
# kubectl exec nginx - 1 ls / tmp # list files
under this directory
lost + found

# kubectl exec nginx - 1 touch / tmp / statefulse t
# add a tempoty file named statefulse t
```

```
# kubectl exec nginx - 1 ls / tmp
lost + found
statefulse t
```

Remove the pod to verify the data persistence:

```
# kubectl delete pod nginx - 1
pod "nginx - 1 " deleted

# kubectl exec nginx - 1 ls / tmp #
data persistenc e storage
lost + found
statefulse t
```

In addition, you can also find that after you delete a pod, the pod automatically restarts after a period of time, which indicates the high availability of the StatefulSet application.

1.8.3 Create a Job application by using an image

By running a Kubernetes cluster with Alibaba Cloud Container Service, you can create a Job application through the Web interface. This example creates a Job application named busybox to describe features of the Job application features.

Prerequisites

You have created a Kubernetes cluster. For more information, see [Create a Kubernetes cluster](#).

Context

A Job processes short-lived one-off tasks in batches to guarantee that one or multiple pods in the batch tasks successfully terminate.

Kubernetes supports the following types of Jobs:

- **Non-parallel Job:** A Job of this type creates only one pod. The Job is completed when the pod terminates successfully.
- **Job with a fixed completion count:** A Job of this type has `. spec . completion s` set to create multiple pods. The Job is completed when the number of these pods reaches the `. spec . completion s` value.
- **Parallel Job with a work queue:** A Job of this type has `. spec . Parallelis m` set but has `. spec . completion s` not set. The Job is completed when at least one pod has terminated with success, and all pods are terminated.

- **Parallel Job with a fixed completion count:** A Job of this type has both `.spec.completion s` and `.spec.Parallelis m` set. Multiple pods of the Job process the work queue at the same time.

According to the `.spec.completion s` and `.spec.Parallelis m` settings, Jobs can be classified into the following patterns.



Note:

The Job created in this example is a parallel Job with a fixed completion count.

Job pattern	Usage example	Action	Completion	Parallelism
One-off Job	Database migration	A Job creates a pod and the Job is completed when the pod terminates successfully.	1	1
Job with a fixed completion count	Pod that processes the work queue	A Job creates pods one by one. When the pods terminate successfully and the number of the terminated pods reaches the completion s value, the Job is completed.	2+	1
Parallel Job with a fixed completion count	Multiple pods process work queues at the same time	A Job creates pods one by one. When the number of pods reaches the completion s value, the Job is completed.	2+	2+

Job pattern	Usage example	Action	Completion	Parallelism
Parallel Job	Multiple pods process work queues at the same time	A Job creates one or multiple pods. When at least one pod terminates successfully, the Job is completed.	1	2+

Procedure

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Kubernetes, choose Application > Job, and then click Create by Image in the upper-right corner.
3. Set the basic parameters and then click Next.
 - Name: Enter a name for the application.
 - Cluster: Select a cluster to which the application is deployed.
 - Namespace: Select a namespace in which the application deployment is located. You can also choose to use the default namespace.
 - Type: Select the Job type.



Note:

In this example, select the Job type.

The screenshot shows the 'Basic Information' step of the Kubernetes Job creation process. The form includes the following fields and values:

- Name:** busybox
- Cluster:** k8s-test
- Namespace:** default
- Type:** Job

A note below the Name field states: "The name should be 1-64 characters long, and can contain numbers, lower case English letters and hyphens, but cannot start with a hyphen." At the bottom right of the form are "Back" and "Next" buttons.

4. Configure containers.



Note:

You can configure multiple containers for the pods of the application.

a) Set the container parameters.

- **Image Name:** Click Select image to select an image in the displayed dialog box and then click OK. In this example, select the busybox image.

You can also enter a private registry in the format of `domainname / namespace / imagename : tag` to specify an image.

- **Image Version:** Click Select image version to select a version. If you do not specify any image version, the system uses the latest version by default.
- **Always pull image:** Container Service caches the image to improve deployment efficiency. During deployment, if the tag of the newly specified image is the same as that of the cached image, Container Service reuses the cached image rather than pulls the same image again. Therefore, if you do not modify the image tag during scenarios where you are changing your code and image, the earlier image in the local cache is used in the application deployment. If you select this check box, Container Service ignores the cached image and re-pulls the image when deploying the application to make sure the latest image and code are always used.
- **Image pull secret:** If you use a private image, we recommend that you use a secret to guarantee the security of your image. For more information, see [Use an image Secret](#).
- **Resource Limit:** Specify the upper limit for the resources (CPU and memory) that can be used by this application to avoid occupying excessive resources. CPU is measured in millicores, that is, one thousandth of one core. Memory is measured in bytes, which can be Gi, Mi, or Ki.
- **Resource Request:** Specify how many resources (CPU and memory) are reserved for the application (that is, these resources become exclusive to the container). If you do not set this parameter, other services or processes will

compete for resources, which means the application may become unavailable due to resource shortage.

- **Init Container:** Select this check box to create an Init Container that contains useful tools. For more information, see <https://kubernetes.io/docs/concepts/workloads/pods/init-containers/>.

The screenshot shows the 'Container1' configuration page. At the top, there is a tab labeled 'Container1' and a button '+ Add Container'. Below this, the 'General' tab is selected. The 'Image Name' field is set to 'busybox' and the 'Image Version' field is set to 'latest'. A red box highlights these two fields. To the right of the 'Image Name' field is a link 'Select image', and to the right of the 'Image Version' field is a link 'Select image version'. Below these fields, there is a checkbox labeled 'Always pull image' and a link 'Image pull secret'. Further down, there are sections for 'Resource Limit' and 'Resource Request'. Each section has input fields for CPU (e.g., 500m), Core, and Memory (e.g., 128Mi, MiB). At the bottom, there is a checkbox labeled 'Init Container'.

b) Optional: Set Environment .

You can use key-value pairs to set environment variables for the pods.

Environment variables are used to add environment labels or pass configurations for the pods. For more information, see [Pod variable](#).

c) Optional: Set Health Check.

You can set liveness probes and readiness probes. Liveness probes are used to detect when to restart the container. Readiness probes determine if the container is ready to receive traffic. For more information about health check,

see <https://kubernetes.io/docs/tasks/configure-pod-container/configure-liveness-readiness-probes>.

Health Check

Liveness

☒ Enable

HTTP

TCP

Command

▼

Protocol

HTTP

▼

path

Port

Http Header

name

value

Initial Delay

3

Period

10

Timeout

1

Success Threshold

1

Failure Threshold

3

Readiness

☒ Enable

HTTP

TCP

Command

▼

Protocol

HTTP

▼

path

Port

Http Header

name

value

Initial Delay

3

Period

10

Timeout

1

Request method	Description
HTTP request	<p>With this health check method, you can send an HTTP GET request to the container. The following parameters are supported:</p> <ul style="list-style-type: none"> • Protocol: HTTP/HTTPS. • Path: path to access the HTTP server. • Port: number or name of the access port exposed by the container. The port number must be in the range of 1 to 65535. • HTTP Header: custom headers in the HTTP request. HTTP allows repeated headers. You can use a key-value pair to set an HTTP Header. • Initial Delay (in seconds): namely, the initialDelaySeconds, indicating the number of seconds for which the first probe must wait after the container is started. The default value is 3. • Period (in seconds): namely, the periodseconds, indicating the interval at which probes are performed. The default value is 10. The minimum value is 1. • Timeout (in seconds): namely, the timeoutSeconds, indicating the number of time that the probe has timed out. The default value is 1 and the minimum value is 1. • Success Threshold: The minimum number of consecutive successful probes needed for determining a probe success after a failed probe . The default value is 1 and the minimum value is 1. It must be set to 1 for a liveness probe. • Failure Threshold: The minimum number of consecutive failed probes needed for determining a probe failure after a successful probe. The default value is 3. The minimum value is 1.

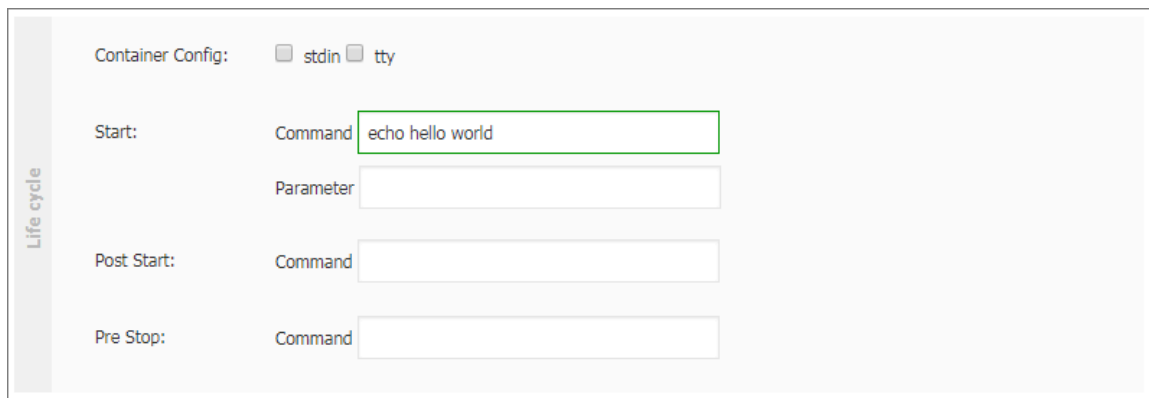
Request method	Description
TCP connection	<p>If you use this health check method, a TCP socket is sent to the container. The kubelet then attempts to open the socket of the container on a specified port. If a connection can be established, the container is considered healthy. If not, it is considered unhealthy. The following parameters are supported:</p> <ul style="list-style-type: none">• Port: number or name of the access port exposed by the container. The port number must be in the range of 1 to 65535.• Initial Delay (in seconds): namely, the <code>initialDelaySeconds</code>, indicating the seconds for the first liveness or readiness probe must wait for after the container is started. The default is 15.• Period (in seconds): namely, the <code>periodSeconds</code>, indicating the interval at which probes are performed. The default value is 10. The minimum value is 1.• Timeout (in seconds): namely, the <code>timeoutSeconds</code>, indicating the number of time that the probe has timed out. The default value is 1 and the minimum value is 1.• Success Threshold: The minimum number of consecutive successful probes needed for determining a probe success after a failed probe. The default value is 1 and the minimum value is 1. It must be set to 1 for a liveness probe.• Failure Threshold: The minimum number of consecutive failed probes needed for determining a probe failure after a successful probe. The default value is 3. The minimum value is 1.

Request method	Description
Command line	<p>With this health check method, you can detect the container health by executing a probe detection command in the container. The following parameters are supported:</p> <ul style="list-style-type: none">· Command: a probe command used to detect the health of the container· Initial Delay (in seconds): namely, the <code>initialDelaySeconds</code>, indicating the number of seconds for which the first liveness or readiness probe must wait after the container is started. The default value is 5.· Period (in seconds): namely, the <code>periodSeconds</code>, indicating the interval at which probes are performed. The default value is 10. The minimum value 1.· Timeout (in seconds): namely, the <code>timeoutSeconds</code>, indicating the number of time that the probe has timed out. The default value is 1 and the minimum value is 1.· Success Threshold: The minimum number of consecutive successful probes needed for determining a probe success after a failed probe. The default value is 1 and the minimum value is 1. It must be set to 1 for a liveness probe.· Failure Threshold: The minimum number of consecutive failed probes needed for determining a probe failure after a successful probe. The default value is 3. The minimum value is 1.

d) Optional: Set the life cycle.

You can set the following parameters for the container life cycle: container config, start, post start, and pre-stop. For more information, see <https://kubernetes.io/docs/tasks/configure-pod-container/attach-handler-lifecycle-event/>.

- **Container Config:** You can select the stdin check box to enable standard input for the container, or select the tty check box to assign a virtual terminal to the container to send signals to the container. You can also select the two options at the same time. That is, you can bind the terminal (tty) to the container standard input (stdin). For example, an interactive program can obtain standard input from you and then display the obtained standard input in the terminal.
- **Start:** Set a pre-start command and parameter for the container.
- **Post Start:** Set a post-start command for the container.
- **Pre Stop:** Set a pre-stop command for the container.



Container Config:	<input type="checkbox"/> stdin <input type="checkbox"/> tty
Life cycle	Start: Command <input type="text" value="echo hello world"/>
	Parameter <input type="text"/>
	Post Start: Command <input type="text"/>
	Pre Stop: Command <input type="text"/>

e) Optional: Set data volumes.

You can configure local storage and cloud storage.

- **Local storage:** Supported storage types include HostPath, ConfigMap, Secret, and EmptyDir. By setting a type of local storage, you can mount its mount source to the container path. For more information, see [Volumes](#).
- **Cloud storage:** Supported types of cloud storage include cloud disks, Network Attached Storage (NAS), and Object Storage Service (OSS).

f) Optional: Set Log Service. You can set collection parameters and customize tags.



Note:

Make sure that you have deployed a Kubernetes cluster and installed the log plugin on the cluster.

Set the following log collection parameters:

- **Log Store:** Set a Logstore. After you specify the Logstore name, the Logstore is generated in Log Service to store collected logs.
- **Log path in the container:** You can set this parameter to stdout or set a log path.
 - **stdout:** If you set a log path to stdout, you can collect the standard output logs of the container.
 - **text log:** If you set a container log path, you can collect the text logs of the path. Wildcards can be used in setting the log file name for a log path.

You can also set custom tags. The custom tags are collected to the container output logs. A custom tag can help you tag container logs, making it easy to collect log statistics, filter logs, and analyze logs by using other methods.

5. After you complete the container configuration, click Next.

6. Configure advanced settings.

You can configure Job Settings.

Parameter	Description
Completions	Number of pods that must be run successfully by the configured Job. The default value is 1.
Parallelism	Number of pods that must be run in parallel by the configured Job at any time. The default value is 1.
ActiveDeadlineSeconds	Operating time limit of the configured Job. If the Job is not completed within the time limit, the system tries to terminate the Job.

Parameter	Description
BackoffLimit	Number of retries performed by the configured Job to create pods after a failure. The default is 6. Each time the Job fails, the failed pods associated with the Job are recreated with time delay . The time delay grows exponentially each time. The upper limit of the time delay is six minutes.
Restart	Only Never and OnFailure restart policies are supported.

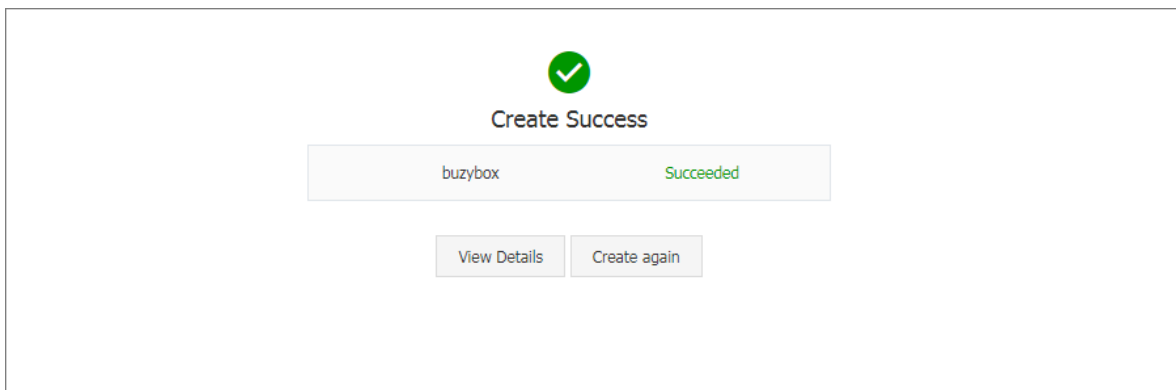
The screenshot shows the 'Advanced' configuration tab for a Kubernetes Job. The 'Job Settings' sidebar on the left contains the following fields:

- Completions: 6
- Parallelism: 2
- ActiveDeadlineSeconds: 600
- BackoffLimit: 6
- Restart: never (dropdown menu)

The main configuration area is currently empty. At the bottom right, there are 'Prev' and 'Create' buttons, and a 'Contact Us' link.

7. Click Create.

8. After you create the Job application, a new page is displayed by default to prompt that you have created the application with the objects included.



You can click View Details to view the Job details.

During the creation process, you can view the creation status of the pods in the Status column. In this example, two pods are created in parallel according to the Job definition.

Job busybox [← Back to List](#) [Refresh](#)

Overview

Name:	busybox
Namespace:	default
Time Created:	11/21/2018,15:56:54
Label:	job-name:busybox controller-uid:0340c7b1-ed63-11e8-89e5-00163e15887a
annotation:	
Status:	Active2 , Succeeded3 Failed0

Pods

Events

Name	Status	Image	Action
busybox-c2k6s	● Succeeded	busybox:latest	Details More
busybox-f5pfs	● Pending	busybox:latest	Details More
busybox-gnjkg	● Pending	busybox:latest	Details More
busybox-pgq7f	● Succeeded	busybox:latest	Details More
busybox-phnvc	● Succeeded	busybox:latest	Details More

Wait until all pods are created.

Job busybox
[← Back to List](#)
Refresh

Overview

Name:	busybox
Namespace:	default
Time Created:	11/21/2018,15:56:54
Label:	job-name:busybox controller-uid:0340c7b1-ed63-11e8-89e5-00163e15887a
annotation:	
Status:	Active0, Succeeded6, Failed0

Pods
Events

Name	Status	Image	Action
busybox-c2k6s	● Succeeded	busybox:latest	Details More ▼
busybox-f5pfs	● Succeeded	busybox:latest	Details More ▼
busybox-gnjkg	● Succeeded	busybox:latest	Details More ▼
busybox-pgq7f	● Succeeded	busybox:latest	Details More ▼
busybox-phnvc	● Succeeded	busybox:latest	Details More ▼
busybox-wdrfj	● Succeeded	busybox:latest	Details More ▼

9. In the upper-left corner, click Back to List. On the Jog page, the Job completion time is displayed.



Note:

If the Job has not created all the pods, the page does not display the Job completion time.

Job
Refresh
Create by Image
Create by Template

Help: [How to use private images](#) [Create applications](#) [Schedule a pod to the specified node](#) [Create a Layer-4 Ingress](#) [Create a Layer-7 Ingress](#) [Configure pod auto scaling](#)

Container monitoring [Blue-green release](#)

Clusters Namespace default

Name	Tag	Status	Pod Status	Image	Time Created	Completion Time	Action
busybox	job-name:busybox controller-uid:0340c7b1-ed63-11e8-89e5-00163e15887a	Succeeded	Active0 Succeeded6 Failed0	busybox:latest	11/21/2018,15:56:54	11/21/2018,15:57:15	Details More ▼

1.8.4 Create an application in Kubernetes dashboard

You can create an application in the Kubernetes dashboard.

Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Clusters in the left-side navigation pane.

3. Click Dashboard at the right of the cluster to enter the Kubernetes dashboard.
4. In the Kubernetes dashboard, click CREATE in the upper-right corner to create an application.
5. The Resource creation page appears. Configure the application information.

Create an application in any of the following three ways:

- **CREATE FROM TEXT INPUT:** Directly enter the orchestration codes in the YAML or JSON format to create an application. You must know the corresponding orchestration format.
- **CREATE AN APP:** Complete the following configurations to create an application.
 - **App name:** Enter the name of the application you are about to create. In this example, enter `nginx - test`.
 - **Container image:** Enter the URL of the image to be used. In this example, use Docker [Nginx](#).
 - **Number of pods:** Configure the number of pods for this application.
 - **Service:** Select External or Internal. External indicates to create a service that can be accessed from outside the cluster. Internal indicates to create a service that can be accessed from within the cluster.
 - **Advanced options:** To configure the information such as labels and environment variables, click SHOW ADVANCED OPTIONS. This configuration distributes the traffic load evenly to three pods.
- **CREATE FROM FILE:** Upload an existing YAML or JSON configuration file to create an application.

6. Click UPLOAD or DEPLOY to deploy the containers and services.

You can also click SHOW ADVANCED OPTIONS to configure more parameters.

What's next

After clicking UPLOAD or DEPLOY, you can view the services and containers of the application.

Click Pods in the left-side navigation pane. You can check the status of each Kubernetes object according to the icon on the left. indicates the object is still being deployed. indicates the object has completed the deployment.

1.8.5 Create a Linux application by using an orchestration template

In a Container Service Kubernetes orchestration template, you must define resource objects required for running an application, and combine the resource objects into a complete application by using label selector.

Prerequisites

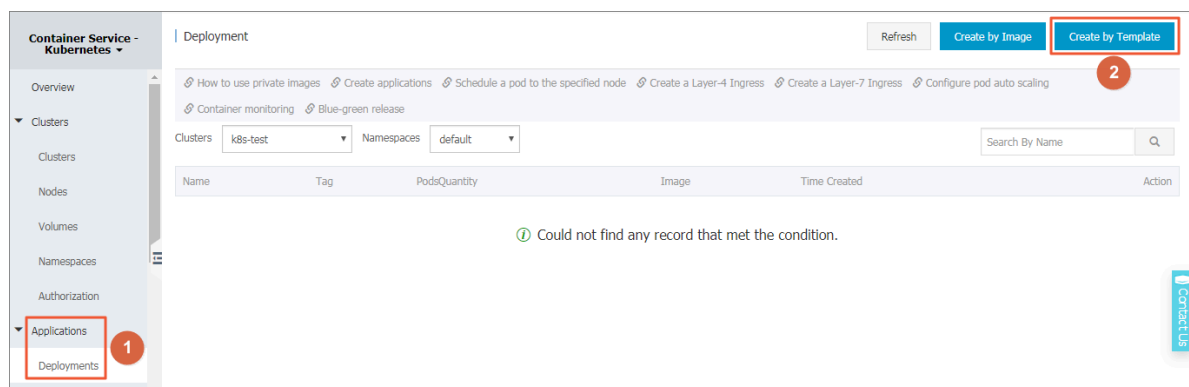
Create a Kubernetes cluster. For more information, see [Create a Kubernetes cluster](#).

Context

Create an Nginx application in this example. Firstly, create a backend pod resource object by creating the deployment. Then, deploy the service to bind it to the backend pod, forming a complete Nginx application.

Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Application > Deployment in the left-side navigation pane.
3. Click Create by Template in the upper-right corner.



4. Configure the template and then click DEPLOY.

- Clusters: Select the cluster in which the resource object is to be deployed.
- Namespace: Select a namespace to which resource object belongs. The default namespace is default. Except for the underlying computing resources such as

nodes and persistent storage volumes, most of the resource objects must act on a namespace.

- **Resource Type:** Alibaba Cloud Container Service provides Kubernetes YAML sample templates of many resource types for you to deploy resource objects quickly. You can write your own template based on the format requirements of Kubernetes YAML orchestration to describe the resource type you want to define.
- **Add Deployment:** You can quickly define a YAML template with this feature.
- **Deploy with exist template:** You can import an existing template into the template configuration page.

Clusters

k8s-test

Namespace

default

Resource Type

Custom

Template

```

1  apiVersion: apps/v1beta2 # for versions before 1.8.0 use apps/v1beta1
2  kind: Deployment
3  metadata:
4    name: nginx-deployment
5    labels:
6      app: nginx
7  spec:
8    replicas: 2
9    selector:
10     matchLabels:
11       app: nginx
12   template:
13     metadata:
14       labels:
15         app: nginx
16     spec:
17       containers:
18       - name: nginx
19         image: nginx:1.7.9 # replace it with your exactly <image_name:tags>
20         ports:
21         - containerPort: 80
22   ---
23
24  apiVersion: v1 # for versions before 1.8.0 use apps/v1beta1
25  kind: Service
26  metadata:
27    name: my-service1 #T000: to specify your service name
28    labels:
29      app: nginx
30  spec:
31    selector:

```

Add Deployment

Deploy with exist template

Deployed successfully. Go to Dashboard to see the deployment progress: [Kubernetes Dashboard](#)

Save Template

DEPLOY

The following is a sample orchestration for an Nginx application. The orchestration is based on an orchestration template built in Container Service. By using this orchestration template, you can create a deployment that belongs to an Nginx application quickly.



Note:

Container Service supports Kubernetes YAML orchestration in which you can use the `---` symbol to separate resource objects so as to create multiple resource objects through a single template.

```

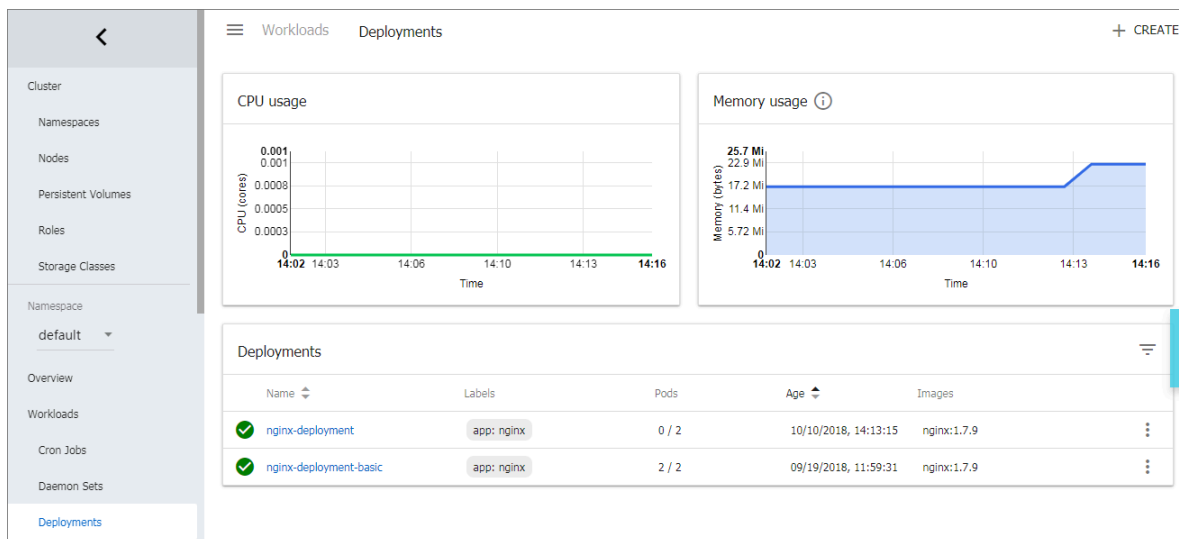
apiVersion : apps / v1beta2 # for versions before 1 . 8 .
0 use apps / v1beta1
kind : Deployment
metadata :
  name : nginx - deployment
  labels :
    app : nginx
spec :
  replicas : 2
  selector :
    matchLabels :
      app : nginx
  template :
    metadata :
      labels :
        app : nginx
    spec :
      containers :
        - name : nginx
          image : nginx : 1 . 7 . 9 # replace it with your
exactly < image_name : tags >
          ports :
            - containerPort : 80
---

apiVersion : v1 # for versions before 1 . 8 . 0 use
apps / v1beta1
kind : Service
metadata :
  name : my - service1 # TODO : to specify your
service name
  labels :
    app : nginx
spec :
  selector :
    app : nginx # TODO : change label selector
to match your backend pod
  ports :
    - protocol : TCP
      name : http
      port : 30080 # TODO : choose an unique
port on each node to avoid port conflict
      targetPort : 80

```

```
type : LoadBalancer ## In this example , change
the type from NodePort to LoadBalancer .
```

5. After you click DEPLOY, a message indicating the deployment status is displayed. After the deployment succeeds, click Kubernetes Dashboard in the message to go to the dashboard and check the deployment progress.



6. In the Kubernetes dashboard, you can see that the service named my-service1 is successfully deployed and its external endpoint is exposed. Click the access address under External endpoints.

Services					
Name	Labels	Cluster IP	Internal endpoints	External endpoints	Age
my-service1	app: nginx	172.19.0.114	my-service1:30080 TCP my-service1:30327 TCP	172.19.0.1:30080	10/10/2018, 14:13:15

7. You can access the Nginx service welcome page in the browser.



What's next

You can also go back to the home page of Container Services and then click Discovery and Load Balancing > Services in the left-side navigation pane to view the Nginx service.

1.8.6 Create a Windows application by using an orchestration template

This topic describes how to create a Windows application by using an orchestration template. Such a template is used to customize the resources required by a Windows application to operate.

Prerequisites

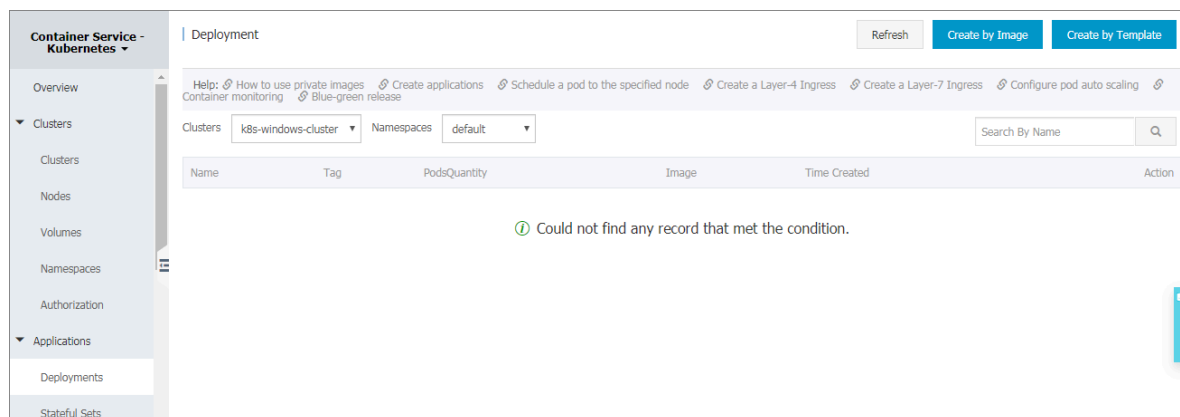
A Kubernetes cluster that supports Windows is created. For more information, see [Create a Windows application by using an orchestration template](#).

Context

In this topic, an application named `aspnet` is created by using an orchestration template. This application contains a deployment and a service. On the backend, the deployment creates pods according to settings. Then, the service is associated with the pods.

Procedure

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Container Service-Kubernetes, choose **Application > Deployment**.
3. In the upper-right corner, click **Create by Template**.



4. To set the orchestration template, set the following:

- **Clusters:** Select the target cluster. The resources required by the application are deployed in a cluster.
- **Namespace:** Select the target namespace. The default namespace is preset. Except for nodes, persistent volumes, and other underlying resource types, most resources required by the application are deployed in a namespace.
- **Sample Template:** Select the target sample template. Alibaba Cloud Container Service for Kubernetes provides many built-in YAML orchestration templates for

different types of resources. You can customize an orchestration template to set a type of resource according to the YAML orchestration requirements.

- **Add Deployment:** Edit a YAML template quickly by using this function.
- **Using Existing Template:** Import an existing template to the template setting area.

Then, click **DEPLOY**.

Namespaces: default

Sample Template: Custom

Template

```

1  apiVersion: v1
2  kind: Service
3  metadata:
4    name: aspnet-svc
5  spec:
6    ports:
7      - port: 80
8        protocol: TCP
9        targetPort: 80
10   selector:
11     app: aspnet
12   type: LoadBalancer
13   ---
14  apiVersion: apps/v1beta2
15  kind: Deployment
16  metadata:
17    name: aspnet
18  spec:
19    selector:
20      matchLabels:
21        app: aspnet
22    template:
23      metadata:
24        labels:
25          app: aspnet
26      spec:
27        containers:
28          - image: 'microsoft/dotnet-samples:aspnetapp'
29            name: aspnet

```

Deployed successfully. Go to Dashboard to see the deployment progress: [Kubernetes Dashboard](#)

Save Template **DEPLOY**

The following is an orchestration template for the Windows application named `aspnet`. With such an orchestration template, you can quickly create a deployment and a service for the application.



Note:

If you want to create multiple resources in a template, you can use `---` to separate different resources.

```

apiVersion : v1
kind : Service
metadata :
  name : aspnet - svc
spec :
  ports :
    - port : 80

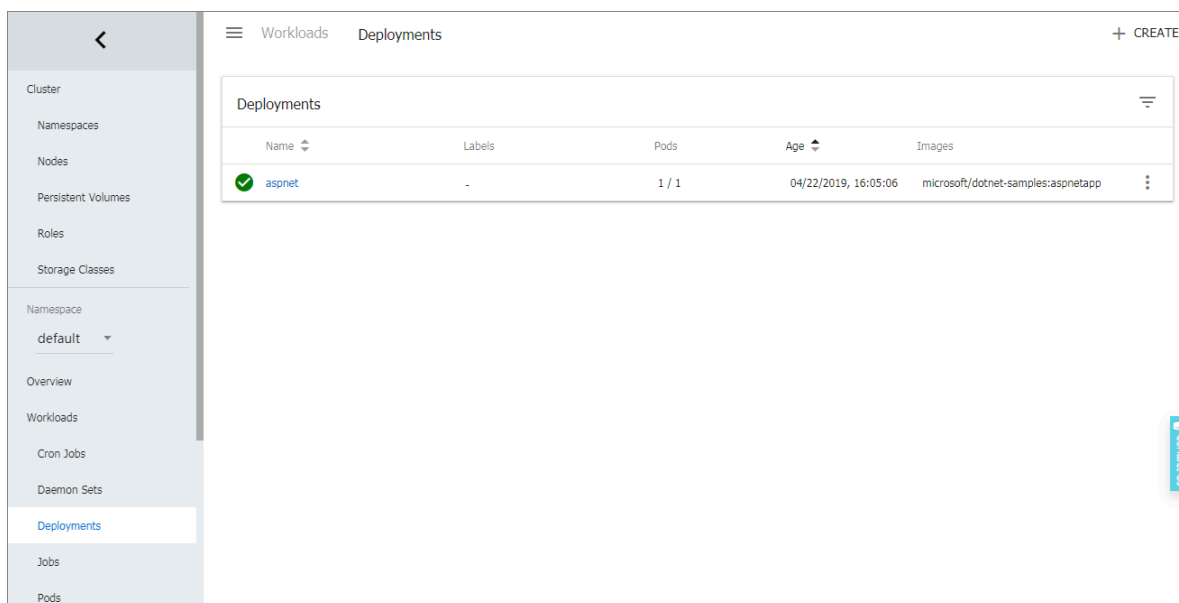
```

```

      protocol : TCP
      targetPort : 80
    selector :
      app : aspnet
    type : LoadBalancer
  ---
  apiVersion : apps / v1beta2
  kind : Deployment
  metadata :
    name : aspnet
  spec :
    selector :
      matchLabels :
        app : aspnet
    template :
      metadata :
        labels :
          app : aspnet
      spec :
        containers :
          - image : 'microsoft / dotnet - samples : aspnetapp '
            name : aspnet

```

5. A message is displayed on the bottom of the Template area to show the result. If a success message is displayed (shown in the preceding figure), click **Kubernetes Dashboard** at the end of the message to view the progress.

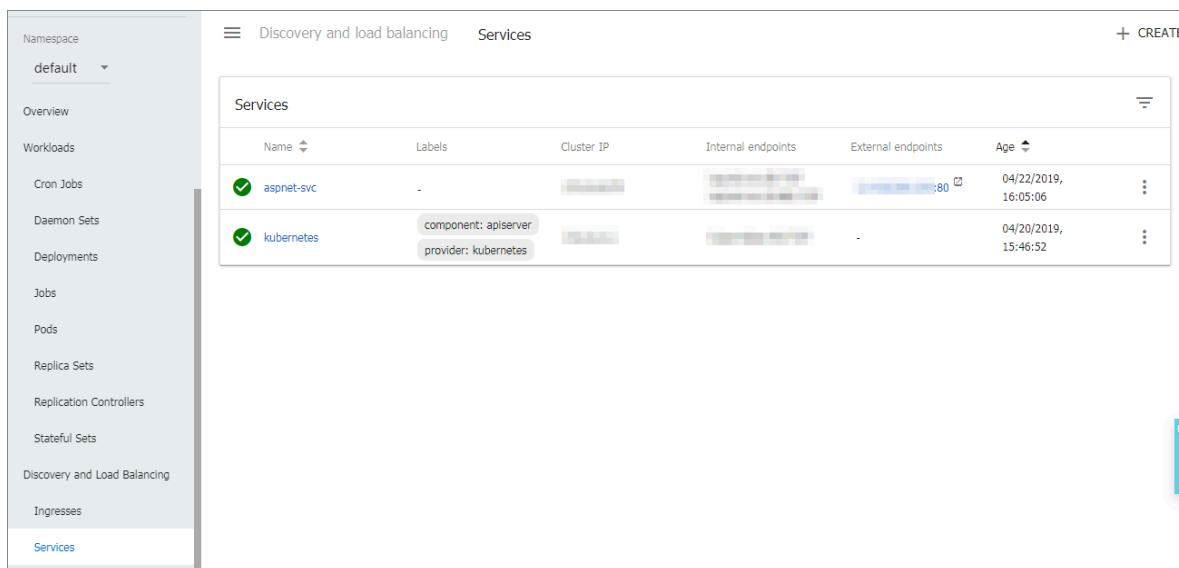


6. In the left-side navigation pane, choose **Discovery and Load Balancing > Service**. Then, in the **External endpoints** column, click the IP address of the created service named `aspnet - svc` to visit the home page of the `aspnet` application in your browser.

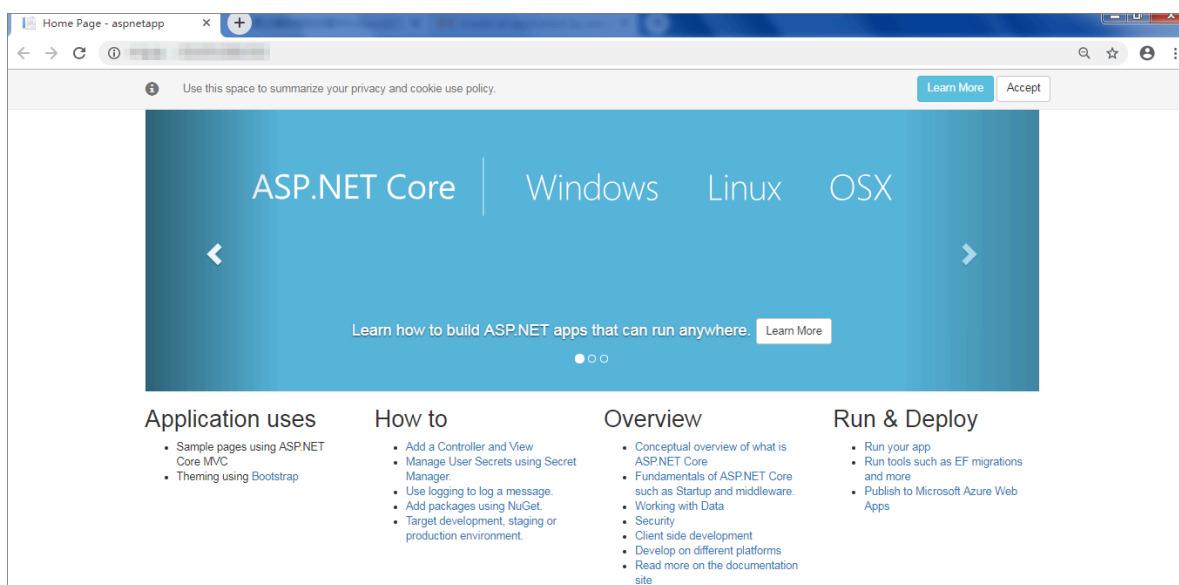


Note:

In the Kubernetes dashboard, you can view that a service named `aspnet - svc` and its external endpoint are created.



The following figure shows the home page of the `aspnet` application.



What's next

You can also return to the home page of Container Service-Kubernetes, and then choose **Discovery and Load Balancing > Service** in the left-side navigation pane to view the service of the `aspnet` application.

1.8.7 Manage applications by using commands

You can create applications or view containers in applications by using commands.

Prerequisites

Before using commands to manage applications, [#unique_57](#).

Create an application by using commands

Run the following statements to run a simple container (a Nginx Web server in this example).

```
root @ master # kubectl run -it nginx -- image = registry .
aliyuncs . com / spacexnice / netdia : latest
```

This command creates a service portal for this container. Specify `-- type =`

`LoadBalanc er` and an Alibaba Cloud Server Load Balancer route will be created to the Nginx container.

```
root @ master # kubectl expose deployment nginx -- port = 80
-- target - port = 80 -- type = LoadBalanc er
```

View containers by using commands

Run the following command to list all the running containers in the default namespaces.

```
root @ master # kubectl get pods
NAME      READY   STATUS    RESTARTS   AGE
nginx - 2721357637 - dvwq3    1 / 1     Running    1      9h
```

1.8.8 Simplify Kubernetes application deployment by using Helm

In Kubernetes, app management is the most challenging and in demand field.

The Helm project provides a uniform software packaging method which supports version control and greatly simplifies Kubernetes app distribution and deployment complexity.

Alibaba Cloud Container Service integrates the app catalog management function with the Helm tool, extends the functions, and supports official repository, allowing you to deploy the application quickly. You can deploy the application in the Container Service console or by using command lines.

This document introduces the basic concepts and usage of Helm and demonstrates how to use Helm to deploy the sample applications WordPress and Spark on an Alibaba Cloud Kubernetes cluster.

Basic concepts of Helm

Helm is an open-source tool initiated by Deis and helps to simplify the deployment and management of Kubernetes applications.

You can understand Helm as a Kubernetes package management tool that facilitates discovery, sharing and use of apps built for Kubernetes. It involves several basic concepts.

- **Chart:** A Helm package containing the images, dependencies, and resource definitions required for running an application. It may also contain service definitions in a Kubernetes cluster, similar to the formula of Homebrew, the dpkg of APT, or the rpm file of Yum.
- **Release:** A chart running on a Kubernetes cluster. A chart can be installed multiple times on the same cluster. A new release will be created every time a chart is installed. For example, to run two databases on the server, you can install the MySQL chart twice. Each installation will generate its own release with its own release name.
- **Repository:** The repository for publishing and storing charts.

Helm components

Helm adopts a client/server architecture composed of the following components:

- **Helm CLI** is the Helm client and can be run locally or on the master nodes of the Kubernetes cluster.
- **Tiller** is the server component and runs on the Kubernetes cluster. It manages the lifecycles of Kubernetes applications.
- **Repository** is the chart repository. The Helm client accesses the chart index files and packages in the repository by means of the HTTP protocol.

Use Helm to deploy applications

Prerequisites

- Before using Helm to deploy an application, create a Kubernetes cluster in Alibaba Cloud Container Service. For more information, see [Create a Kubernetes cluster](#).

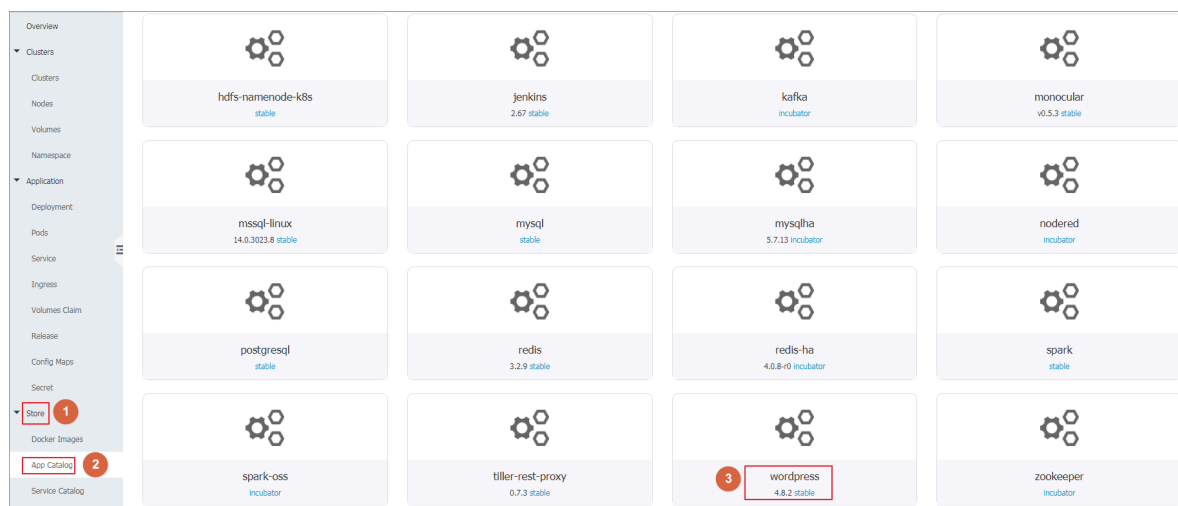
Tiller is automatically deployed to the cluster when the Kubernetes cluster is created. Helm CLI is automatically installed on all the master nodes and the configuration points to the Alibaba Cloud chart repository.

- Check the Kubernetes version of your cluster.

Only clusters whose Kubernetes version is 1.8.4 or later are supported. For clusters whose Kubernetes version is 1.8.1, upgrade the cluster on the Cluster List page.

Deploy applications in Container Service console

- Log on to the [Container Service console](#).
- Under Kubernetes, click Store > App Catalog in the left-side navigation pane.
- On the App Catalog page, click a chart (WordPress in this example) to enter the chart details page.



4. Enter the basic information for the deployment on the right.

- **Clusters:** Select the cluster in which the application is to be deployed.
- **Namespace:** Select the namespace. default is selected by default.
- **Release Name:** Enter the release name for the application. Enter test in this example.

Readme
Values

WordPress

WordPress is one of the most versatile open source content management systems on the market. A publishing platform for building blogs and websites.

TL;DR;

```
$ helm install stable/wordpress
```

Introduction

This chart bootstraps a [WordPress](#) deployment on a [Kubernetes](#) cluster using the [Helm](#) package manager.

It also packages the [Bitnami MariaDB](#) chart which is required for bootstrapping a MariaDB deployment for the database requirements of the WordPress application.

Deploy

Only Kubernetes versions 1.8.4 and above are supported. For clusters of version 1.8.1, you can perform "upgrade cluster" operation in the cluster list

Clusters
k8s-cluster

Namespace
default

Release Name
wordpress-default

DEPLOY

5. Click the Values tab to modify the configurations.

In this example, bind dynamic data volumes of the cloud disk to a persistent storage volume claim (PVC). For more information, see .



Note:

You need to create a persistent storage volume (PV) of cloud disk in advance. The capacity of the PV cannot be less than the value defined by the PVC.

Readme
Values

```

72  ##
73  mariadbDatabase: bitnami_wordpress
74
75  ## Create a database user
76  ## ref: https://github.com/bitnami/bitnami-docker-mariadb/blob/master/README.md#creating-a
77  ## -database-user-on-first-run
78  mariadbUser: bn_wordpress
79
80  ## Password for mariadbUser
81  ## ref: https://github.com/bitnami/bitnami-docker-mariadb/blob/master/README.md#creating-a
82  ## -database-user-on-first-run
83  ## mariadbPassword:
84
85  ## Enable persistence using Persistent Volume Claims
86  ## ref: http://kubernetes.io/docs/user-guide/persistent-volumes/
87  ##
88  persistence:
89    enabled: true
90    ## mariadb data Persistent Volume Storage Class
91    ## If defined, storageClassName: <storageClass>
92    ## If set to "-", storageClassName: "", which disables dynamic provisioning
93    ## If undefined (the default) or set to null, no storageClassName spec is
94    ## set, choosing the default provisioner. (gp2 on AWS, standard on
95    ## GKE, AWS & OpenStack)
96    ##
97    storageClass: "alicloud-disk-efficiency"
98    accessMode: ReadWriteOnce
99    size: 20Gi
100
101  ## Kubernetes configuration
102  ## For minikube, set this to NodePort, elsewhere use LoadBalancer
103  ##
104  serviceType: LoadBalancer
105

```

Deploy

Only Kubernetes versions 1.8.4 and above are supported. For clusters of version 1.8.1, you can perform "upgrade cluster" operation in the cluster list

Clusters
k8s-cluster

Namespace
default

Release Name
wordpress-default

DEPLOY

Version

0.6.13

Project Homepage

6. Click **DEPLOY** after completing the configurations. After the successful deployment, you are redirected to the release page of this application.

Resource	Kind	Values
wordpress-default-mariadb	Secret	View YAML
wordpress-default-wordpress	Secret	View YAML
wordpress-default-mariadb	ConfigMap	View YAML
wordpress-default-mariadb	PersistentVolumeClaim	View YAML
wordpress-default-wordpress	PersistentVolumeClaim	View YAML
wordpress-default-mariadb	Service	View YAML
wordpress-default-wordpress	Service	View YAML
wordpress-default-mariadb	Deployment	View YAML
wordpress-default-wordpress	Deployment	View YAML

7. Click **Application > Service** in the left-hand navigation pane. Select the target cluster and namespace and find the corresponding service. You can obtain the **HTTP/HTTPS** external endpoint address.

Name	Type	Time Created	ClustersIP	internalendpoint	externalendpoint	Action
kubernetes	ClusterIP	06/27/2018,17:53:49		kubernetes:443 TCP	-	Details Update View YAML Delete
wordpress-default-mariadb	ClusterIP	07/10/2018,17:36:16		wordpress-default-mariadb:3306 TCP	-	Details Update View YAML Delete
wordpress-default-wordpress	LoadBalancer	07/10/2018,17:36:16		wordpress-default-wordpress:80 TCP wordpress-default-wordpress:32109 TCP wordpress-default-wordpress:443 TCP wordpress-default-wordpress:32094 TCP	kubernetes:443	Details Update View YAML Delete

8. Click the preceding access address to enter the WordPress blog publishing page.

Deploy applications by using command lines

You can use SSH to log on to the master node of the Kubernetes cluster when deploying applications by using command lines (Helm CLI is automatically installed and has configured the repository). For more information, see [Access Kubernetes clusters by using SSH](#). You can also install and configure the kubectl and Helm CLI locally.

In this example, install and configure the kubectl and Helm CLI locally and deploy the applications WordPress and Spark.

Install and configure kubectl and Helm CLI

1. Install and configure kubectl on a local computer.

For more information, see [Connect to a Kubernetes cluster by using kubectl](#).

To view information of the target Kubernetes cluster, enter the command `kubectl cluster - info`.

2. Install Helm on a local computer.

For the installation method, see [Install Helm](#).

3. Configure the Helm repository. Here the charts repository provided by Alibaba Cloud Container Service is used.

```
helm init --client-only --stable-repo-url https://aliacs-app-catalog.oss-cn-hangzhou.aliyuncs.com/charts/  
helm repo add incubator https://aliacs-app-catalog.oss-cn-hangzhou.aliyuncs.com/charts-incubator/  
helm repo update
```

Basic operations of Helm

- To view the list of charts installed on the cluster, enter the following command:

```
helm list
```

Or you can use the abbreviated version:

```
helm ls
```

- To view the repository configurations, enter the following command:

```
helm repo list
```

- To view or search for the Helm charts in the repository, enter one of the following commands:

```
helm search  
helm search repository name # For example, stable or incubator.  
helm search chart name # For example, wordpress or spark.
```

- To update the chart list to get the latest version, enter the following command:

```
helm repo update
```

For more information about how to use Helm, see [Helm document](#).

Deploy WordPress by using Helm

Use Helm to deploy a WordPress blog website.

Enter the following command.

```
helm install --name wordpress --test stable/wordpress
```



Note:

The Alibaba Cloud Kubernetes service provides the support for dynamic storage volumes of block storage (cloud disk). You need to create a storage volume of cloud disk in advance.

The result is as follows:

```
NAME : wordpress --test
LAST DEPLOYED : Mon Nov 20 19 : 01 : 55 2017
NAMESPACE : default
STATUS : DEPLOYED
...
```

Use the following command to view the release and service of WordPress.

```
helm list
kubectl get svc
```

Use the following command to view the WordPress related pods and wait until the status changes to Running.

```
kubectl get pod
```

Use the following command to obtain the WordPress access address:

```
echo http://$(kubectl get svc wordpress --test -o jsonpath='{.status.loadBalancer.ingress[0].ip}')
```

Access the preceding URL in the browser, and you can see the familiar WordPress website.

You can also follow the chart instructions and use the following command to obtain the administrator account and password of the WordPress website:

```
echo Username : user
```

```
echo Password : $( kubectl get secret -- namespace default  
wordpress - test - wordpress - o jsonpath="{. data . wordpress -  
password }" | base64 -- decode )
```

To completely delete the WordPress application, enter the following command:

```
helm delete -- purge wordpress - test
```

Deploy Spark by using Helm

Use Helm to deploy Spark for processing big data.

Enter the following command:

```
helm install -- name myspark stable / spark
```

The result is as follows:

```
NAME : myspark  
LAST DEPLOYED : Mon Nov 20 19 : 24 : 22 2017  
NAMESPACE : default  
STATUS : DEPLOYED  
...
```

Use the following commands to view the release and service of Spark.

```
helm list  
kubectl get svc
```

Use the following command to view the Spark related pods and wait until the status changes to Running. Pulling images takes some time because the Spark related images are large.

```
kubectl get pod
```

Use the following command to obtain the Spark Web UI access address:

```
echo http ://$( kubectl get svc myspark - webui - o  
jsonpath='{. status . loadBalancer . ingress [ 0 ]. ip }'): 8080
```

Access the preceding URL in the browser, and you can see the Spark Web UI, on which indicating currently three worker instances exist.

Then, use the following command to use Helm to upgrade the Spark application and change the number of worker instances from three to four. The parameter name is case sensitive.

```
helm upgrade myspark --set "Worker.Replicas=4" stable/spark
```

The result is as follows:

```
Release "myspark" has been upgraded. Happy Helming!
LAST DEPLOYED: Mon Nov 20 19:27:29 2017
NAMESPACE: default
STATUS: DEPLOYED
...
```

Use the following command to view the newly added pods of Spark and wait until the status changes to Running.

```
kubectl get pod
```

Refresh the Spark Web UI in the browser. The number of worker instances changes to four.

To completely delete the Spark application, enter the following command:

```
helm delete --purge myspark
```

Use third-party chart repository

Besides the preset Alibaba Cloud chart repository, you can also use the third-party chart repository (make sure the network is accessible). Add the third-party chart repository in the following command format:

```
helm repo add repository name repository URL
helm repo update
```

For more information about the Helm related commands, see [Helm document](#).

References

Helm boosts the growth of communities. More and more software providers, such as Bitnami, have begun to provide high-quality charts. You can search for and discover existing charts at <https://k8sapps.com/>.

1.8.9 Use an application trigger

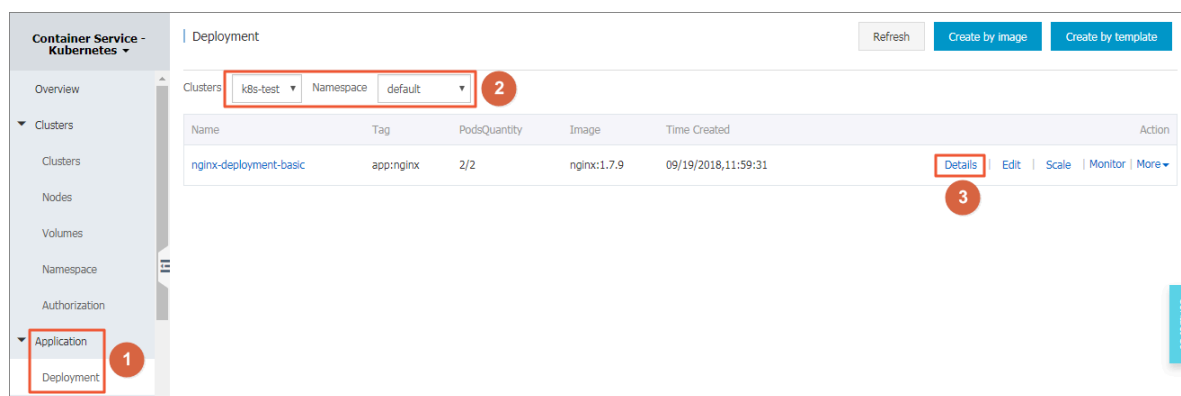
Alibaba Cloud Container Service Kubernetes supports the application trigger function. You can use an application trigger in many ways.

Prerequisites

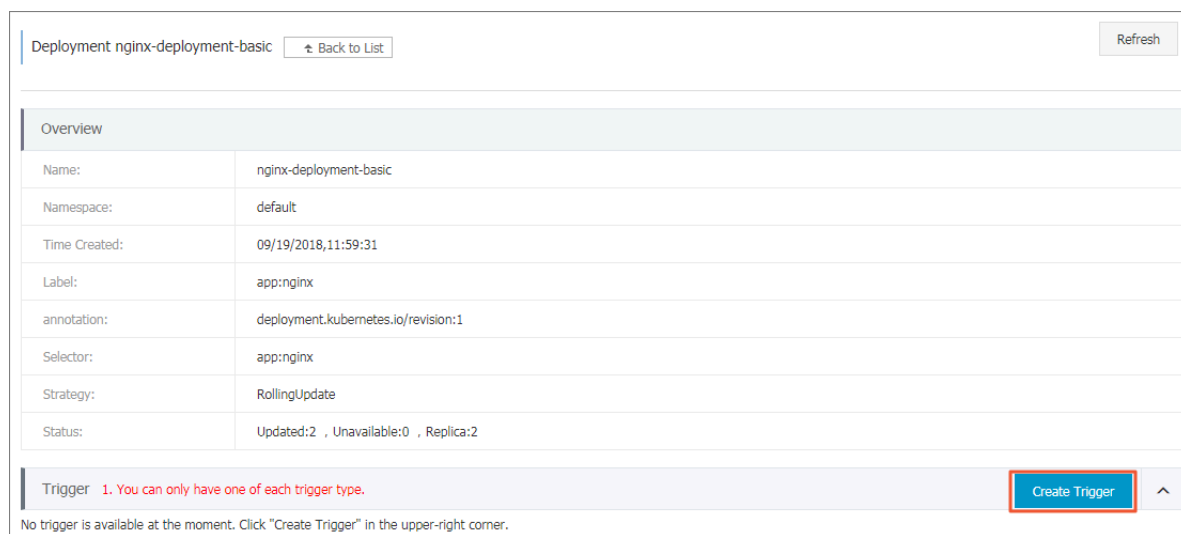
- You have created a Kubernetes cluster. For more information, see [Create a Kubernetes cluster](#).
- You have created an application that is used to create an application trigger and test the trigger. In this example, create an nginx application.

Procedure

1. Log on to the [Container Service console](#).
2. Click Application > Deployment and select a cluster and namespace. Click Details at the right of the target nginx application.



3. On the nginx application details page, click Create Trigger on the right side of the trigger bar.



4. In the pop-up dialog box, click Redeploy and click Confirm.



Note:

Currently, only the redeploy action is supported.

Create Trigger

* Action :

Redeploy

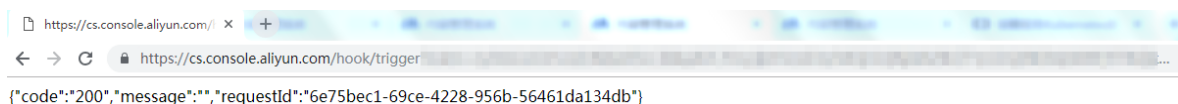
Confirm

Cancel

After the trigger is created, a trigger link is displayed in the trigger bar on the **nginx application detail page**.

[illegible]

5. Copy the trigger link and visit it in the browser. A message is returned on the web page, containing information such as the request ID.



6. Back to the nginx application detail page, you can see that a new pod appears.

Pods	Access	Events	Horizontal Pod Autoscaler
Name		Status	Image
nginx-deployment-basic-6898cc69fb-9726v		<div><div></div>Running</div>	nginx:1.7.9
nginx-deployment-basic-6898cc69fb-9nlms		<div><div></div>Running</div>	nginx:1.7.9

After a period of time, the nginx application removes the old pod and keeps only the new pod.

What's next

You can call a trigger by using GET or POST in a third-party system. For example, you can run the `curl` command to call a trigger.

Call the redeploy trigger as follows:

```
curl https://cs.console.aliyun.com/hook/trigger?token=
xxxxxxx
```

1.8.10 View pods

You can view the pods of a Kubernetes cluster in the Container Service console or in the Kubernetes dashboard.

View pods in Container Service console

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click **Application > Pods** in the left-side navigation pane to go to the Pods page.
3. Select the target cluster and namespace, the target pod, and click **Details** on the right.



Note:

You can update or delete a pod. For pods created by using deployments, we recommend that you manage these pods by using deployments.

4. View the pod details.

View pods in Kubernetes dashboard

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click **Clusters** in the left-side navigation pane.
3. Click **Dashboard** at the right of the cluster to enter the Kubernetes dashboard.
4. In the Kubernetes dashboard, click **Pods** in the left-side navigation pane to view the pods in the cluster.

You can also click **Services** in the left-side navigation pane and then click the service name to view the pods in this service.

5. You can check the status of each Kubernetes object according to the icon on the left. indicates the object is still being deployed. indicates the object has completed the deployment.
6. Click the pod name to view the details, CPU usage, and memory usage of the pod.
7. Click LOGS in the upper-right corner to view the pod logs.
8. You can also click the icon at the right of the pod and then select Delete to delete the pod.

1.8.11 Create a service

This topic describes how to create a Kubernetes service in Alibaba Cloud Container Service for Kubernetes.

A Kubernetes service, known as a service in this and related topics of Alibaba Cloud Container Service for Kubernetes, is an abstract object that defines a logical set of pods and a policy through which to access the pods. Usually, a label selector determines which set of pods are targeted by a service.

In a Kubernetes cluster, each pod has its own IP address, and the pods of a deployment can be removed at any time. However, this action changes the IP addresses of the pods. As a result, directly using IP addresses of pods is ineffective as the scenario does not provide high availability. By comparison, a Kubernetes service decouples the relationship between the frontend and the backend. Specifically, a Kubernetes service is a loose coupling service solution where the operations of the backend do not impact the frontend.

For more information, see [Kubernetes service](#).

Prerequisite

A Kubernetes cluster is created. For more information, see [Create a Kubernetes cluster](#).

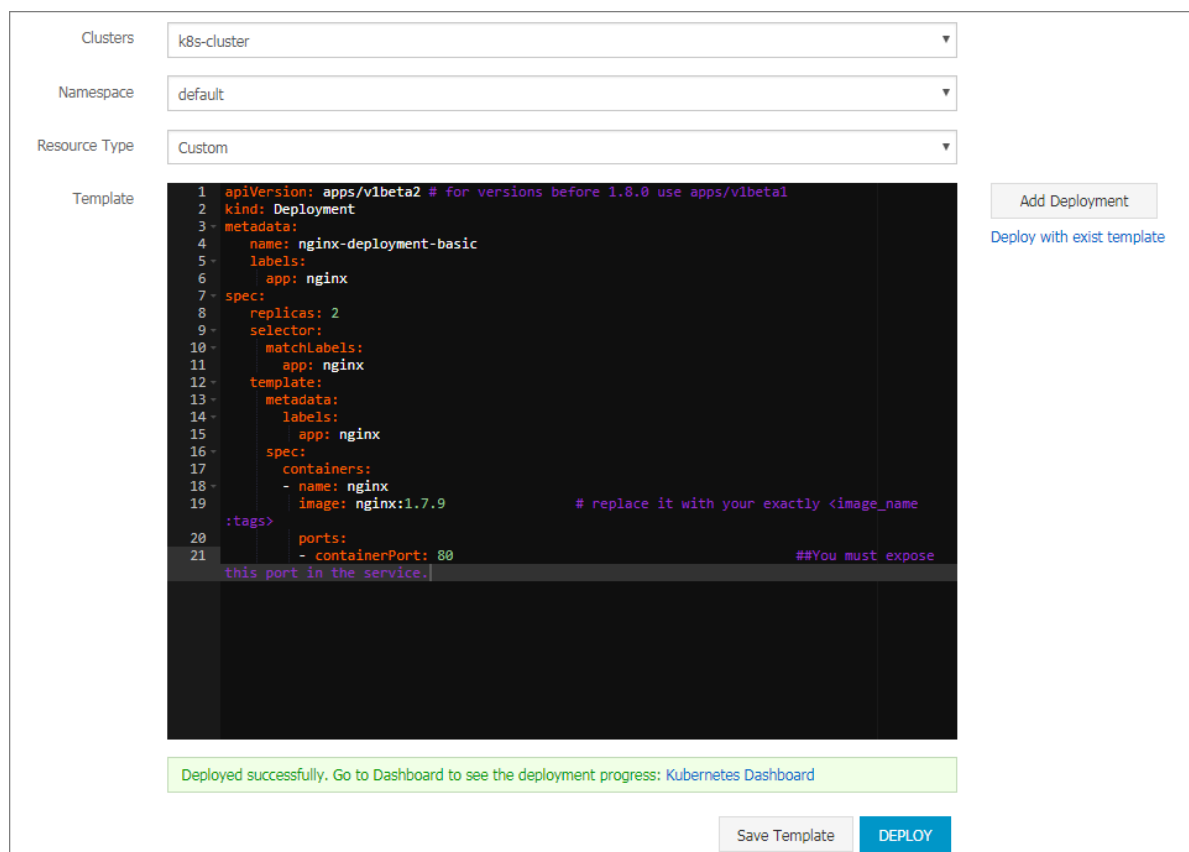
Step 1: Create a deployment

1. Log on to the [Container Service console](#).

2. In the left-side navigation pane under Kubernetes, choose Applications > Deployments. Then click Create by Template in the upper-right corner.



3. Select the target cluster and namespace, and select a custom template or a sample template from the Resource Type drop-down list. Then, click DEPLOY.



In this example, the sample template specifies an Nginx deployment.

```

apiVersion : apps / v1beta2 # for versions before 1 . 8 .
0 use apps / v1beta1
kind : Deployment
metadata :
  name : nginx - deployment - basic
  labels :
    app : nginx
spec :
  replicas : 2
  selector :
  
```

```

matchLabels:
  app: nginx
template:
  metadata:
    labels:
      app: nginx
  spec:
    containers:
      - name: nginx
        image: nginx:1.7.9 # replace it
        ports:
          - containerPort: 80
  ## This port must be exposed in a service.

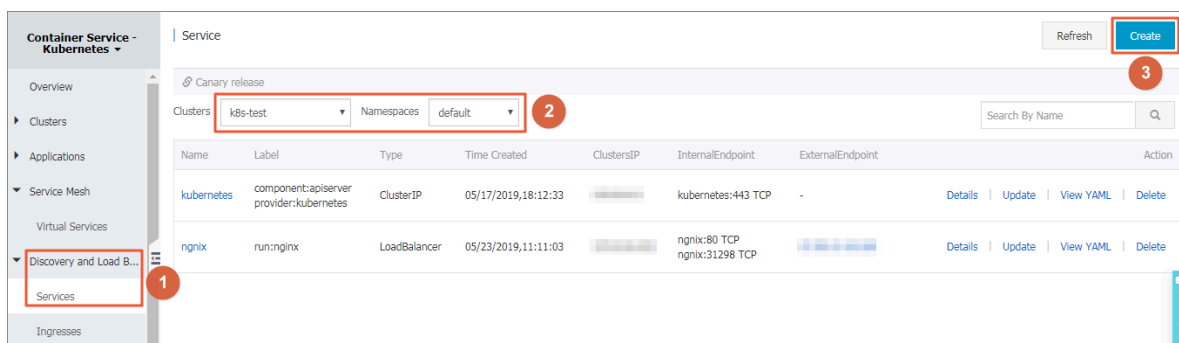
```

4. Click Kubernetes Dashboard to view the running status of this deployment.



Step 2: Create a service

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Container Service-Kubernetes, choose **Discovery and Load Balancing > Services**.
3. Select the target cluster and namespace. Then, click **Create** in the upper-right corner.



4. In the displayed dialog box, set service parameters.

Create Service

Name:

nginx-svc

Type:

Server Load Balancer

public

Related:

nginx-deployment-basic

Port Mapping:

+ Add

service port	Container Port	Protocol	
80	80	TCP	-

annotation:

+ Add Annotations for load balancer

Name	Value	
service.beta.kubernetes.io	20	-

Tag:

+ Add

Name	Value	
app	nginx	-

Create

Cancel

- **Name:** Enter the service name. In this example, the service name is set to nginx-svc.
- **Type:** Select the service type, namely, the service access method.
 - **Cluster IP:** Exposes the service by using the internal IP address of your cluster. If you select this service type, the service is accessible only within the cluster. This is the default service type.

- **Node port:** Exposes the service by using the IP address and the static port (NodePort) of each node. A node port service routes to a cluster IP service that is automatically created. You can access the node port service from outside the cluster by requesting `< NodeIP > : < NodePort >`.
- **Server Load Balancer:** Alibaba Cloud Server Load Balancer service. With this type of service, you can set an Internet or intranet access method for your application. SLB can route to a node port service and a cluster IP service.
- **Related:** Select the backend object to associate with the service. In this example, the nginx-deployment-basic deployment created in the preceding step is associated with the service. If you do not associate the service with any objects, the system does not create any corresponding endpoint objects. In this case, you can manually associate the service with your own specific endpoints. For more information, see [Services without selectors](#).
- **Port Mapping:** Add a service port number and a container port number. The container port number that you set must be the same as the port number of the container exposed by the pod.
- **annotation:** Add an annotation to the service. You can set SLB parameters. For example, to control the service traffic, you can set the peak bandwidth of the service to 20 Mbit/s by setting this parameter as `service.beta.kubernetes.io/alibaba-cloud-loadbalancer-bandwidth: 20`. For more information, see [Access services by using Server Load Balancer](#).
- **Tag:** Add a tag to the service to identify the service.

5. Click Create. The nginx-svc service is then displayed in the service list.

Service						Refresh	Create
Help: Canary release							
Clusters	k8s-cluster	Namespace	default	Search By Name		Q	
Name	Type	Time Created	ClustersIP	InternalEndpoint	ExternalEndpoint	Action	
kubernetes	ClusterIP	02/05/2019,15:58:07		kubernetes:443 TCP	-	Details	Update View YAML Delete
nginx-svc	LoadBalancer	02/16/2019,15:17:13		nginx-svc:80 TCP nginx-svc:31200 TCP		Details	Update View YAML Delete

6. Enter the external endpoint of the nginx-svc service in your browser to access the service.

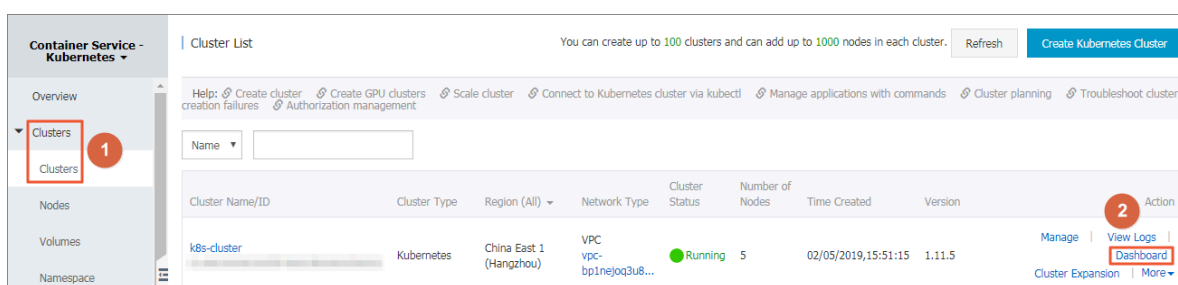


1.8.12 Scale a service

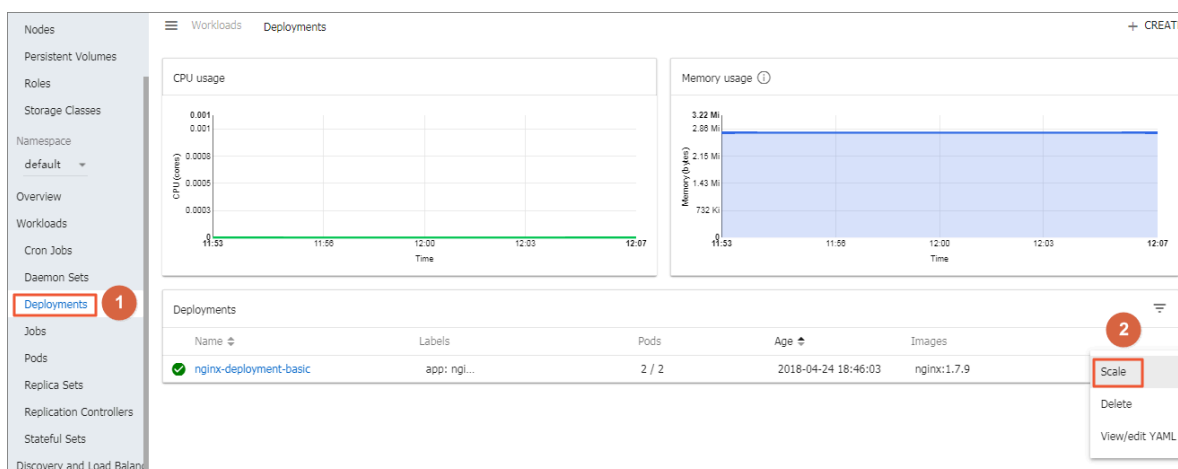
This topic describes how to scale out or scale in an application service as needed after an application is created.

Procedure

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Kubernetes, choose Clusters > Clusters.
3. On the right of the target cluster, click Dashboard.

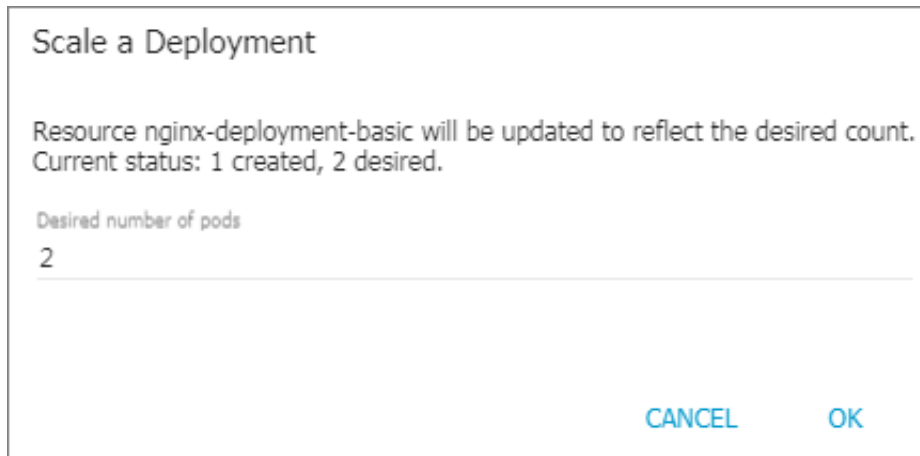


4. In the left-side navigation pane, click Deployments.
5. Click the  icon on the right of the target deployment, and then click Scale.





6. In the displayed dialog box, change the value of Desired number of pods to the number you require. Here, the example number of desired pods is 2. Then, click OK.

This action adds a new pod. The number of replicas becomes 2.



A dialog box titled "Scale a Deployment". The text inside says: "Resource nginx-deployment-basic will be updated to reflect the desired count. Current status: 1 created, 2 desired." Below this, there is a label "Desired number of pods" and a text input field containing the number "2". At the bottom right, there are two buttons: "CANCEL" and "OK".

What's next

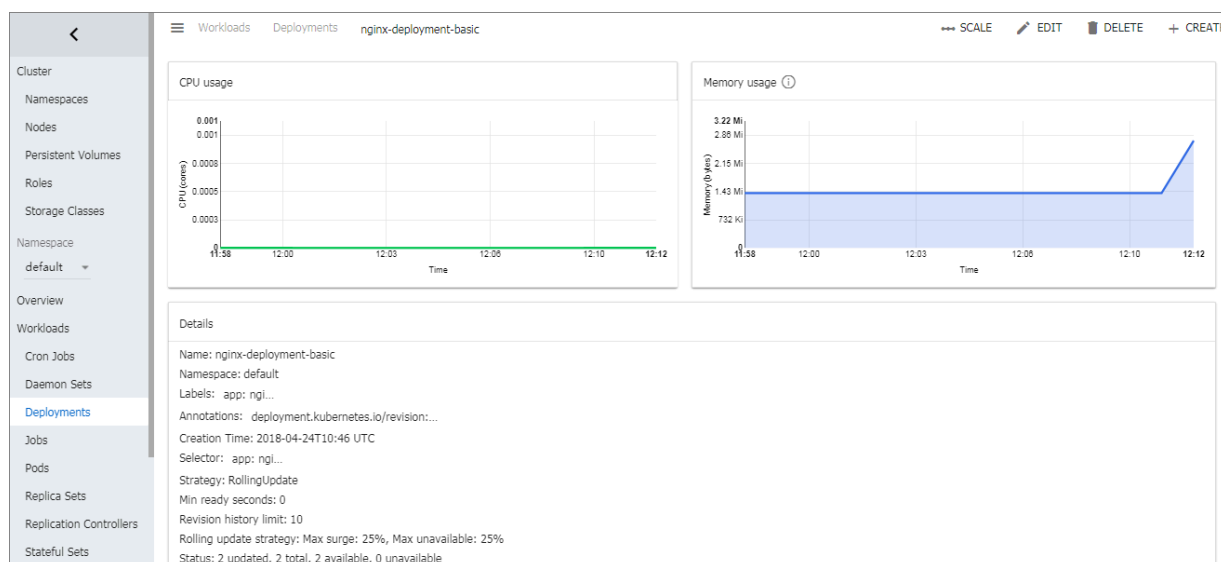
You can check the status of each Kubernetes object according to the icon on the left of the deployment list.  indicates the object is being deployed.  indicates the object has been deployed.

Additionally, you can click a deployment name to view the details of the running Web service. Specifically, you can view the replica sets included in the deployment, and the CPU usage and memory usage of these replica sets.



Note:

If no resources are displayed, we recommend that you wait a few minutes and then refresh the page.



1.8.13 View a service

This topic describes how to view a Kubernetes service in Alibaba Cloud Container Service for Kubernetes.

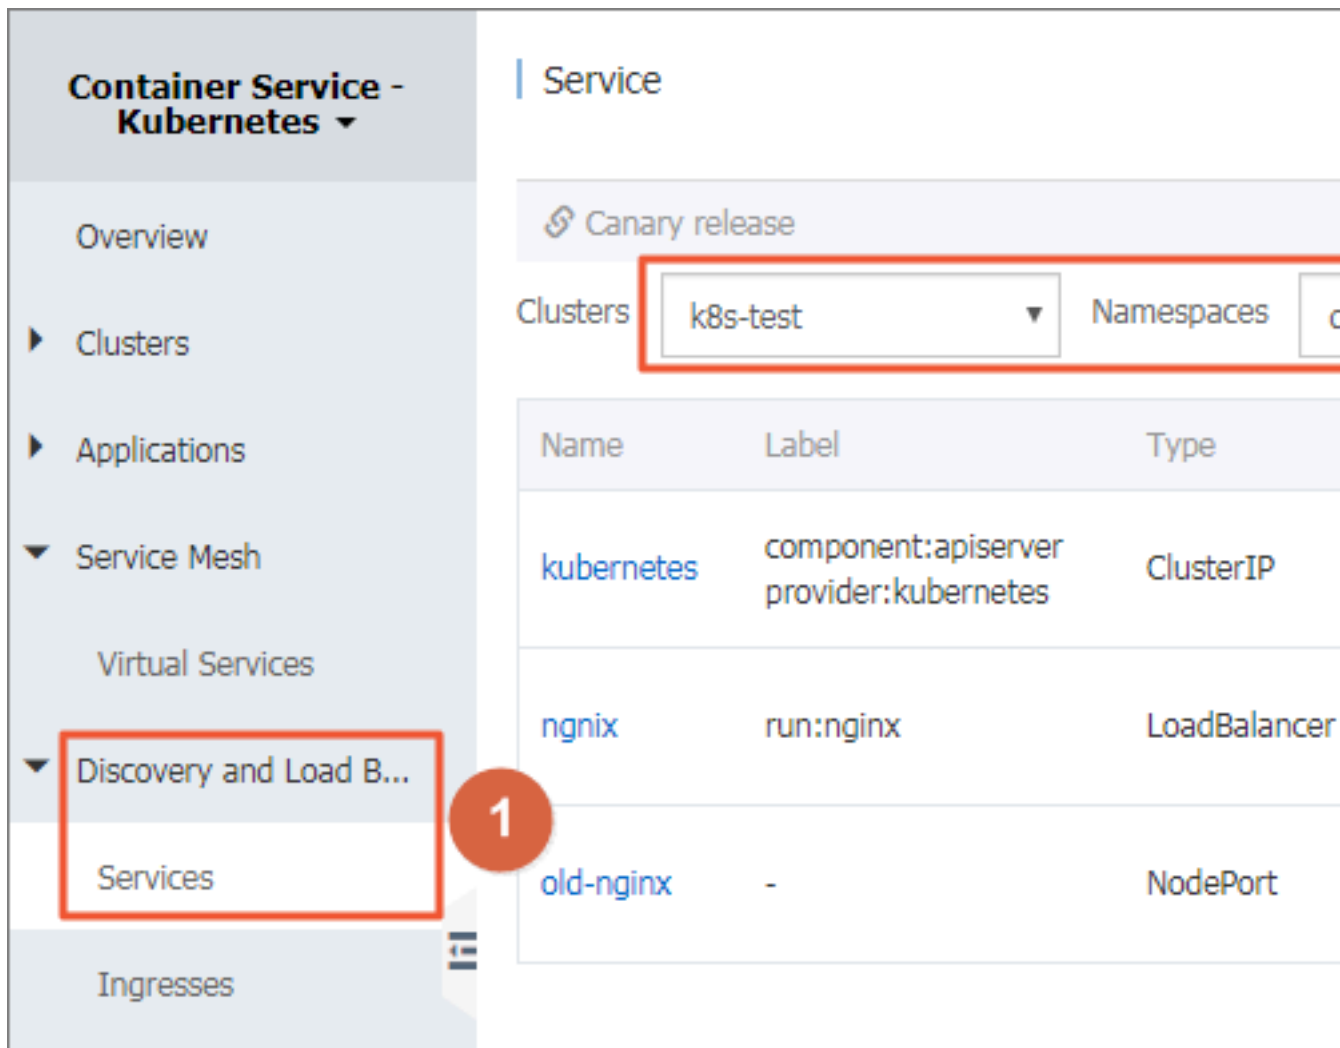
Context

If you set an external service when you create an application, the Kubernetes dashboard creates the external service, in addition to running containers. The service is used to pre-set a Server Load Balancer to distribute traffic to the containers.

Procedure

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Container Service-Kubernetes, choose **Discovery and Load Balancing > Services**.

3. Select the target cluster and namespace, and then click Details on the right of the target service.



The screenshot shows the 'Container Service - Kubernetes' console. On the left sidebar, the 'Discovery and Load B...' menu item is highlighted with a red box and a red circle with the number 1. The 'Clusters' dropdown is also highlighted with a red box and contains 'k8s-test'. The 'Namespaces' dropdown is partially visible. The table below shows three services:

Name	Label	Type
kubernetes	component:apiserver provider:kubernetes	ClusterIP
nginx	run:nginx	LoadBalancer
old-nginx	-	NodePort

You can view the service name, service type, service creation time, cluster IP address, external endpoint, and other information. In this example, the external

endpoint (IP address) assigned to the service is displayed. To access the Nginx application, click this IP address.

nginx-svc ↶ Back to List	
Overview	
Name:	nginx-svc
Namespace:	default
Time Created:	02/16/2019,16:41:42
Label:	app:nginx-v2
annotation:	service.beta.kubernetes.io/alibaba-cloud-loadbalancer-bandwidth:20
Type:	LoadBalancer
ClustersIP:	
InternalEndpoint:	nginx-svc:80 TCP nginx-svc:32397 TCP
ExternalEndpoint:	

You can also open the Kubernetes dashboard of the target cluster and click Services in the left-side navigation pane to view all services of the cluster.

1.8.14 Update a service

This topic describes how to update a Kubernetes service in the Container Service console or the Kubernetes dashboard.

Update a service in the Container Service console

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Kubernetes, choose Discovery and Load Balancing > Service.

3. Select the target cluster and namespace, and then click Update on the right of the target service. In this example, the target service is named nginx-svc.

The screenshot displays the 'Service' management interface in the Aliyun Container Service console. The left-hand navigation pane shows the 'Service' option selected under the 'Discovery and Load Balancing' category, indicated by a red circle with the number 1. The main content area features a table of services. At the top of this area, the 'Clusters' dropdown is set to 'k8s-cluster' and the 'Namespace' dropdown is set to 'default', both highlighted with a red rectangle and a red circle with the number 3. Below the table, the 'nginx-svc' service is highlighted, and its 'Update' button is circled in red, marked with a red circle with the number 4. The table lists the following services:

Name	Type	Time Created	ClustersIP	InternalEndpoint	ExternalEndpoint	Action
kubernetes	ClusterIP	02/05/2019,15:58:07	[IP]	kubernetes:443 TCP	-	Details Update View YAML Delete
new-nginx	NodePort	02/17/2019,09:26:21	[IP]	new-nginx:80 TCP new-nginx:32580 TCP	-	Details Update View YAML Delete
nginx-svc	LoadBalancer	02/16/2019,16:41:42	[IP]	nginx-svc:80 TCP nginx-svc:32397 TCP	[IP]	Details Update View YAML Delete
old-nginx	NodePort	02/17/2019,09:25:20	[IP]	old-nginx:80 TCP old-nginx:30675 TCP	-	Details Update View YAML Delete

4. In the displayed dialog box, update the service parameters. Then, click Update. In this example, the service tag is updated from app:nginx-v1 to app:nginx-v2.

Update Service

Name: nginx-svc

Type:

Server Load Balancer

public

Port Mapping:

+ Add

service port	Container Port	Protocol	
80	80	TCP	-

annotation:

+ Add Annotations for load balancer

Name	Value	
service.beta.kubernetes.io	20	-

Tag:

+ Add

Name	Value	
app	nginx-v1	-

Update

Cancel

5. In the service list, click **Details** on the right of the target service to view the service updates. In this example, the updated service tag `app:nginx-v2` is displayed.

nginx-svc ⬆ Back to List	
Overview	
Name:	nginx-svc
Namespace:	default
Time Created:	02/16/2019,16:41:42
Label:	app:nginx-v2
annotation:	service.beta.kubernetes.io/alibaba-loadbalancer-bandwidth:20
Type:	LoadBalancer
ClustersIP:	
InternalEndpoint:	nginx-svc:80 TCP nginx-svc:32397 TCP
ExternalEndpoint:	:80

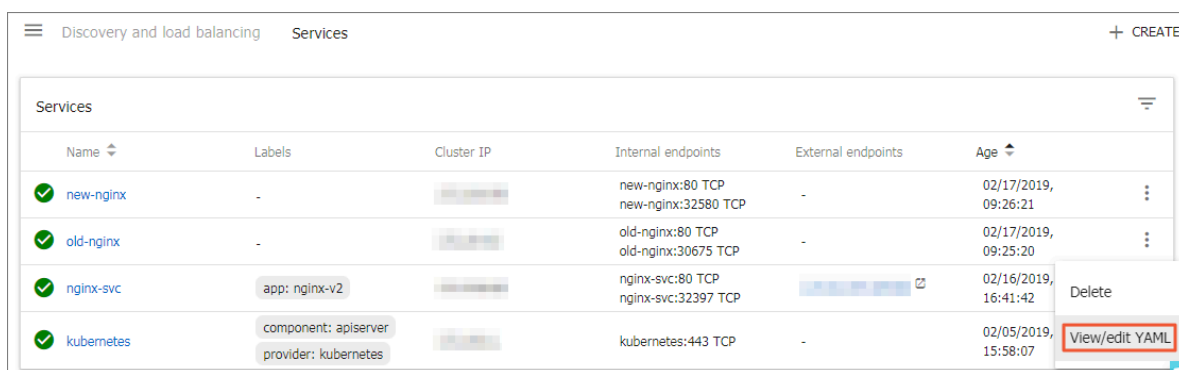
Update a service in the Kubernetes dashboard

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

Container Service - Kubernetes Overview Clusters Nodes Volumes Namespace	Cluster List							
	You can create up to 100 clusters and can add up to 1000 nodes in each cluster. Refresh Create Kubernetes Cluster Help: Create cluster creation failures Create GPU clusters Scale cluster Connect to Kubernetes cluster via kubectl Manage applications with commands Cluster planning Troubleshoot cluster							
	Name							
	Cluster Name/ID	Cluster Type	Region (All)	Network Type	Cluster Status	Number of Nodes	Time Created	Version
	k8s-cluster	Kubernetes	China East 1 (Hangzhou)	VPC vpc-bp1nejq3u8...	Running	5	02/05/2019,15:51:15	1.11.5
								Manage View Logs Dashboard Cluster Expansion More

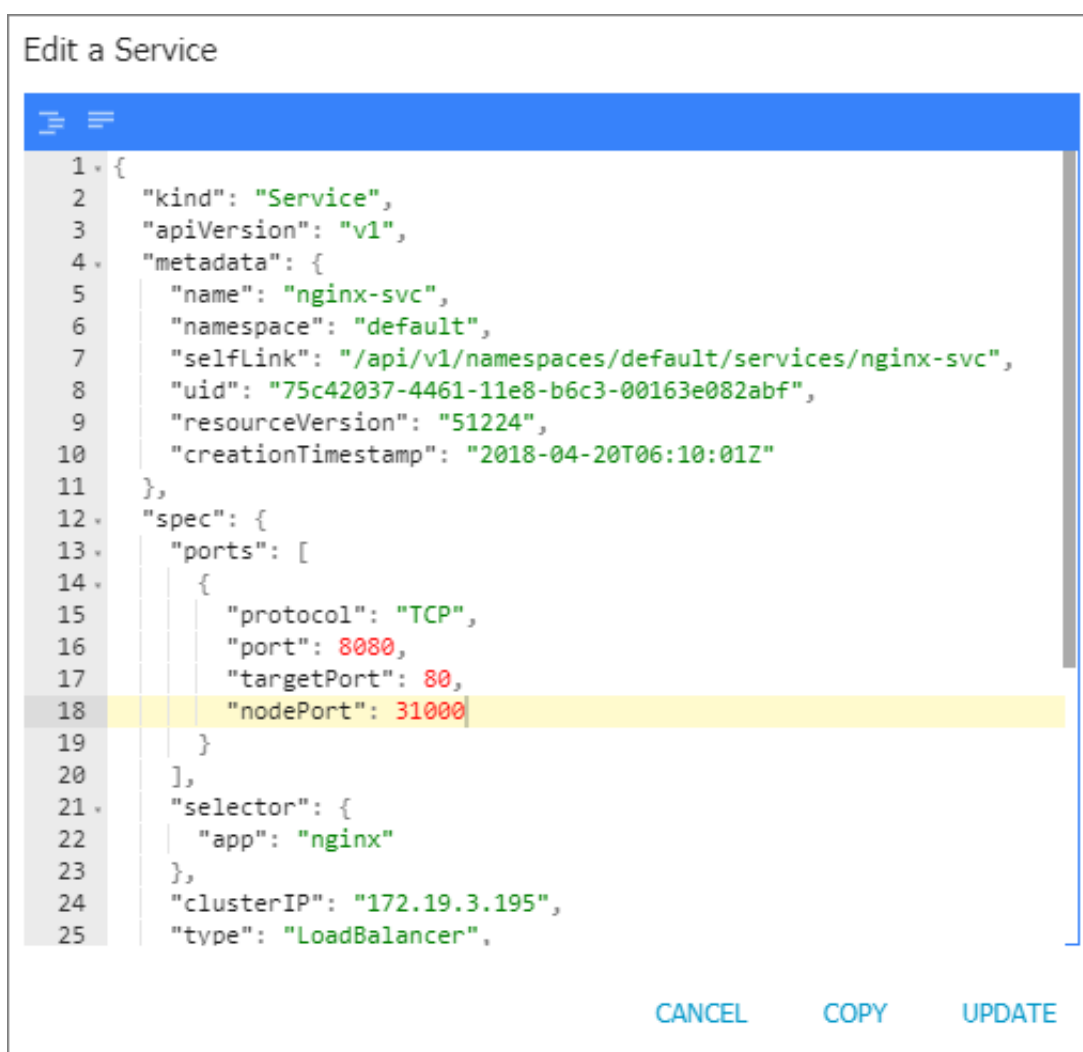
4. In the Kubernetes dashboard, select the target namespace and then click **Services** in the left-side navigation pane.

5. Click the icon on the right of the target service and then click View/edit YAML.



Name	Labels	Cluster IP	Internal endpoints	External endpoints	Age	
new-nginx	-		new-nginx:80 TCP new-nginx:32580 TCP	-	02/17/2019, 09:26:21	⋮
old-nginx	-		old-nginx:80 TCP old-nginx:30675 TCP	-	02/17/2019, 09:25:20	⋮
nginx-svc	app: nginx-v2		nginx-svc:80 TCP nginx-svc:32397 TCP		02/16/2019, 16:41:42	⋮ Delete
kubernetes	component: apiserver provider: kubernetes		kubernetes:443 TCP	-	02/05/2019, 15:58:07	⋮ View/edit YAML

6. In the displayed dialog box, modify the service settings. Then, click UPDATE. In this example, the nodePort is changed to 31000.



```

1. {
2.   "kind": "Service",
3.   "apiVersion": "v1",
4.   "metadata": {
5.     "name": "nginx-svc",
6.     "namespace": "default",
7.     "selfLink": "/api/v1/namespaces/default/services/nginx-svc",
8.     "uid": "75c42037-4461-11e8-b6c3-00163e082abf",
9.     "resourceVersion": "51224",
10.    "creationTimestamp": "2018-04-20T06:10:01Z"
11.  },
12.  "spec": {
13.    "ports": [
14.      {
15.        "protocol": "TCP",
16.        "port": 8080,
17.        "targetPort": 80,
18.        "nodePort": 31000
19.      }
20.    ],
21.    "selector": {
22.      "app": "nginx"
23.    },
24.    "clusterIP": "172.19.3.195",
25.    "type": "LoadBalancer",

```

CANCEL COPY UPDATE

1.8.15 Delete a service

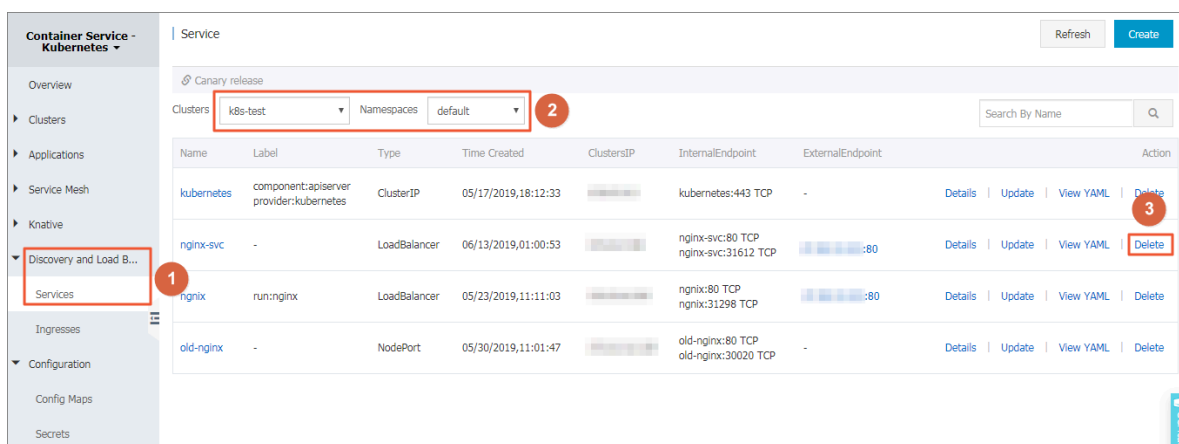
This topic describes how to delete a Kubernetes service in the Alibaba Cloud Container Service console.

Prerequisites

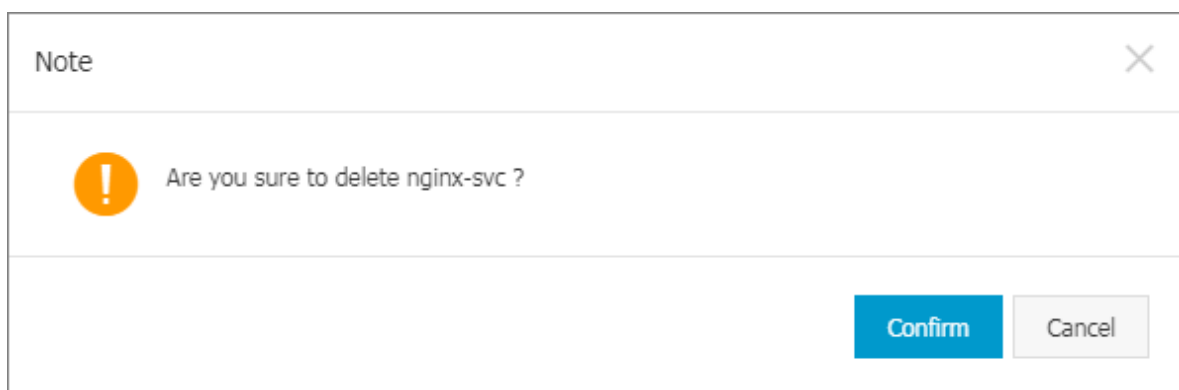
- A Kubernetes cluster is created. For more information, see [Create a Kubernetes cluster](#).
- A service is created. For more information, see [Create a service](#).

Procedure

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Kubernetes, choose Discovery and Load Balancing > Service.
3. Select the target cluster and namespace, and then click Delete on the right of the target service. In this example, the target service is named nginx-svc.



4. In the displayed dialog box, click Confirm.



1.8.16 Change container configurations

You can change the container configurations in the Container Service console.

Procedure

1. Log on to the [Container Service console](#).
2. Click Kubernetes > Clusters in the left-side navigation pane.

3. Click Dashboard at the right of the cluster to enter the Kubernetes dashboard.
4. In the Kubernetes dashboard, click Pods in the left-side navigation pane.
5. Click the icon at the right of the pod and then select View/edit YAML.
6. The Edit a Pod dialog box appears. Change the container configurations and then click UPDATE.

1.8.17 Schedule a pod to a specified node

You can add a node label and then configure the `nodeSelect` or to schedule a pod to a specified node. For more information about the implementation principle of nodeSelector, see [nodeselector](#).

For business scenario needs, to deploy a service used for management and control to a master node, or deploy services to a machine with an SSD disk, you can use this method to schedule pods to specified nodes.

Prerequisites

You have successfully created a Kubernetes cluster. For more information, see [Create a Kubernetes cluster](#).

Step 1 Add a node label

1. Log on to the [Container Service console](#).
2. Under the Kubernetes menu, click Clusters > Nodes in the left-side navigation pane.
3. Select the cluster from the Cluster drop-down list and then click Label Management in the upper-right corner.
4. Select one or more nodes by selecting the corresponding check boxes and then click Add Tag. In this example, select a worker node.

5. Enter the name and value of the label in the displayed dialog box and then click OK.

The node label `group : worker` is displayed on the Label Management page.

You can also add a node label by running the command `kubectl label nodes < node - name > < label - key >=< label - value >`.

Step 2 Deploy a pod to a specified node

1. Log on to the [Container Service console](#).
2. Under the Kubernetes menu, click Applications > Deployment in the left-side navigation pane.
3. Click Create by template in the upper-right corner.
4. Configure the template to deploy a pod. After completing the configurations, click DEPLOY.
 - Clusters: Select a cluster.
 - Namespace: Select the namespace to which the resource object belongs. In this example, use default as the namespace.
 - Resource Type: Select Custom in this example.

The orchestration template in this example is as follows:

```
apiVersion : v1
kind : Pod
metadata :
  labels :
    name : hello - pod
  name : hello - pod
spec :
  containers :
    - image : nginx
      imagePullPolicy : IfNotPresent
      name : hello - pod
      ports :
        - containerPort : 8080
          protocol : TCP
      resources : {}
      securityContext :
        capabilities : {}
        privileged : false
        terminationMessagePath : /dev/termination-log
      dnsPolicy : ClusterFirst
      restartPolicy : Always
      nodeSelector :
```

```
group : worker ## The same as the node label
configured in the preceding step .
status :{}
```

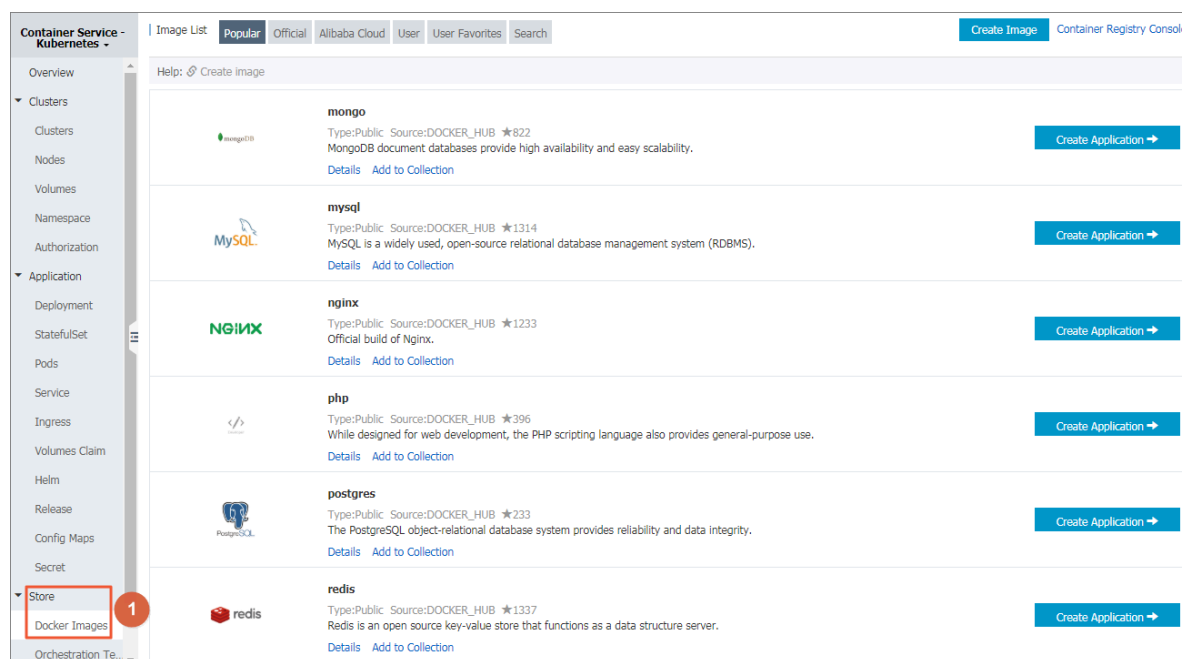
5. A message indicating the deployment status is displayed after you click **DEPLOY** .
After the successful deployment, click **Kubernetes Dashboard** in the message to go to the dashboard and check the deployment status.
6. Click the pod name to view the pod details.

You can view the information such as the pod label and node ID, which indicates the pod is successfully deployed to a node with the label `group : worker` .

1.8.18 View image list

Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click **Store > Docker Images** in the left-side navigation pane.



You can view the image category.

- **Popular:** Some common images recommended by Container Service.
- **Official:** Official images provided by Docker Hub.

1.8.19 Use an image Secret

Container Service Kubernetes clusters support using image secrets through the web interface. You can create an image secret and use an existing image secret.

Prerequisites

- You have created a Kubernetes cluster. For more information, see [Create a Kubernetes cluster](#).
- You have built a private image repository and uploaded your image to the repository. In this example, use Alibaba Cloud Container Registry. For more information, see [Use a private image repository to create an application](#).

Context

When you use a private image to create an application, you have to configure a secret for the image to secure the image. In the Container Service console, you can deliver the identity authentication information of the private image repository to Kubernetes through a secret of the docker-registry type.

Procedure

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Container Service-Kubernetes, choose Applications > Deployments. Then, click Create by Image in the upper-right corner.
3. Configure Name, Cluster, Namespace, Replicas, and Type. The configured value of the replicas parameter specifies the number of pods contained in the application. Click Next.



Note:

In this example, select the Deployment type.

If you do not configure Namespace, the system uses the default namespace by default.

The screenshot shows a configuration interface with four tabs: Basic Information, Container, Advanced, and Done. The Basic Information tab is active. It contains the following fields:

- Name:** A text input field containing "tomcat". Below it, a note states: "The name should be 1-64 characters long, and can contain numbers, lower case English letters and hyphens, but cannot start with a hyphen."
- Cluster:** A dropdown menu with a small icon on the left.
- Namespace :** A dropdown menu containing the value "default".
- Replicas:** A text input field containing the value "2".
- Type:** A dropdown menu containing the value "Deployment".

At the bottom right of the form are two buttons: "Back" and "Next".

4. Configure containers.



Note:

This example describes only the configuration of the container image secret. For more information about container configuration, see [Create a deployment application by using an image](#).

5. On the container configuration page, configure the image name first. Enter the private image address in the Image Name box. The format is `domainname / namespace / imagename`.



Note:

Public images do not require image secrets.

6. In the image version box, enter the private image address version.

The screenshot shows a configuration interface for containers. At the top, there is a tab labeled "Container1" and a button with a plus sign and the text "Add Container". Below this, there are two rows of configuration fields:

- Image Name:** A text input field containing "registry.cn-hangzhou.aliyuncs.com/dev-". To its right is a blue link labeled "Select image".
- Image Version:** A text input field containing "V1". To its right is a blue link labeled "Select image version".

7. Click Image pull secret.

- Select Create secret.
 - **Name:** Specifies the secret name. You can define it by yourself.
 - **Repository Domain Name:** Specified the Docker repository address. If you enter the Alibaba Cloud Container Service image repository in the image name box, the system automatically adds the repository address by default.
 - **Username:** Specifies the user name of the Docker repository. If you use Alibaba Cloud Container Registry, the username is your Alibaba Cloud account name.
 - **Password:** Specifies the logon password of the Docker repository. If you use Alibaba Cloud Container Registry, the password is the independent logon password for Container Registry.
 - **Email:** Specifies an email address. This is optional.

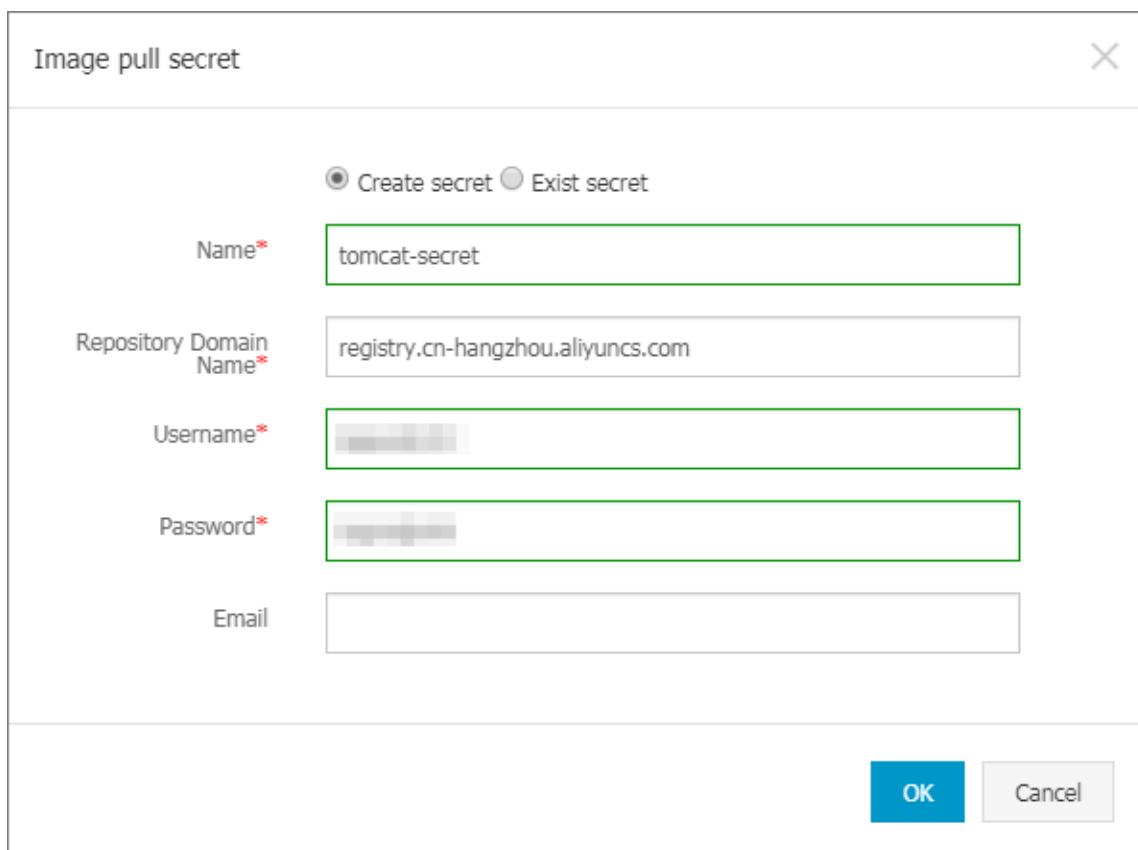


Image pull secret

☒ Create secret ☐ Exist secret

Name* tomcat-secret

Repository Domain Name* registry.cn-hangzhou.aliyuncs.com

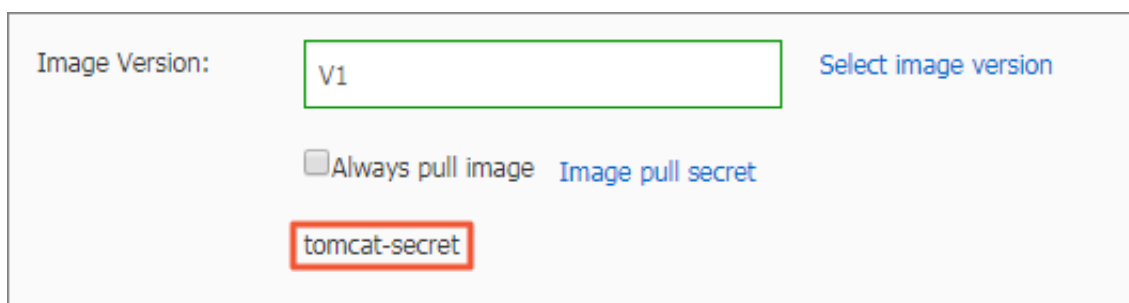
Username*

Password*

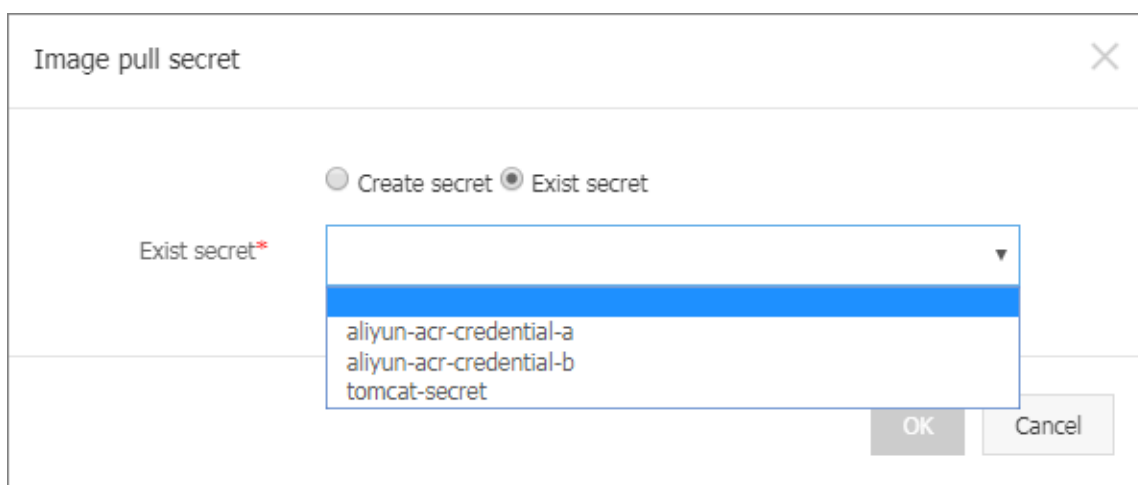
Email

OK Cancel

Click OK. The created secret is displayed on the page.



- You can also click **Exist secret**. You can pre-create a container image secret by using command lines or a YAML file. For information, see [How to use private images in Kubernetes clusters](#) and [Use a private image repository to create an application](#).

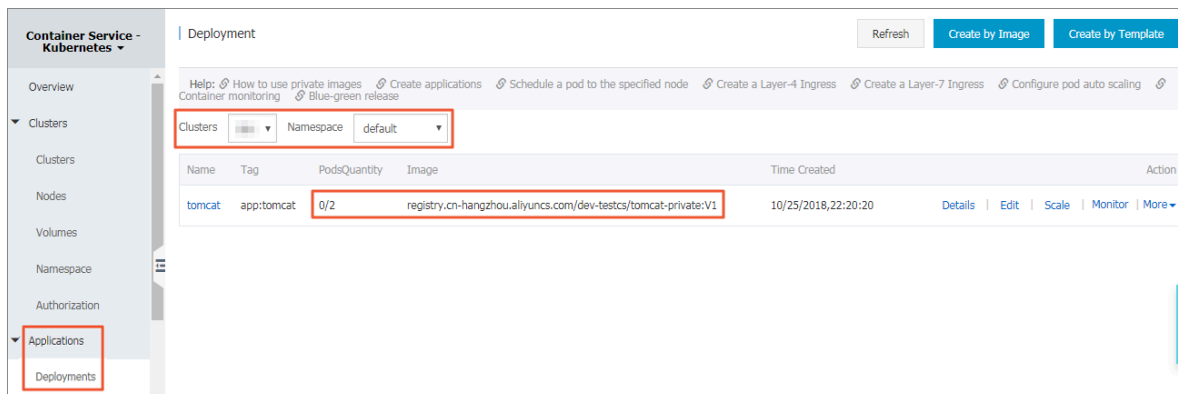


8. After you complete the container configuration, click **Next**.
9. Follow the page guide to complete other configurations, and then click **Create**.
10. Click **Applications > Deployments** in the left-side navigation pane, and select the cluster and namespace in which the application is created to view the status of the tomcat application.



Note:

The system shows that the tomcat application runs properly, which indicates that you have used the tomcat private image through the secret.



1.8.20 Pull a private image without a password

This topic describes how to pull a private image without a password from the Alibaba Cloud container image repository.

Prerequisites

You have created a Kubernetes cluster. For more information, see [Create a Kubernetes cluster](#).

Context

Function overview

- You can only pull a private image from an Alibaba Cloud container image repository that belongs to your account.
- You can pull a private image from a cross-region Alibaba Cloud container image repository.
- You can only perform this operation in the default namespace.
- Kubernetes clusters that support this function include:
 - Dedicated Kubernetes clusters
 - Managed Kubernetes clusters
 - Serverless Kubernetes clusters

- The following are Kubernetes cluster versions that support this function:
 - Dedicated Kubernetes cluster versions that are not earlier than v1.11.2 support this function by default. If the dedicated Kubernetes cluster version is earlier than v1.11.2, follow the procedures described in this topic.
 - All versions of managed Kubernetes clusters support this function.
 - All versions of serverless Kubernetes clusters support this function.

Procedure

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

Cluster Resource	
ROS	
Internet SLB	
VPC	
NAT Gateway	
Master RAM Role	
Worker RAM Role	



Note:

This topic uses the latest version of the RAM console.

If you use an earlier version of the RAM console, you can modify the target policy document by using either of the following two methods:

Method 1

- a. In the left-side navigation pane, click Roles, and then enter the Worker RAM Role name in the Role Name box. Click the target Role Name.

Role Management			
		Create Role	Refresh
Role Name	K	Search	
Role Name	Created At	Actions	
	2018-12-28 15:46:49	Manage	Authorize Delete

b. In the Basic Information area, click Edit Basic Information in the upper-right corner.

Basic Information		Edit Basic Information
Role Name	Ku	Description
Created At	2018-12-28 15:46:49	Arn

Method 2

In the lower-right corner of the RAM dashboard page, click New Version to switch to the latest version of the RAM console. In the Container Service console, click Worker RAM Role to log on to the RAM console.

The screenshot shows the RAM dashboard with a sidebar on the left containing links to Dashboard, Users, Groups, Policies, Roles, Settings, and ActionTrail. The main content area displays a welcome message and four overview cards: User Overview (6 users), Group Overview (0 groups), Permission Policy Overview (15 custom authorization policies), and Role Overview (28 roles). Below these is an 'Operations Guide' section with a numbered list of steps: 1. Manage authorization policies, 2. Create and authorize groups, 3. Create users and add them to groups, and 4. Authorization complete. To the right of the guide is a 'Policy Management' section with a 'System Policy' tab selected. A 'New Authorization Policy' button is circled in red, and a 'New Version' button is also circled in red in the bottom right corner.

5. On the RAM Roles page, click the policy name in the Permission area to view the policy details.

6. On the Policies page, click Modify Policy Document in the Policy Document area.

The screenshot shows the RAM Policies page for a specific policy named 'k8sWorkerRolePolicy-8b896e4c-64bf-4222-b7a9-9bb0d75b2fae'. The left sidebar shows the 'Policies' link selected. The main content area has tabs for 'Policy Document', 'Versions', and 'References'. The 'Policy Document' tab is active, and a 'Modify Policy Document' button is circled in red. Below the button, the policy document is displayed in a code editor with the following content:

```

1 {
2   "Version": "1",
3   "Statement": [
4     {
5       "Action": [
6         "ecs:AttachDisk",
7         "ecs:DetachDisk",
8         "ecs:DescribeDisks",
9         "ecs:CreateDisk",

```

7. In the Policy Document area, add the following fields and then click OK.

```
{
  " Action ": [
    " cr : Get *",
    " cr : List *",
    " cr : PullRepository "
  ],
  " Resource ": "*",
  " Effect ": " Allow "
}
```

Modify Policy Document

Policy Name

k8sWork

Policy Document

```
96      "Resource": [
97      |      "*"
98      ],
99      "Effect": "Allow"
100    },
101    {
102      "Action": [
103      |      "cr:Get*",
104      |      "cr:List*",
105      |      "cr:PullRepository"
106      ],
107      "Resource": "*",
108      "Effect": "Allow"
109    }
110  ]
111 }
```

OKClose

8. Create the `aliyun - acr - credential - helper` service to refresh the temporary token of Container Registry at intervals.

```

apiVersion : v1
kind : ServiceAccount
metadata :
  name : aliyun - acr - credential - helper
  namespace : kube - system
---
apiVersion : rbac . authorization . k8s . io / v1beta1
kind : ClusterRoleBinding
metadata :
  name : aliyun - acr - credential - helper - rolebinding
  namespace : kube - system
roleRef :
  apiGroup : rbac . authorization . k8s . io
  kind : ClusterRole
  name : cluster - admin
subjects :
  - kind : ServiceAccount
    name : aliyun - acr - credential - helper
    namespace : kube - system
---
# kubectl create secret docker - registry acr - image -
pull - secret - public -- docker - server = cr - tmp - xxx --
docker - username = cr - temp - xxx -- docker - password = cr -
temp - xxx -- docker - email = cr - temp - xxx
apiVersion : v1
data :
  .dockerconfig json : eyJhdXRocy I6eyJjci10 bXAteHh4Ij
p7InVzZXJu YW1lIjoiY3 ItdGVtcC14 eHg1LCJwYX Nzdz29yZCI6
ImNyLXRlbX AteHh4Iiw ZW1haWwiOi Jjci10ZW1w LXh4eCI6Im
F1dGgiOiJZ M0l0ZEdWdG NDMTRLSGc2 WTNJdGRHVn RjQzE0ZUhn PSJ9fX0
=
kind : Secret
metadata :
  name : aliyun - acr - credential - a
  namespace : default
type : kubernetes . io / dockerconfig json
---
# kubectl create secret docker - registry acr - image -
pull - secret - vpc -- docker - server = cr - tmp - xxx -- docker
- username = cr - temp - xxx -- docker - password = cr - temp -
xxx -- docker - email = cr - temp - xxx
apiVersion : v1
data :
  .dockerconfig json : eyJhdXRocy I6eyJjci10 bXAteHh4Ij
p7InVzZXJu YW1lIjoiY3 ItdGVtcC14 eHg1LCJwYX Nzdz29yZCI6
ImNyLXRlbX AteHh4Iiw ZW1haWwiOi Jjci10ZW1w LXh4eCI6Im
F1dGgiOiJZ M0l0ZEdWdG NDMTRLSGc2 WTNJdGRHVn RjQzE0ZUhn PSJ9fX0
=
kind : Secret
metadata :
  name : aliyun - acr - credential - b
  namespace : default
type : kubernetes . io / dockerconfig json
---
apiVersion : apps / v1beta2
kind : Deployment
metadata :
  name : aliyun - acr - credential - helper
  namespace : kube - system

```

```

labels :
  app : aliyun - acr - credential - helper
spec :
  replicas : 1
  selector :
    matchLabels :
      app : aliyun - acr - credential - helper
  template :
    metadata :
      labels :
        app : aliyun - acr - credential - helper
    spec :
      serviceAccount : aliyun - acr - credential - helper
      containers :
        - name : aliyun - acr - credential - helper
          image : registry.cn-shanghai.aliyuncs.com/acs-aliyun-acr-credential-helper:v18.10.29.0-1a28f02
          imagePullPolicy : Always
          terminationGracePeriodSeconds : 0

```

1.9 Network management

1.9.1 Networks supported by Alibaba Cloud Container Service for Kubernetes

This topic describes the networks supported by Alibaba Cloud Container Service for Kubernetes.

Container networks

Container Service provides a stable and high-performance container network through its deep integration of the Kubernetes network and Alibaba Cloud Virtual Private Cloud (VPC). Container Service supports the following types of interconnections:

- Pods within a container cluster can access each other.
- A pod can access a service within a container cluster.
- An Elastic Compute Service (ECS) instance can access a service within a container cluster.
- A pod can directly access an ECS instance (*) in the same VPC.
- An ECS instance can directly access a pod (*) in the same VPC.



Note:

The asterisk (*) indicates that you need to set a valid security group rule.

1.9.2 Access services by using Server Load Balancer

This topic describes how to access services by using Alibaba Cloud Server Load Balancer (SLB).

Check the cloud-controller-manager version

If you specify an existing SLB in a cluster that has a cloud-controller-manager component of v1.9.3 or later versions, the system does not process listeners for this SLB by default. You must manually configure listeners for this SLB.

To view the cloud-controller-manager version, run the following command:

```
root @ master # kubectl get po -n kube-system -o yaml
| grep image :| grep cloud-controller-manager | uniq

image : registry-vpc.cn-hangzhou.aliyuncs.com/acs/cloud-controller-manager-amd64:v1.9.3
```

Use a command-line tool

Method 1

1. Create an Nginx application by using a command-line tool.

```
root @ master # kubectl run nginx --image=registry.aliyuncs.com/acs/netdia:latest
root @ master # kubectl get po
```

NAME	READY	STATUS
nginx - 2721357637 - dvwq3	1 / 1	Running
1	6s	

2. Create an SLB service for the Nginx application and specify `type = LoadBalancer` to expose the Nginx service to the Internet.

```
root @ master # kubectl expose deployment nginx --port = 80 --target-port = 80 --type = LoadBalancer
root @ master # kubectl get svc
```

NAME	CLUSTER-IP	EXTERNAL-IP
nginx	172.19.10.209	101.37.192.20
80 : 31891 / TCP	4s	

3. Visit `http://101.37.192.20` in a browser to access your Nginx service.

Method 2

1. Save the following yml code to the `nginx-svc.yml` file:

```
apiVersion : v1
kind : Service
metadata :
  labels :
    run : nginx
```

```

name : nginx - 01
namespace : default
spec :
  ports :
  - port : 80
    protocol : TCP
    targetPort : 80
  selector :
    run : nginx
  type : LoadBalancer

```

2. Run the `kubectl apply -f nginx - svc . yml` command.

```

root @ master # kubectl apply -f nginx - svc . yml
root @ master # kubectl get service
NAME          TYPE          CLUSTER - IP          EXTERNAL - IP
ngi - 01nx    LoadBalancer 172 . 19 . 9 . 243    101 .
37 . 192 . 129 80 : 32325 / TCP    3h

```

3. Visit `http://101.37.192.129` in a browser to access your Nginx service.

Use the Kubernetes dashboard

1. Save the following yml code to the `nginx - svc . yml` file:

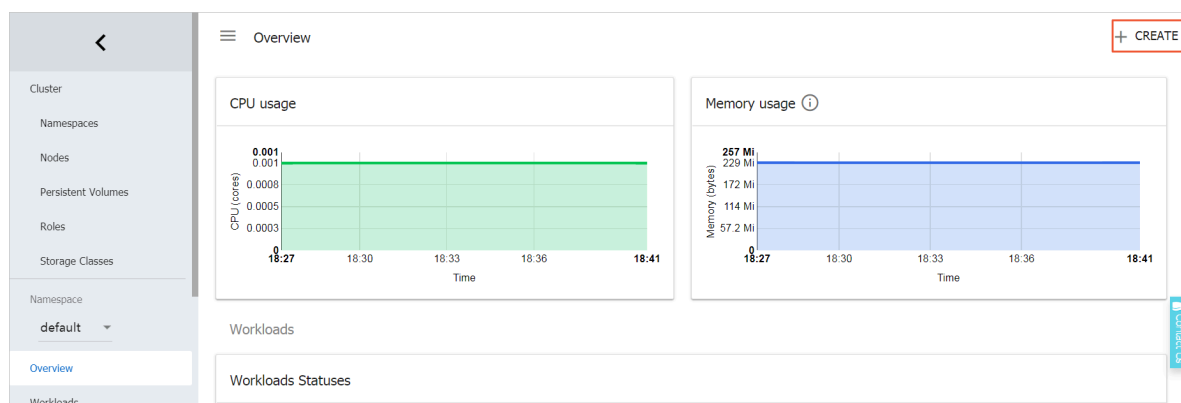
```

apiVersion : v1
kind : Service
metadata :
  labels :
    run : nginx
  name : http - svc
  namespace : default
spec :
  ports :
  - port : 80
    protocol : TCP
    targetPort : 80
  selector :
    run : nginx
  type : LoadBalancer

```

2. Log on to the [Container Service console](#) and click Dashboard on the right of the target cluster.

3. Click **CREATE** in the upper-right corner to create an application.



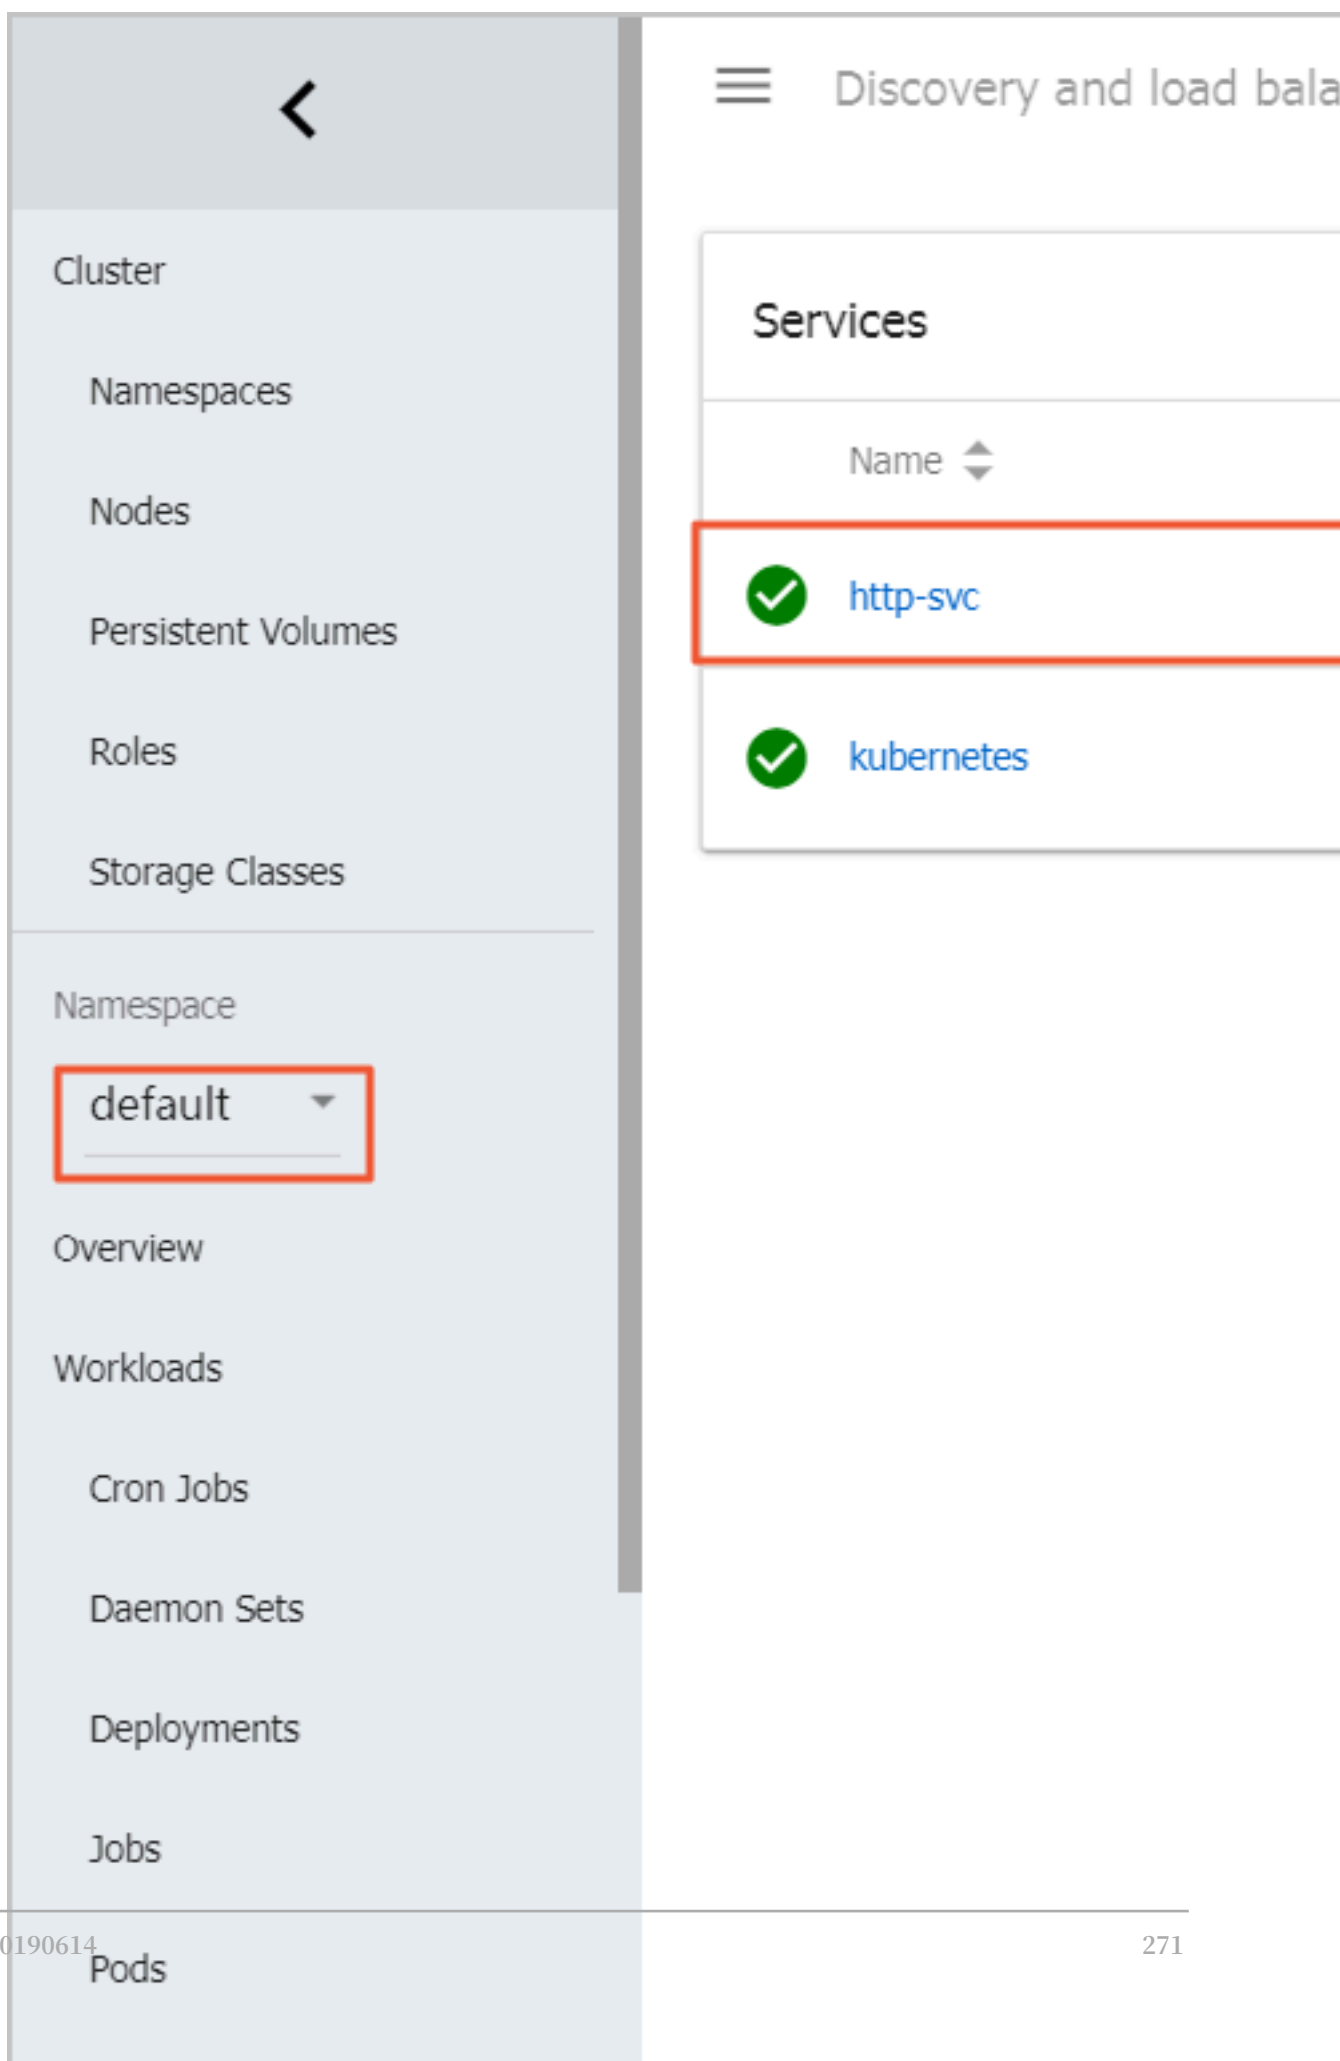
4. Click the **CREATE FROM FILE** tab. Select the `nginx - svc . yml` file you saved.

5. Click **UPLOAD**.

An SLB instance that points to the created Nginx application is created. The service name is `http - svc`.

6. In the left-side navigation pane on the dashboard page, select the default namespace, and then click `Services` .

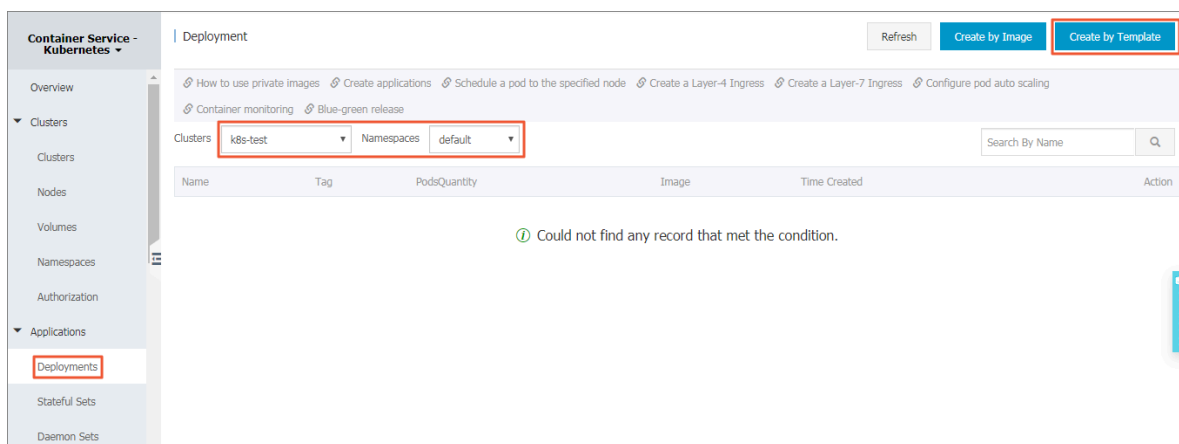
You can view the created Nginx service `http - svc` and the SLB address `http:// 114 . 55 . 79 . 24 : 80` .



7. Open this address in your browser to access the service.

Use the Container Service console

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Kubernetes, choose Applications > Deployments.
3. Select the target cluster and namespace, and then click Create by Template in the upper-right corner.

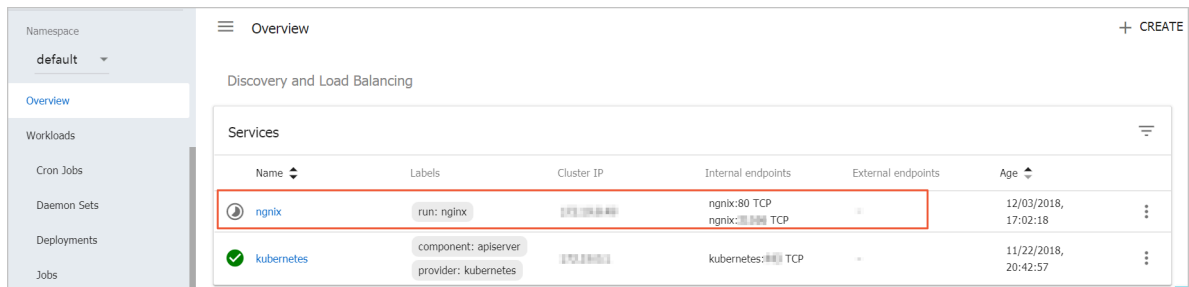


4. Select the custom Resource Type and then copy the following code to the Template.

```
apiVersion : v1
kind : Service
metadata :
  labels :
    run : nginx
  name : nginx
  namespace : default
spec :
  ports :
    - port : 80
      protocol : TCP
      targetPort : 80
  selector :
    run : nginx
  type : LoadBalancer
```

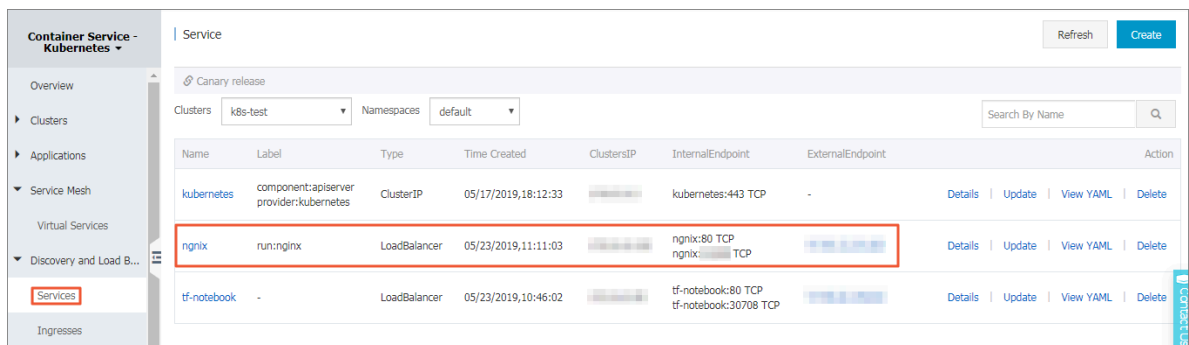
5. Click DEPLOY.

6. Click Kubernetes Dashboard to check the deployment progress on the dashboard page.



Name	Labels	Cluster IP	Internal endpoints	External endpoints	Age
nginx	run: nginx	172.17.0.1	nginx:80 TCP nginx:443 TCP		12/03/2018, 17:02:18
kubernetes	component: apiserver provider: kubernetes	10.0.0.1	kubernetes:443 TCP		11/22/2018, 20:42:57

Alternatively, choose **Discovery and Load Balancing > Services** in the left-side navigation pane, and select the target cluster and namespace to view the deployed service.



Name	Label	Type	Time Created	ClustersIP	InternalEndpoint	ExternalEndpoint	Action
kubernetes	component:apiserver provider:kubernetes	ClusterIP	05/17/2019,18:12:33	10.0.0.1	kubernetes:443 TCP		Details Update View YAML Delete
nginx	run:nginx	LoadBalancer	05/23/2019,11:11:03	172.17.0.1	nginx:80 TCP nginx:443 TCP		Details Update View YAML Delete
tf-notebook	-	LoadBalancer	05/23/2019,10:46:02	172.17.0.1	tf-notebook:80 TCP tf-notebook:30708 TCP		Details Update View YAML Delete

More information

Alibaba Cloud SLB also supports a lot of parameters such as health checks, billing methods, and SLB types. For more information, see [SLB configuration parameters](#).

Annotations

Alibaba Cloud supports plenty of SLB features by using annotations.

Use an existing intranet SLB instance

You must specify three annotations. Replace "your-loadbalancer-id" with your SLB instance ID.



Note:

Multiple Kubernetes services can reuse the same SLB instance.

- The SLB instances created by Kubernetes through a service cannot be reused. Otherwise, the reused SLB instances may be removed incidentally. Only the SLB instances manually created in the console or created by calling API can be reused.

- The multiple services that reuse the same SLB instance cannot have the same frontend listening port. Otherwise, port conflicts will occur.
- If you reuse an SLB instance, the listener name and the virtual server group name are used to identify the SLB instance in Kubernetes. We recommend that you do not modify the listener name or virtual server group name.
- You can modify SLB instance names.
- SLB instances cannot be reused across clusters.

```

apiVersion : v1
kind : Service
metadata :
  annotations :
    service.beta.kubernetes.io/alibaba-cloud-loadbalancer-address-type : "intranet"
    service.beta.kubernetes.io/alibaba-cloud-loadbalancer-id : "your-loadbalancer-id"
    service.beta.kubernetes.io/alibaba-cloud-loadbalancer-force-override-listeners : "true"
  labels :
    run : nginx
  name : nginx
  namespace : default
spec :
  ports :
    - port : 80
      protocol : TCP
      targetPort : 80
  selector :
    run : nginx
  sessionAffinity : None
  type : LoadBalancer

```

Create an HTTP-type SLB instance

```

apiVersion : v1
kind : Service
metadata :
  annotations :
    service.beta.kubernetes.io/alibaba-cloud-loadbalancer-protocol-port : "http : 80"
  name : nginx
  namespace : default
spec :
  ports :
    - port : 80
      protocol : TCP
      targetPort : 80
  selector :
    run : nginx
  type : LoadBalancer

```

Create an HTTPS-type SLB instance

You must first create a certificate in the Alibaba Cloud console before creating an HTTPS-type SLB instance by using the following template (the certificate ID is required by the annotations in the template):

```
apiVersion : v1
kind : Service
metadata :
  annotation s :
    service . beta . kubernetes . io / alicloud - loadbalanc er -
    cert - id : " your - cert - id "
    service . beta . kubernetes . io / alicloud - loadbalanc er -
    protocol - port : " https : 443 "
  name : nginx
  namespace : default
spec :
  ports :
    - port : 443
      protocol : TCP
      targetPort : 443
  selector :
    run : nginx
  sessionAff inity : None
  type : LoadBalanc er
```

Limit SLB instance bandwidth

```
apiVersion : v1
kind : Service
metadata :
  annotation s :
    service . beta . kubernetes . io / alicloud - loadbalanc er -
    charge - type : " paybybandw idth "
    service . beta . kubernetes . io / alicloud - loadbalanc er -
    bandwidth : " 100 "
  name : nginx
  namespace : default
spec :
  ports :
    - port : 443
      protocol : TCP
      targetPort : 443
  selector :
    run : nginx
  type : LoadBalanc er
```

Specify the SLB instance specification

```
apiVersion : v1
kind : Service
metadata :
  annotation s :
    service . beta . kubernetes . io / alicloud - loadbalanc er -
spec : " slb . s1 . small "
  name : nginx
  namespace : default
spec :
  ports :
    - port : 443
      protocol : TCP
```

```
targetPort : 443
selector :
  run : nginx
type : LoadBalancer
```

Use an existing SLB instance

By default, listeners are not overridden if you use an existing SLB instance. To forcibly override the existing listeners, set `service.beta.kubernetes.io/alibabacloud-loadbalancer-force-override-listeners` to `true`.



Note:

Multiple Kubernetes services can reuse the same SLB instance.

- The SLB instances created by Kubernetes through a service cannot be reused. Otherwise, the reused SLB instances may be removed incidentally. Only the SLB instances manually created in the console or created by calling API can be reused.
- The multiple services that reuse the same SLB instance cannot have the same frontend listening port. Otherwise, port conflicts will occur.
- If you reuse an SLB instance, the listener name and the virtual server group name are used to identify the SLB instance in Kubernetes. We recommend that you do not modify the listener name or virtual server group name.
- You can modify SLB instance names.
- SLB instances cannot be reused across clusters.

```
apiVersion : v1
kind : Service
metadata :
  annotations :
    service.beta.kubernetes.io/alibabacloud-loadbalancer-id : "your_loadbalancer_id"
  name : nginx
  namespace : default
spec :
  ports :
    - port : 443
      protocol : TCP
      targetPort : 443
  selector :
    run : nginx
  type : LoadBalancer
```

Use an existing SLB instance and forcibly override existing listeners

If you forcibly override the existing listeners, they are removed.



Note:

Multiple Kubernetes services can reuse the same SLB instance.

- The SLB instances created by Kubernetes through a service cannot be reused. Otherwise, the reused SLB instances may be removed incidentally. Only the SLB instances manually created in the console or created by calling API can be reused.
- The multiple services that reuse the same SLB instance cannot have the same frontend listening port. Otherwise, port conflicts will occur.
- If you reuse an SLB instance, the listener name and the virtual server group name are used to identify the SLB instance in Kubernetes. We recommend that you do not modify the listener name or virtual server group name.
- You can modify SLB instance names.
- SLB instances cannot be reused across clusters.

```
apiVersion : v1
kind : Service
metadata :
  annotations :
    service.beta.kubernetes.io/alibaba-cloud-loadbalancer-id : "your_loadbalancer_id"
    service.beta.kubernetes.io/alibaba-cloud-loadbalancer-force-override-listeners : "true"
    name : nginx
    namespace : default
spec :
  ports :
    - port : 443
      protocol : TCP
      targetPort : 443
  selector :
    run : nginx
  type : LoadBalancer
```

Use the Worker node with specified labels as a backend server

Use a comma (,) to separate two labels, for example, K1: V1, K2: V2.

The relationship between multiple labels is `and`.

```
apiVersion : v1
kind : Service
metadata :
  annotations :
    service.beta.kubernetes.io/alibaba-cloud-loadbalancer-backend-label : "failure-domain.beta.kubernetes.io/zone : ap-southeast-5a"
    name : nginx
    namespace : default
spec :
  ports :
    - port : 443
      protocol : TCP
      targetPort : 443
  selector :
```

```
run : nginx
type : LoadBalancer
```

Set the session persistence timeout for a TCP-type SLB instance

The parameter `service.beta.kubernetes.io/alibabacloud-loadbalancer-persistence-timeout` applies only to TCP listeners.

If the SLB instance is configured with multiple TCP listener ports, this parameter setting applies to all the ports by default.

```
apiVersion : v1
kind : Service
metadata :
  annotations :
    service.beta.kubernetes.io/alibabacloud-loadbalancer-persistence-timeout : "1800"
  name : nginx
  namespace : default
spec :
  ports :
    - port : 443
      protocol : TCP
      targetPort : 443
  selector :
    run : nginx
    type : LoadBalancer
```

Set session persistence for HTTP-type and HTTPS-type SLB instances (insert cookie)

Only HTTP-type and HTTPS-type SLB instances support this setting.

If an instance is configured with multiple HTTP or HTTPS listener ports, the session persistence setting applies to all the HTTP or HTTPS listener ports by default.

```
apiVersion : v1
kind : Service
metadata :
  annotations :
    service.beta.kubernetes.io/alibabacloud-loadbalancer-sticky-session : "on"
    service.beta.kubernetes.io/alibabacloud-loadbalancer-sticky-session-type : "insert"
    service.beta.kubernetes.io/alibabacloud-loadbalancer-cookie-timeout : "1800"
  name : nginx
  namespace : default
spec :
  ports :
    - port : 80
      protocol : TCP
      targetPort : 80
  selector :
    run : nginx
```

```
type : LoadBalancer
```

Set session persistence for HTTP-type and HTTPS-type SLB instances (server cookie)

Only HTTP-type and HTTPS-type SLB instances support this setting.

If an instance is configured with multiple HTTP or HTTPS listener ports, the session persistence setting applies to all the HTTP or HTTPS listener ports by default.

```
apiVersion : v1
kind : Service
metadata :
  annotations :
    service.beta.kubernetes.io/alibaba-cloud-loadbalancer-sticky-session : "on"
    service.beta.kubernetes.io/alibaba-cloud-loadbalancer-sticky-session-type : "server"
    service.beta.kubernetes.io/alibaba-cloud-loadbalancer-cookies : "your_cookie"
    service.beta.kubernetes.io/alibaba-cloud-loadbalancer-protocol-port : "http:80"
  name : nginx
  namespace : default
spec :
  ports :
    - port : 80
      protocol : TCP
      targetPort : 80
  selector :
    run : nginx
  type : LoadBalancer
```

Specify the primary and secondary zones when creating an SLB instance

Support for the primary and secondary zones varies according to region, for example, the ap-southeast-5.

Once created, the primary and secondary zones cannot be changed.

```
apiVersion : v1
kind : Service
metadata :
  annotations :
    service.beta.kubernetes.io/alibaba-cloud-loadbalancer-master-zoneid : "ap-southeast-5a"
    service.beta.kubernetes.io/alibaba-cloud-loadbalancer-slave-zoneid : "ap-southeast-5a"
  name : nginx
  namespace : default
spec :
  ports :
    - port : 80
      protocol : TCP
      targetPort : 80
  selector :
    run : nginx
```

```
type : LoadBalancer
```

Use the node where the pod is located as a backend server

```
apiVersion : v1
kind : Service
metadata :
  name : nginx
  namespace : default
spec :
  externalTrafficPolicy : Local
  ports :
  - port : 80
    protocol : TCP
    targetPort : 80
  selector :
    run : nginx
  type : LoadBalancer
```






Note:


The annotations are case sensitive.

Annotation	Description	Default value
service.beta.kubernetes.io/alibaba-cloud-loadbalancer-protocol-port	Use a comma (,) to separate two values, for example, https:443,http:80.	None
service.beta.kubernetes.io/alibaba-cloud-loadbalancer-address-type	Valid values: internet or intranet.	internet
service.beta.kubernetes.io/alibaba-cloud-loadbalancer-slb-network-type	SLB instance network type. Valid values: classic or vpc.	classic
service.beta.kubernetes.io/alibaba-cloud-loadbalancer-charge-type	Valid values: paybytraffic or paybybandwidth.	paybytraffic


Annotation	Description	Default value
<code>service.beta.kubernetes.io/alibabacloud-loadbalancer-id</code>	SLB instance ID. You can specify an existing SLB instance by using <code>service.beta.kubernetes.io/alibabacloud-loadbalancer-id</code> , and existing listeners will be overridden. Note that the SLB instance will not be deleted if you delete the service.	None
<code>service.beta.kubernetes.io/alibabacloud-loadbalancer-backend-label</code>	Use labels to specify the Worker nodes to be mounted to the backend of the SLB instance.	None
<code>service.beta.kubernetes.io/alibabacloud-loadbalancer-spec</code>	SLB instance specification. For more information, see #unique_115 .	None
<code>service.beta.kubernetes.io/alibabacloud-loadbalancer-persistence-timeout</code>	Session persistence timeout (in seconds). This parameter setting applies only to TCP listeners and the value can be 0 to 3600. The default value is 0, indicating that the session remains disabled. For more information, see #unique_116 .	0

Annotation	Description	Default value
<code>service . beta . kubernetes . io / alicloud - loadbalanc er - sticky - session</code>	<p>Whether to enable session persistence. Valid value: on off .</p> <div>  Note: It applies only to HTTP and HTTPS listeners. </div> <p>For more information, see #unique_117 and #unique_118.</p>	off
<code>service . beta . kubernetes . io / alicloud - loadbalanc er - sticky - session - type</code>	<p>Method used to handle the cookie. Valid values:</p> <ul style="list-style-type: none"> insert : Insert the cookie. server : Rewrite the cookie. <div>  Note: <ul style="list-style-type: none"> It applies only to HTTP and HTTPS listeners. If you set the value of the parameter <code>service . beta . kubernetes . io / alicloud - loadbalanc er - sticky - session to on</code> , you must specify this parameter. </div> <p>For more information, see #unique_117 and #unique_118.</p>	None

Annotation	Description	Default value
<code>service.beta.kubernetes.io/alibaba-cloud-loadbalancer-cookie-timeout</code>	<p>Cookie timeout period (in seconds). Value range: 1 to 86400 .</p> <div> Note: If the <code>service.beta.kubernetes.io/alibaba-cloud-loadbalancer-sticky-session</code> parameter is set to <code>on</code> and the <code>service.beta.kubernetes.io/alibaba-cloud-loadbalancer-sticky-session-type</code> parameter is set to <code>insert</code>, this parameter is mandatory.</div> <p>For more information, see #unique_117 and #unique_118.</p>	None


Annotation	Description	Default value
<code>service . beta . kubernetes . io / alicloud - loadbalanc er - cookie</code>	<p>Cookie configured on the server.</p> <p>The cookie must be a string of 1 to 200 characters and can only contain ASCII letters and numeric characters. It cannot contain commas (,), semicolons (;), or spaces, and it cannot start with a dollar sign (\$).</p> <div>  Note: <p>If the <code>service . beta . kubernetes . io / alicloud - loadbalanc er - sticky - session</code> parameter is set to <code>on</code> and the <code>service . beta . kubernetes . io / alicloud - loadbalanc er - sticky - session - type</code> parameter is set to <code>server</code>, this parameter is mandatory.</p> <p>For more information, see #unique_117 and #unique_118.</p> </div>	None

Annotation	Description	Default value
<code>service.beta.kubernetes.io/alibabacloud-loadbalancer-master-zoneid</code>	Zone ID of the primary backend server.	None
<code>service.beta.kubernetes.io/alibabacloud-loadbalancer-slave-zoneid</code>	Zone ID of the secondary backend server.	None
<code>externalTrafficPolicy</code>	Nodes that can be used as backend servers. Valid values: <ul style="list-style-type: none"> Cluster : Use all backend nodes as backend servers. Local : Use the nodes where pods are located as backend servers. 	Cluster
<code>service.beta.kubernetes.io/alibabacloud-loadbalancer-force-override-listeners</code>	Determines whether to override the listeners when you specify an existing SLB instance.	false : Do not override.
<code>service.beta.kubernetes.io/alibabacloud-loadbalancer-region</code>	Region where the SLB instance is located.	None
<code>service.beta.kubernetes.io/alibabacloud-loadbalancer-bandwidth</code>	SLB instance bandwidth.	50
<code>service.beta.kubernetes.io/alibabacloud-loadbalancer-cert-id</code>	ID of a certificate on Alibaba Cloud. You must upload a certificate first.	None

Annotation	Description	Default value
<code>service.beta.kubernetes.io/alibabacloud-loadbalancer-health-check-flag</code>	Valid values: <code>on</code> <code>off</code> .	The default value is <code>off</code> . Modifying this parameter is not required for TCP, because the health check function is enabled for TCP by default and this parameter cannot be set.
<code>service.beta.kubernetes.io/alibabacloud-loadbalancer-health-check-type</code>	Health check type. Valid values: <code>tcp</code> <code>http</code> . For more information, see #unique_116 .	<code>tcp</code>
<code>service.beta.kubernetes.io/alibabacloud-loadbalancer-health-check-uri</code>	URI used for health checks.  Note: If the health check type is TCP, you do not need to set this parameter. For more information, see #unique_116 .	None
<code>service.beta.kubernetes.io/alibabacloud-loadbalancer-health-check-connect-port</code>	Port used for health checks. Valid values: <ul style="list-style-type: none"><code>520</code> : The backend port configured for the listener is used by default.<code>1 - 65535</code> : The port opened on the backend server for health checks is used. For more information, see #unique_116 .	None

Annotation	Description	Default value
<code>service.beta.kubernetes.io/alibabacloud-loadbalancer-healthy-threshold</code>	For more information, see #unique_116 .	None
<code>service.beta.kubernetes.io/alibabacloud-loadbalancer-unhealthy-threshold</code>	The number of consecutive health check successes before the backend server is determined healthy (from failure to success). Value range: 2 to 10 For more information, see #unique_116 .	None
<code>service.beta.kubernetes.io/alibabacloud-loadbalancer-health-check-interval</code>	Time interval between two consecutive health checks (seconds). Value range: 1 to 50. For more information, see #unique_116 .	None

Annotation	Description	Default value
<code>service.beta.kubernetes.io/alibabacloud-loadbalancer-health-check-connect-timeout</code>	<p>Time period required by waiting for a health check response (in seconds). If the backend ECS instance does not send a valid response within a specified period of time, the system determines that the health check has failed.</p> <p>Value range: 1 to 300 .</p> <div>  Note: If the value of the parameter <code>service.beta.kubernetes.io/alibabacloud-loadbalancer-health-check-connect-timeout</code> is less than the value of the parameter <code>service.beta.kubernetes.io/alibabacloud-loadbalancer-health-check-interval</code>, <code>service.beta.kubernetes.io/alibabacloud-loadbalancer-health-check-connect-timeout</code> is invalid and the timeout period equals the value of <code>service</code> </div>	None

Annotation	Description	Default value
<code>service.beta.kubernetes.io/alibaba-cloud-loadbalancer-health-check-timeout</code>	<p>Time period required by waiting for a health check response (in seconds). If the backend ECS instance does not send a valid response within a specified period of time, the system determines that the health check has failed.</p> <p>Value range: 1 to 300 .</p> <div>  Note: If the value of the parameter <code>service.beta.kubernetes.io/alibaba-cloud-loadbalancer-health-check-timeout</code> is less than that of the parameter <code>service.beta.kubernetes.io/alibaba-cloud-loadbalancer-health-check-interval</code>, <code>service.beta.kubernetes.io/alibaba-cloud-loadbalancer-health-check-timeout</code> is invalid, and the timeout period equals the value of the parameter <code>service.beta.kubernetes.io/alibaba-cloud-loadbalancer-health-check-interval</code>. </div>	None
Issue: 20190614	For more information, see #unique_117 .	289

1.9.3 Support for Ingress

In Kubernetes clusters, Ingress is a collection of rules that authorize inbound connection to the cluster services and provides you with Layer-7 Server Load Balancer capabilities. You can provide the Ingress configuration with externally accessible URL, Server Load Balancer, SSL, and name-based virtual host.

Prerequisites

To test the complex routing service, create an Nginx application in this example. You must create the Nginx deployment and multiple services in advance to observe the routing effect. Replace with your own service in the actual test. In the actual test enter your own service.

```
root @ master # kubectl run nginx -- image = registry . cn -
hangzhou . aliyuncs . com / acs / netdia : latest

root @ master # kubectl expose deploy nginx -- name = http -
svc -- port = 80 -- target - port = 80
root @ master # kubectl expose deploy nginx -- name = http -
svc1 -- port = 80 -- target - port = 80
root @ master # kubectl expose deploy nginx -- name = http -
svc2 -- port = 80 -- target - port = 80
root @ master # kubectl expose deploy nginx -- name = http -
svc3 -- port = 80 -- target - port = 80
```

Simple routing service

Create a simple Ingress service by using the following commands. All the accesses to the `/ svc` path are routed to the Nginx service. `nginx . ingress . kubernetes . io / rewrite - target : /` redirects the path `/ svc` to the path `/` that can be recognized by backend services.

```
root @ master # cat << EOF | kubectl create - f -
apiVersion : extensions / v1beta1
kind : Ingress
metadata :
  name : simple
  annotation s :
    nginx . ingress . kubernetes . io / rewrite - target : /
spec :
  rules :
  - http :
    paths :
    - path : / svc
      backend :
        serviceNam e : http - svc
        servicePor t : 80
EOF
root @ master # kubectl get ing
NAME HOSTS ADDRESS PORTS AGE
```

```
simple      *      101 . 37 . 192 . 211      80
lls
```

Now visit `http://101.37.192.211/svc` to access the Nginx service.

Simple fanout routing based on domain names

If you have multiple domain names providing different external services, you can generate the following configuration to implement a simple fanout effect based on domain names:

```
root @ master # cat << EOF | kubectl create -f -
apiVersion : extensions / v1beta1
kind : Ingress
metadata :
  name : simple - fanout
spec :
  rules :
    - host : foo . bar . com
      http :
        paths :
          - path : / foo
            backend :
              serviceNam e : http - svc1
              servicePor t : 80
          - path : / bar
            backend :
              serviceNam e : http - svc2
              servicePor t : 80
    - host : foo . example . com
      http :
        paths :
          - path : / film
            backend :
              serviceNam e : http - svc3
              servicePor t : 80
EOF
root @ master # kubectl get ing
NAME          HOSTS          ADDRESS          PORTS          AGE
simple - fanout *      101 . 37 . 192 . 211      80
lls
```

Then, you can access the `http - svc1` service by using `http://foo.bar.com/foo`, access the `http - svc2` service by using `http://foo.bar.com/bar`, and access the `http - svc3` service by using `http://foo.example.com/film`.



Note:

- In a production environment, point the domain name to the preceding returned address `101.37.192.211`.

- - In a testing environment, you can modify the `hosts` file to add a domain name mapping rule.

```
101 . 37 . 192 . 211    foo . bar . com
101 . 37 . 192 . 211    foo . example . com
```

Default domain name of simple routing

It does not matter if you do not have the domain name address. Container Service binds a default domain name for Ingress service. You can use this default domain name to access the services. The domain name is in the format of `*.[cluster - id].[region - id].alicontainer.com`. You can obtain the address on the cluster Basic Information page in the console.

Use the following configuration to expose two services with the default domain name.

```
root @ master # cat << EOF | kubectl create -f -
apiVersion : extensions / v1beta1
kind : Ingress
metadata :
  name : shared - dns
spec :
  rules :
    - host : foo .[ cluster - id ].[ region - id ]. alicontainer .
      com ## Replace with the default service access domain
        name of your cluster .
      http :
        paths :
          - path : /
            backend :
              serviceName : http - svc1
              servicePort : 80
    - host : bar .[ cluster - id ].[ region - id ]. alicontainer .
      com ## Replace with the default service access domain
        name of your cluster .
      http :
        paths :
          - path : /
            backend :
              serviceName : http - svc2
              servicePort : 80
EOF
root @ master # kubectl get ing
NAME          HOSTS          ADDRESS          PORTS          AGE
shared - dns   foo .[ cluster - id ].[ region - id ]. alicontainer .
er . com , bar .[ cluster - id ].[ region - id ]. alicontainer .
com           47 . 95 . 160 . 171      80              40m
```

Then, you can access the `http - svc1` service by using `http :// foo .[cluster - id].[region - id]. alicontainer . com` /and access the `http - svc2` service by using `http :// bar .[cluster - id].[region - id]. alicontainer . com` .

Configure a safe routing service

Management of multiple certificates is supported to provide security protection for your services.

1. Prepare your service certificate.

If no certificate is available, generate a test certificate in the following method:



Note:

The domain name must be consistent with your Ingress configuration.

```
root @ master # openssl req - x509 - nodes - days 365
- newkey rsa : 2048 - keyout tls . key - out tls . crt -
subj "/ CN = foo . bar . com / O = foo . bar . com "
```

The above command generates a certificate file `tls . crt` and a private key file `tls . key`.

Create a Kubernetes secret named `foo . bar` using the certificate and private key. The secret must be referenced when you create the Ingress.

```
root @ master # kubectl create secret tls foo . bar --
key tls . key -- cert tls . crt
```

2. Create a safe Ingress service.

```
root @ master # cat << EOF | kubectl create - f -
apiVersion : extensions / v1beta1
kind : Ingress
metadata :
  name : tls - fanout
spec :
  tls :
  - hosts :
    - foo . bar . com
    secretName : foo . bar
  rules :
  - host : foo . bar . com
    http :
      paths :
      - path : / foo
        backend :
          serviceNam e : http - svc1
          servicePor t : 80
      - path : / bar
        backend :
          serviceNam e : http - svc2
          servicePor t : 80
EOF
root @ master # kubectl get ing
NAME HOSTS ADDRESS PORTS
AGE
```

```
tls - fanout      *      101 . 37 . 192 . 211      80
11s
```

3. Follow the notes in Simple fanout routing based on domain names to configure the `hosts` file or set the domain name to access the TLS service.

You can access the `http - svc1` service by using `http :// foo . bar . com / foo` and access the `http - svc2` service by using `http :// foo . bar . com / bar`.

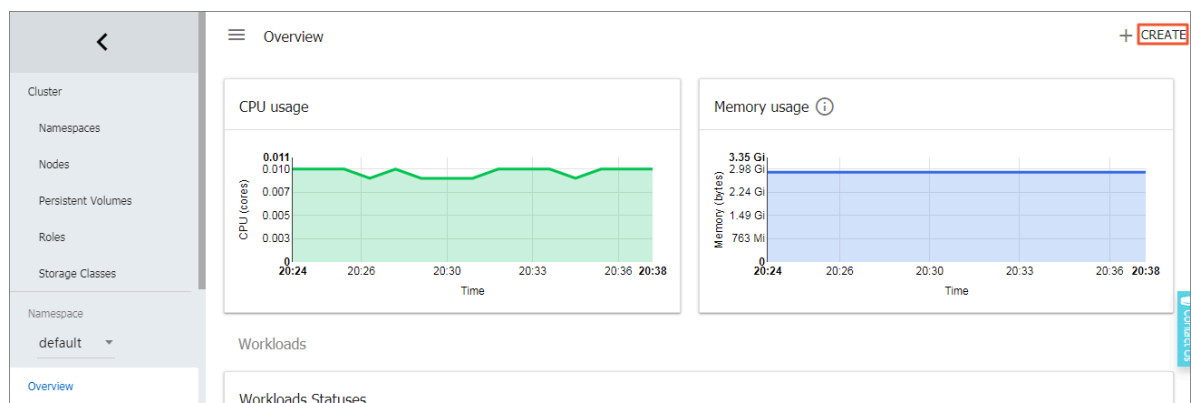
You can also access the HTTPS service by using HTTP. By default, Ingress redirects HTTP access configured with HTTPS to the HTTPS address. Therefore, access to `http :// foo . bar . com / foo` will be automatically redirected to `https :// foo . bar . com / foo`.

Deploy Ingress in Kubernetes dashboard

1. Save the following yml code to the `nginx - ingress . yml` file.

```
apiVersion : extensions / v1beta1
kind : Ingress
metadata :
  name : simple
spec :
  rules :
  - http :
      paths :
      - path : / svc
        backend :
          serviceNam e : http - svc
          servicePor t : 80
```

2. Log on to the [Container Service console](#). In the left-side navigation pane under Kubernetes, click Clusters. Then click Dashboard on the right of the target cluster.
3. Click CREATE in the upper-right corner to create an application.



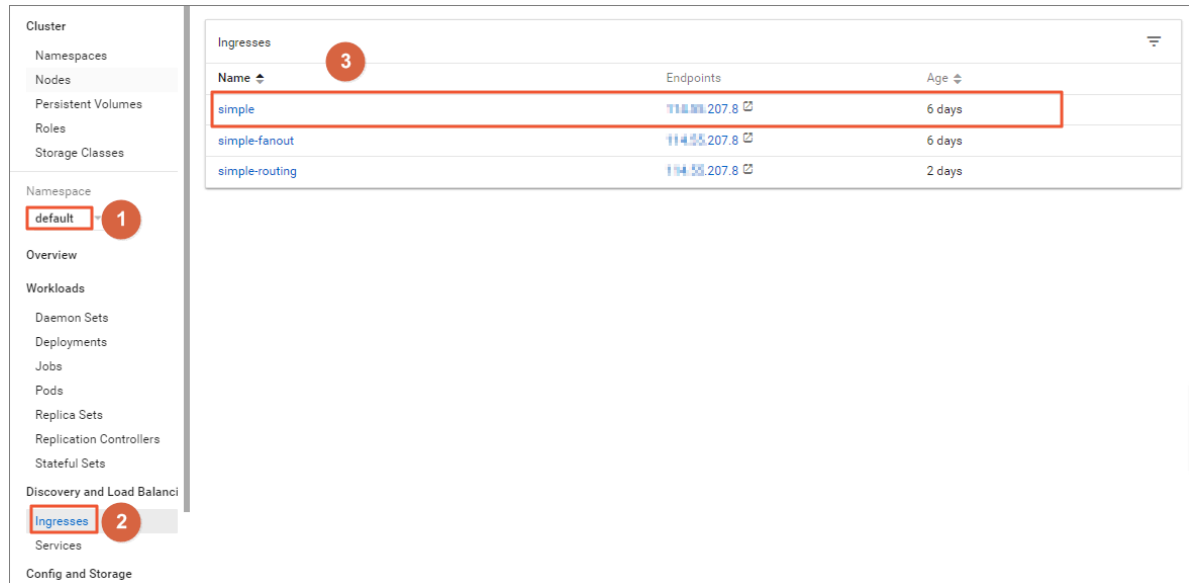
4. Click the CREATE FROM FILE tab. Select the `nginx - ingress . yml` file you saved.

5. Click UPLOAD.

Then an Ingress Layer-7 proxy route will be created to the `http - svc` service.

6. Click default under Namespace in the left-side navigation pane. Click Ingresses in the left-side navigation pane.

You can view the created Ingress resource and its access address `http :// 118 . 178 . 174 . 161 / svc .`



7. Enter the address in the browser to access the created `http - svc` service.

1.9.4 Configure Ingress monitoring

You can view the Ingress monitoring data by enabling the default VTS module of Ingress.

Enable VTS module by running commands

1. Modify the Ingress ConfigMap configuration to add the configuration item `enable`

```
- vts - status : " true ".
```

```
root @ master # kubectl edit configmap nginx - configurat
ion - n kube - system
configmap " nginx - configurat ion " edited
```

After the modification, the contents of the Ingress ConfigMap are as follows:

```
apiVersion : v1
data :
  enable - vts - status : " true " # Enable VTS module
  proxy - body - size : 20m
kind : ConfigMap
metadata :
  Annotation s :
```

```
kubectl . kubernetes . io / last - applied - configurat ion :
|
| {" apiVersion ":" v1 "," data ":{ " proxy - body - size ":"
20m "," kind ":" ConfigMap "," metadata ":{ " annotation s ":"
{}," labels ":{ " app ":" ingress - nginx "," name ":" nginx -
configurat ion "," namespace ":" kube - system "}}
creationTimestamp : 2018 - 03 - 20T07 : 10 : 18Z
labels :
  app : ingress - nginx
  name : nginx - configurat ion
  namespace : kube - system
  selfLink : / api / v1 / namespaces / kube - system / configmaps /
nginx - configurat ion
```

2. Verify if Ingress Nginx has enabled the VTS module normally.

```
root @ master # kubectl get pods -- selector = app = ingress
- nginx - n kube - system
NAME READY STATUS RESTARTS AGE
nginx - ingress - controller - 79877595c8 - 78gq8 1 / 1 Running
0 1h
root @ master # kubectl exec - it nginx - ingress -
controller - 79877595c8 - 78gq8 - n kube - system -- cat
/etc / nginx / nginx . conf | grep vhost_traf fic_status
_display
vhost_traf fic_status _display ;
vhost_traf fic_status _display_f ormat html ;
```

3. Locally access the Ingress Nginx monitoring console.



Note:

By default, the VTS port is not opened for security considerations. Here use the port-forward method to access the console.

```
root @ master # kubectl port - forward nginx - ingress -
controller - 79877595c8 - 78gq8 - n kube - system 18080
Forwarding from 127 . 0 . 0 . 1 : 18080 -> 18080
```

Handling connection for 18080

- Use `http://localhost:18080/nginx_status` to access the VTS monitoring console.

Nginx Vhost Traffic Status

Server main

Host	Version	Uptime	Connections				Requests				Shared memory			
			active	reading	writing	waiting	accepted	handled	Total	Req/s	name	maxSize	usedSize	usedNode
nginx-ingress-controller-79877595c8-78gq8	1.13.7	32m 41s	7	0	1	6	93566	93566	1428	1	vhost_traffic_status	10.0 MiB	2.4 KiB	1

Server zones

Zone	Requests			Responses					Traffic					Cache										
	Total	Req/s	Time	1xx	2xx	3xx	4xx	5xx	Total	Sent	Rcvd	Sent/s	Rcvd/s	Miss	Bypass	Expired	Stale	Updating	Revalidated	Hit	Scarce	Total		
-	660	1	0ms	0	660	0	0	0	660	1.7 MiB	145.4 KiB	1.1 KiB	503 B	0	0	0	0	0	0	0	0	0		
*	660	1	0ms	0	660	0	0	0	660	1.7 MiB	145.4 KiB	1.1 KiB	503 B	0	0	0	0	0	0	0	0	0		

Upstreams

upstream-default-backend

Server	State	Response Time	Weight	MaxFails	FailTimeout	Requests			Responses						Traffic				
						Total	Req/s	Time	1xx	2xx	3xx	4xx	5xx	Total	Sent	Rcvd	Sent/s	Rcvd/s	
172.16.3.6:8080	up	0ms	1	0	0	0	0	0ms	0	0	0	0	0	0	0	0B	0B	0B	0B

update interval: 1 sec

[JSON](#) | [GITHUB](#)

Enable VTS module by using the Kubernetes dashboard

- Log on to the [Container Service console](#).
- On the Cluster List page of Kubernetes clusters, click Dashboard at the right of a cluster to enter the Kubernetes dashboard page.
- Select kube-system under Namespace in the left-side navigation pane. Click Config Maps in the left-side navigation pane. Click the icon at the right of nginx-configuration and then select View/edit YAML. Edit the config map to add the configuration item `enable - vts - status : " true "`.

The contents of the saved Ingress ConfigMap are as follows:

```
{
  " kind ": " ConfigMap ",
  " apiVersion ": " v1 ",
  " metadata ": {
    " name ": " nginx - configurat ion ",
    " namespace ": " kube - system ",
    " selfLink ": "/ api / v1 / namespaces / kube - system /
configmaps / nginx - configurat ion ",
    " creationTi mestamp ": " 2018 - 03 - 20T07 : 10 : 18Z ",
    " labels ": {
      " app ": " ingress - nginx "
    },
    " annotation s ": {
      " kubectl . kubernetes . io / last - applied - configurat
ion ": "{ \" apiVersion \": \" v1 \", \" data \": { \" proxy - body -
size \": \" 20m \", \" kind \": \" ConfigMap \", \" metadata \": { \"
annotation s \": { }, \" labels \": { \" app \": \" ingress - nginx \",
\" name \": \" nginx - configurat ion \", \" namespace \": \" kube -
system \" } } } \" "

```

```

    }
  },
  "data": {
    "proxy - body - size": " 20m ",
    "enable - vts - status": " true "
  }
}

```

4. Locally access the Ingress Nginx monitoring console.



Note:

By default, the VTS port is not opened for security considerations. Here use the port-forward method to access the console.

```

root @ master # kubectl port - forward nginx - ingress -
controller - 79877595c8 - 78gq8 - n kube - system 18080
Forwarding from 127 . 0 . 0 . 1 : 18080 -> 18080
Handling connection for 18080

```

5. Use `http://localhost:18080/nginx_stat` us to access the VTS monitoring console.

Nginx Vhost Traffic Status

Server main

Host	Version	Uptime	Connections				Requests				Shared memory			
			active	reading	writing	waiting	accepted	handled	Total	Req/s	name	maxSize	usedSize	usedNode
nginx-ingress-controller-79877595c8-78gq8	1.13.7	32m 41s	7	0	1	6	93566	93566	1428	1	vhost_traffic_status	10.0 MiB	2.4 KiB	1

Server zones

Zone	Requests			Responses					Traffic					Cache									
	Total	Req/s	Time	1xx	2xx	3xx	4xx	5xx	Total	Sent	Rcvd	Sent/s	Rcvd/s	Miss	Bypass	Expired	Stale	Updating	Revalidated	Hit	Scarce	Total	
-	660	1	0ms	0	660	0	0	0	660	1.7 MiB	145.4 KiB	1.1 KiB	503 B	0	0	0	0	0	0	0	0	0	
	660	1	0ms	0	660	0	0	0	660	1.7 MiB	145.4 KiB	1.1 KiB	503 B	0	0	0	0	0	0	0	0	0	

Upstreams

upstream-default-backend

Server	State	Response Time	Weight	MaxFails	FailTimeout	Requests		Responses							Traffic					
						Total	Req/s	Time	1xx	2xx	3xx	4xx	5xx	Total	Sent	Rcvd	Sent/s	Rcvd/s		
172.16.3.6:8080	up	0ms	1	0	0	0	0	0	0ms	0	0	0	0	0	0	0	0 B	0 B	0 B	0 B

update interval: sec

[JSON](#) | [GITHUB](#)

1.9.5 Ingress configurations

Alibaba Cloud Container Service provides the highly reliable Ingress controller components and integrates with Alibaba Cloud Server Load Balancer to provide the flexible and reliable Ingress service for your Kubernetes clusters.

See the following Ingress orchestration example. You must configure the annotation when creating an Ingress in the Container Service console. Some configurations must create dependencies. For more information, see [Create an Ingress in the Container Service console](#), [Support for Ingress](#), and [Kubernetes Ingress](#). Ingress also supports the

configuration of configmap. For more information, see <https://kubernetes.github.io/ingress-nginx/user-guide/nginx-configuration/configmap/>.

```

apiVersion : extensions / v1beta1
kind : Ingress
metadata :
  annotation s :
    nginx . ingress . kubernetes . io / service - match : ' new -
nginx : header (" foo ", /^ bar $/)' # Gray release
rule . Header is used in this example .
    nginx . ingress . kubernetes . io / service - weight : ' new -
nginx : 50 , old - nginx : 50 ' # Traffic weight
  annotation
    creationTi mestamp : null
    generation : 1
    name : nginx - ingress
    selfLink : / apis / extensions / v1beta1 / namespaces / default /
ingresses / nginx - ingress
spec :
  rules : ##
Ingress rule
  - host : foo . bar . com
    http :
      paths :
        - backend :
            serviceNam e : new - nginx
            servicePor t : 80
          path : /
        - backend :
            serviceNam e : old - nginx
            servicePor t : 80
          path : /
  tls : ##
Enable TLS to set a secure Ingress .
  - hosts :
    - *. xxxxxx . cn - hangzhou . alicontain er . com
    - foo . bar . com
      secretName : nginx - ingress - secret
## Secret name
status :
  loadBalanc er : {}

```

Annotation

You can configure an ingress annotation, specifying the ingress controller to use, rules for routing, such as routing weight rules, grayscale publish, and rewrite rules. For more information, see <https://kubernetes.github.io/ingress-nginx/user-guide/nginx-configuration/annotations/>.

For example, a typical rewrite annotation `nginx . ingress . kubernetes . io / rewrite - target : /` redirects the path `/ path` to the path `/` that can be recognized by the backend services.

Rules

The rules indicate those that authorize the inbound access to the cluster and are generally the HTTP rules, including the domain name (virtual hostname), URL access path, service name, and port.

You must complete the following configurations for each HTTP rule:

- **Host:** Enter the testing domain name of an Alibaba Cloud Kubernetes cluster or a virtual hostname, such as `foo . bar . com`.
- **Path:** Specify the URL path of the service access. Each path is associated with a backend service. Before Alibaba Cloud Server Load Balancer forwards the traffic to the backend, all inbound requests must match with the domain name and path.
- **Backend configuration:** Service configuration that is a combination of `service : port` and traffic weight. The Ingress traffic is forwarded to the matched backend services based on the traffic weight.
 - **Name:** The name of the backend service forwarded by Ingress.
 - **Port:** The port exposed by the service.
 - **Weight:** The weight rate of each service in a service group.



Note:

1. The service weight is calculated in relative values. For example, if both service weights are set to 50, the weight ratio of both services is 50%.
2. A service group (a service with the same Host and Path in the same ingress yaml) has a default weight value of 100 and the weight is not explicitly set.

Grayscale publish

Container Service supports different traffic segmentation methods for grayscale publish and AB test scenarios.



Note:

Currently, the Alibaba Cloud Container Service Kubernetes Ingress Controller requires `0 . 12 . 0 - 5` and above to support the traffic segmentation feature.

1. Traffic segmentation based on the request header.
2. Traffic segmentation based on cookie.
3. Traffic segmentation based on query (request) parameters.

After the grayscale rule is configured, the request that matches the grayscale publish rule can be routed to the set service. If the service sets a weight rate of less than 100%, requests that match the grayscale publish rule continue to be routed to the corresponding service based on the weight rate.

TLS

You can encrypt the Ingress by specifying a secret that contains the TLS private key and certificate to implement the secure Ingress access. The TLS secret must contain the certificate named `tls.crt` and private key named `tls.key`. For more information about the TLS principles, see [TLS](#). For how to create a secret, see [Configure a safe routing service](#).

Label

You can add tags for Ingress to indicate the characteristics of the Ingress.

1.9.6 Create an Ingress in the Container Service console

Alibaba Cloud Container Service console integrates with the Ingress service, which allows you to quickly create an Ingress service in the Container Service console to build the flexible and reliable traffic access layer.

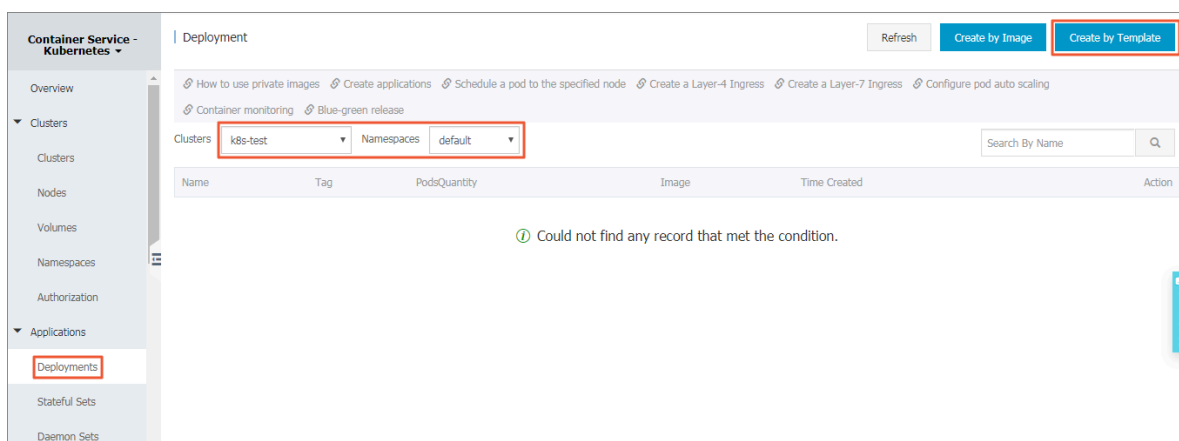
Prerequisites

- You have successfully created a Kubernetes cluster and Ingress controller is running normally in the cluster. For how to create a Kubernetes cluster, see [Create a Kubernetes cluster](#).
- Log on to the master node by using SSH. For more information, see [Access Kubernetes clusters by using SSH](#).

Step 1: Create a deployment and a service

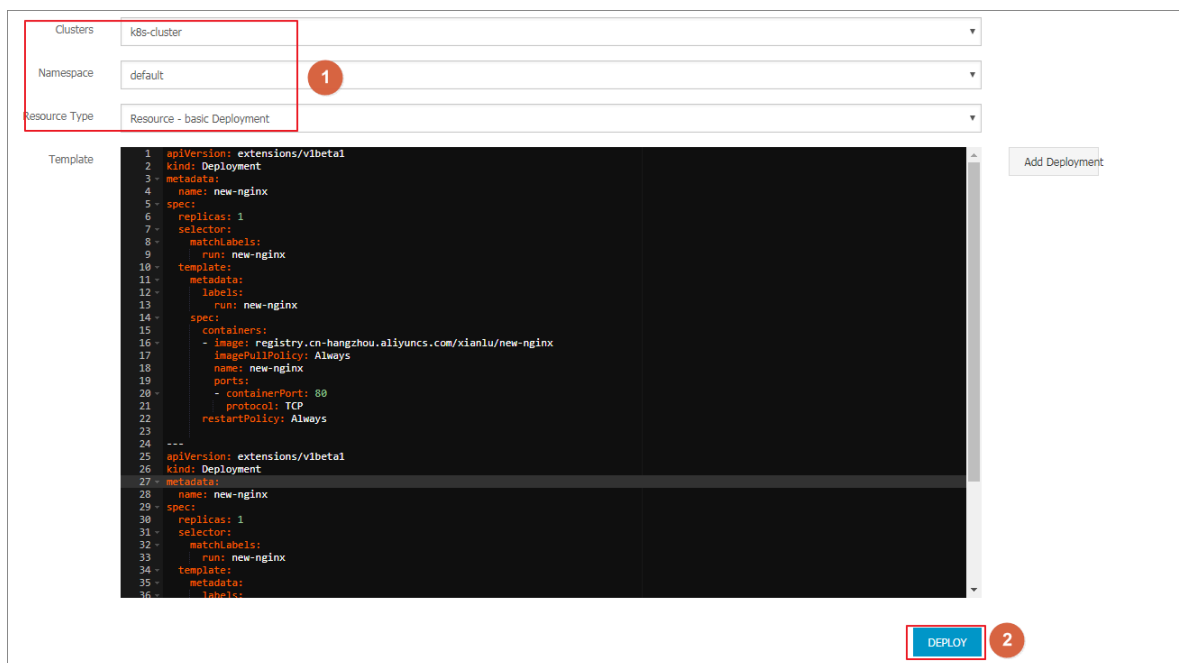
1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Container Service-Kubernetes, choose Applications > Deployments.

3. Click Create by template in the upper-right corner.



4. Select the cluster and namespace from the Clusters and Namespace drop-down lists. Select a sample template or Custom from the Resource Type drop-down list. Click DEPLOY.

In this example, three nginx applications are created. One for the old application (old-nginx), one for the new (new-nginx), and an application for testing the cluster access domain name (domain-nginx).



The orchestration template for old-nginx is as follows:

```
apiVersion : extensions / v1beta1
kind : Deployment
metadata :
  name : old - nginx
spec :
  replicas : 2
  selector :
    matchLabel s :
```

```

        run : old - nginx
    template :
        metadata :
            labels :
                run : old - nginx
        spec :
            containers :
                - image : registry . cn - hangzhou . aliyuncs . com / xianlu
                  / old - nginx
                  imagePullPolicy : Always
                  name : old - nginx
                  ports :
                    - containerPort : 80
                      protocol : TCP
                  restartPolicy : Always
---
apiVersion : v1
kind : Service
metadata :
    name : old - nginx
spec :
    ports :
        - port : 80
          protocol : TCP
          targetPort : 80
    selector :
        run : old - nginx
    sessionAffinity : None
    type : NodePort

```

The orchestration template for new-nginx is as follows:

```

apiVersion : extensions / v1beta1
kind : Deployment
metadata :
    name : new - nginx
spec :
    replicas : 1
    selector :
        matchLabels :
            run : new - nginx
    template :
        metadata :
            labels :
                run : new - nginx
        spec :
            containers :
                - image : registry . cn - hangzhou . aliyuncs . com / xianlu
                  / new - nginx
                  imagePullPolicy : Always
                  name : new - nginx
                  ports :
                    - containerPort : 80
                      protocol : TCP
                  restartPolicy : Always
---
apiVersion : v1
kind : Service
metadata :
    name : new - nginx
spec :
    ports :
        - port : 80

```

```

    protocol : TCP
    targetPort : 80
  selector :
    run : new - nginx
  sessionAffinity : None
  type : NodePort

```

The orchestration template for domain-nginx is as follows:

```

apiVersion : apps / v1beta2 # For versions before 1.8.0
                                use apps / v1beta1
kind : Deployment
metadata :
  name : domain - nginx
  labels :
    app : nginx
spec :
  replicas : 2
  selector :
    matchLabels :
      app : nginx
  template :
    metadata :
      labels :
        app : nginx
    spec :
      containers :
        - name : nginx
          image : nginx : 1.7.9 # replace it with your
                                exactly < image_name : tags >
          ports :
            - containerPort : 80
---
apiVersion : v1
kind : Service
metadata :
  name : domain - nginx
spec :
  ports :
    - port : 80
      protocol : TCP
      targetPort : 80
  selector :
    app : nginx
  sessionAffinity : None

```

```
type : NodePort
```

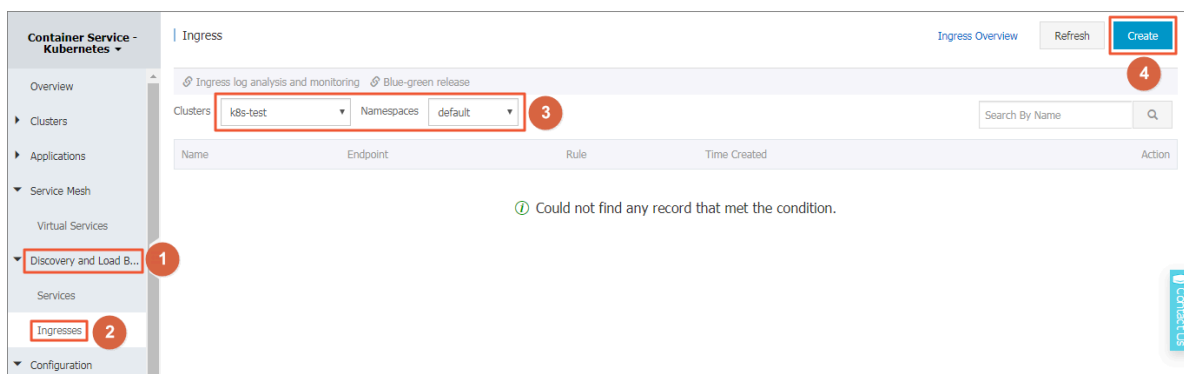
5. In the left-side navigation pane under Container Service-Kubernetes, choose **Discovery and Load Balancing > Services**.

After the service is created, you can see it on the Service List page.

Service List							Refresh	Create
Clusters		k8s-cluster	Namespace		default			
Name	Type	Time Created	ClustersIP	InternalEndpoint	ExternalEndpoint	Action		
domain-nginx	NodePort	07/11/2018,17:43:32		domain-nginx:80 TCP domain-nginx:32347 TCP	-	Details	Update	View YAML Delete
kubernetes	ClusterIP	07/11/2018,17:35:35		kubernetes:443 TCP	-	Details	Update	View YAML Delete
new-nginx	NodePort	07/11/2018,17:37:01		new-nginx:80 TCP new-nginx:32637 TCP	-	Details	Update	View YAML Delete
old-nginx	NodePort	07/11/2018,17:37:01		old-nginx:80 TCP old-nginx:32039 TCP	-	Details	Update	View YAML Delete

Step 2: Create an Ingress

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Container Service-Kubernetes, choose **Discovery and Load Balancing > Ingresses**.
3. Select the cluster and namespace from the Clusters and Namespace drop-down lists. Then click **Create** in the upper-right corner.



4. In the displayed dialog box, enter the Ingress name. In this example, enter **nginx-ingress**.

Name:

5. Configure the rules.

The Ingress rules are the rules that authorize the inbound access to the cluster and are generally the HTTP rules. Configure the domain name (virtual hostname), URL path, service name, and port. For more information, see [Ingress configurations](#).

In this example, add a complicated Ingress rule. Configure the default test domain name and virtual hostname of the cluster to display the Ingress service based on the domain names.

Rule: + Add

Domain ✖

test.[cluster-id].[region-id].alicontainer.com

Select *.alicontainer.com or Custom

path

/

Service + Add

Name	Port	Weight	Percent of Weight
domain-nginx	80	100	100.0%

Domain ✖

foo.bar.com

Select *.alicontainer.com or Custom

path

/

Service + Add

Name	Port	Weight	Percent of Weight
new-nginx	80	50	50.0%
old-nginx	80	50	50.0%

- The simple Ingress based on the default domain name, that is, provide the access service externally by using the default domain name of the cluster.
 - Domain: Enter the default domain name of the cluster. In this example, use `test.[cluster-id].[region-id].alicontainer.com`.

The default domain name of this cluster is displayed in the Create dialog box, in the `*.[cluster - id].[region - id]. alicontain er . com` format. You can also obtain the default domain name on the Basic Information page of the cluster.

- Service: Configure the access path, name, and port of the service.
 - Path: Specify the URL path of the service access. The default is the root path `/`, which is not configured in this example. Each path is associated with a backend service. Before Alibaba Cloud Server Load Balancer forwards the traffic to the backend, all inbound requests must match with the domain name and path.
 - Service configuration: The backend configuration, which is a combination of service name, port, and service weight. The configuration of multiple services in the same access path is supported, and Ingress traffic is split and is forwarded to the matched backend services.
- The simple fanout Ingress based on the domain name. In this example, use a virtual hostname as the testing domain name to provide the access service externally. You can use the recorded domain name in the production environment to provide the access service. You can use the recorded domain name in the production environment to provide the access service.
- Domain: In this example, use the testing domain name `foo . bar . com`.

You must modify the hosts file to add a domain name mapping rule.

```
118 . 178 . 108 . 143   foo . bar . com #   Ingress   IP
address
```

- Service: Configure the access path, name, and port of the service.
 - Path: Specify the URL path of the service access. Path is not configured in this example, and the root path is `/`.
 - Name: In this example, set up both new and old services, `nginx-new` and `nginx-old`.
 - Port: Expose 80 port.
 - Weight settings: Set the weight of multiple services under this path. The service weight is calculated by relative value. The default value is 100. As shown in this example, the service weight values of both the old and new versions are 50, which means that the weight rate of both services is 50%.

6. Grayscale publish configuration.



Note:

Currently, the Alibaba Cloud Container Service Kubernetes Ingress Controller requires 0.12.0-5 and above to support the traffic segmentation feature.

Container Service supports different traffic segmentation methods for grayscale publish and AB test scenarios.

- a. Traffic segmentation based on the request header.
- b. Traffic segmentation based on cookie.
- c. Traffic segmentation based on query (request) parameters.

After the grayscale rule is configured, the request that matches the grayscale publish rule can be routed to the new service version new-nginx. If the service sets a weight rate of less than 100%, requests that match the grayscale publish rule continue to be routed to the corresponding service based on the weight rate.

In this case, set the request header to meet a grayscale publish rule of `foo =^ bar`, only requests with the request header can access the new-nginx service.

Grayscale release:

+ Add After the gray rule is set, the request meeting the rule will set a weight other than 100, the request to satisfy the gamma rule and old version services according to the weights.

Service	Type	Name
new-nginx	Header	foo

- Service: Routing rule configuration service.
- Type: matching request header, cookie, and query (request) parameters are supported.
- Name and match value: User-defined request field, name and match value are key-value pairs.
- Match rules: Regular and exact matches are supported.

7. Configure the annotations.

Click rewrite annotation, a typical redirection annotation can be added to the route. `nginx . ingress . kubernetes . io / rewrite - target : /` indicates that the `/ path` is redirected to the root path `/` that the backend service can recognize.



Note:

In this example, the access path is not configured, so no need to configure rewrite annotations. The purpose of the rewrite annotation is to enable Ingress to forward to the backend as the root path, avoiding 404 errors caused by incorrect access path configuration.

You can also click Add to enter the annotation name and value, which is the annotation key-value pair for Ingress. For more information, see <https://kubernetes.github.io/ingress-nginx/user-guide/nginx-configuration/annotations/>.

annotation:

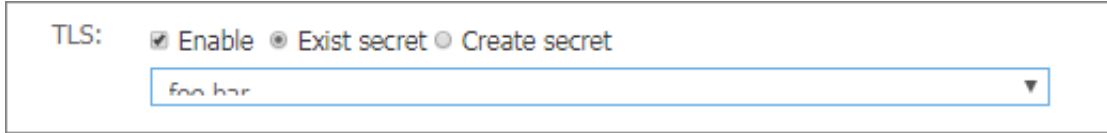
Add

rewrite annotation

Name	Value
<input type="text" value="nginx.ingress.kubernetes.io/rewri"/>	<input type="text" value="/"/>

8. Configure TLS. Select **Enable** and configure the secure Ingress service. For more information, see [Configure a safe routing service](#).

- You can select to use an existing secret.



a. Log on to the master node and create `tls . key` and `tls . crt`.

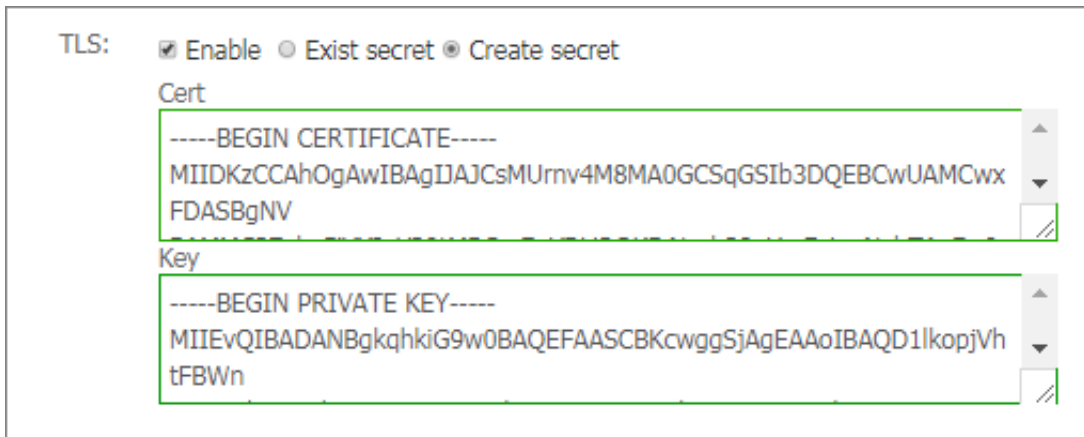
```
openssl req -x509 -nodes -days 365 -newkey rsa :
2048 -keyout tls . key -out tls . crt -subj "/ CN =
foo . bar . com / O = foo . bar . com "
```

b. Create a secret.

```
kubectl create secret tls foo . bar -- key tls .
key -- cert tls . crt
```

c. Run the `kubectl get secret` command to see that secret has been successfully created. You can use the secret that you have created in the Web interface, `foo . bar`.

- You can create the secret with one click by using the created TLS private key and certificate.



a. Log on to the master node and create `tls . key` and `tls . crt`.

```
openssl req -x509 -nodes -days 365 -newkey rsa :
2048 -keyout tls . key -out tls . crt -subj "/ CN =
foo . bar . com / O = foo . bar . com "
```

b. Run the `vim tls . key` and `vim tls . crt` to get the generated private key and certificate.

c. Copy the generated certificate and private key to the Cert and Key fields.

9. Add the tags.

Add the corresponding tags for Ingress to indicate the characteristics of the Ingress.

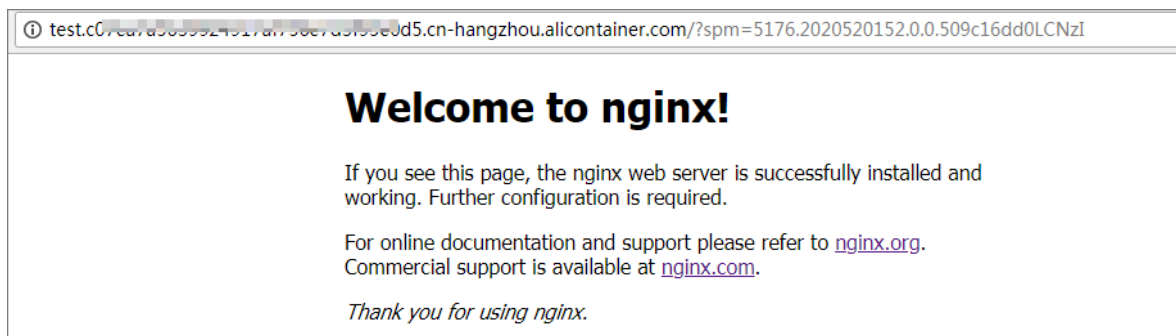
Name	Value
node-role.kubernetes.io/ingress	true

10. Click Create.

The Ingress nginx-ingress is displayed on the Ingress page.

Name	endpoint	Rule	Time Created	Action
nginx-ingress		foo.bar.com/ -> new-nginx foo.bar.com/ -> old-nginx	07/11/2018,17:49:46	Details Update View YAML Delete

11. Click on the access domain name `test.[cluster-id].[region-id].alicontainer.com` in the route, and `foo.bar.com` to access the welcome page of nginx.

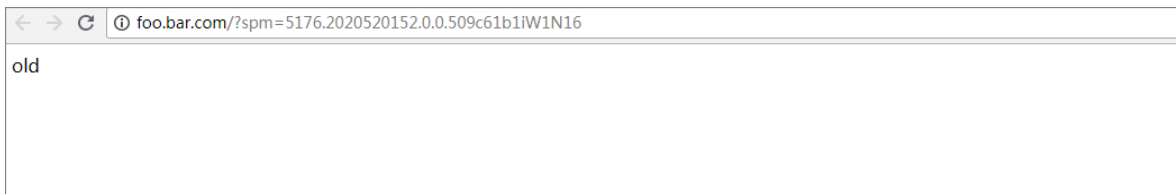


Click on the route address pointing the new-nginx service and find the page that points the old-nginx application.



Note:

Access the route address in the browser. By default, the request header does not have the `foo =^ bar $`, so the traffic is directed to the old-nginx application.



12. Log on to the master node by using SSH. Run the following command to simulate the access result with a specific request header.

```
curl -H "Host : foo . bar . com " http:// 47 . 107 . 20 . 35
old
curl -H "Host : foo . bar . com " http:// 47 . 107 . 20 . 35
old
curl -H "Host : foo . bar . com " http:// 47 . 107 . 20 . 35 # Similar to browser access requests
old
curl -H "Host : foo . bar . com " -H "foo : bar " http:// 47 . 107 . 20 . 35 # Simulate an access request with a unique header , returning results based on routing weight
new
curl -H "Host : foo . bar . com " -H "foo : bar " http:// 47 . 107 . 20 . 35
old
curl -H "Host : foo . bar . com " -H "foo : bar " http:// 47 . 107 . 20 . 35
old
curl -H "Host : foo . bar . com " -H "foo : bar " http:// 47 . 107 . 20 . 35
new
```

1.9.7 View Ingress details

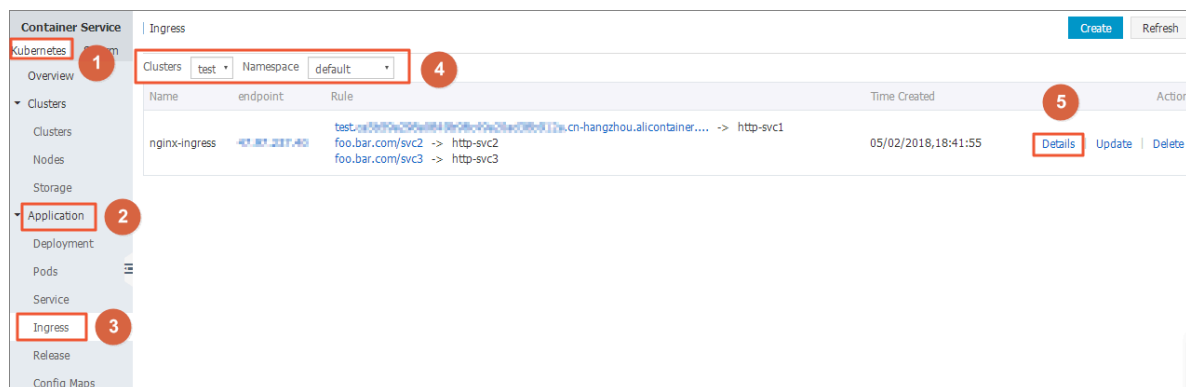
Prerequisites

- You have successfully created a Kubernetes cluster and Ingress controller is running normally in the cluster. For how to create a Kubernetes cluster, see [Create a Kubernetes cluster](#).
- You have successfully created an Ingress. For more information, see [Create an Ingress in the Container Service console](#).

Procedure

1. Log on to the [Container Service console](#).
2. Click Kubernetes Application > Ingress in the left-side navigation pane.

3. Select the cluster and namespace from the Clusters and Namespace drop-down lists. Click Details at the right of the Ingress.



On the details page, you can view the overview and rules of the Ingress.

nginx-ingress

[← Back to List](#)

Refresh

Overview

Name:	nginx-ingress
Namespace:	default
Time Created:	2018-05-02T10:41:55Z
Label:	node-role.kubernetes.io/ingress:true
annotation:	nginx.ingress.kubernetes.io/rewrite-target:/
endpoint:	47.87.237.40

Rule

Domain	path	Name	service port
test.cn-hangzhou.aliyuncs.com	/svc1	http-svc1	80
foo.bar.com	/svc2	http-svc2	80
foo.bar.com	/svc3	http-svc3	80

1.9.8 Update an Ingress

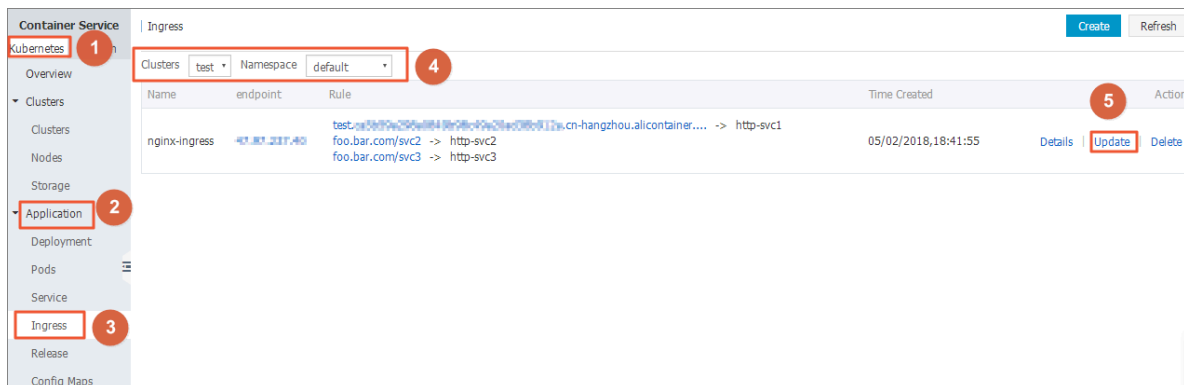
Prerequisites

- You have successfully created a Kubernetes cluster and Ingress controller is running normally in the cluster. For how to create a Kubernetes cluster, see [Create a Kubernetes cluster](#).
- You have successfully created an Ingress. For more information, see [Create an Ingress in the Container Service console](#).

Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Application > Ingress in the left-side navigation pane.

3. Select the cluster and namespace from the Clusters and Namespace drop-down lists. Click Update at the right of the Ingress.



4. Update the Ingress parameters in the displayed dialog box and then click OK.

change `test.[cluster-id].[region-id].alicontainer.com` to `testv2.[cluster-id].[region-id].alicontainer.com`.



What's next

On the Ingress page, you can see a rule of this Ingress is changed.



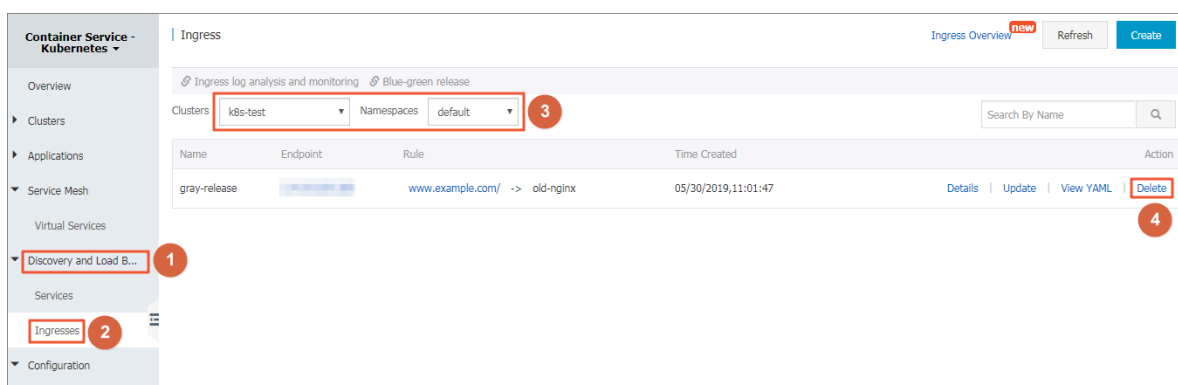
1.9.9 Deleting a route

Prerequisites

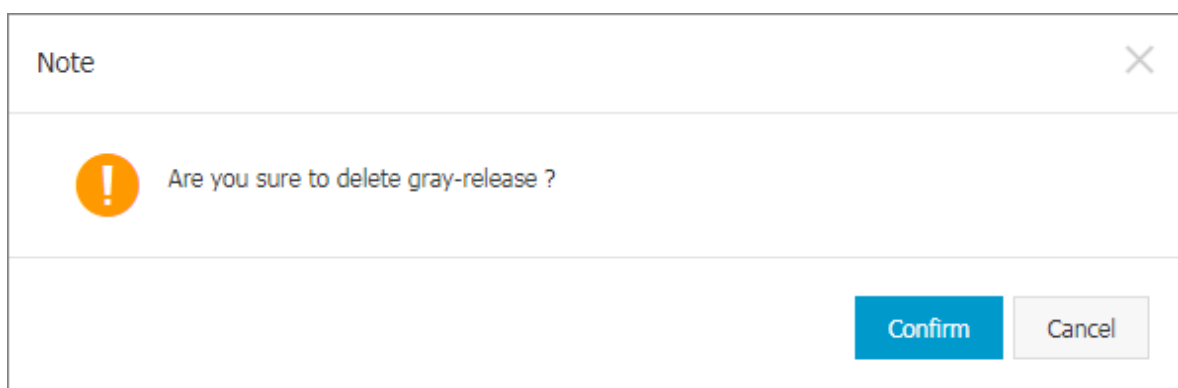
- You have successfully created a Kubernetes cluster and Ingress controller is running normally in the cluster. For how to create a Kubernetes cluster, see [Create a Kubernetes cluster](#).
- You have successfully created an Ingress. For more information, see [Create an Ingress in the Container Service console](#).

Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Application > Ingress in the left-side navigation pane.
3. Select the cluster and namespace from the Clusters and Namespace drop-down lists. Click Delete at the right of the Ingress.



4. Click Confirm in the displayed dialog box.



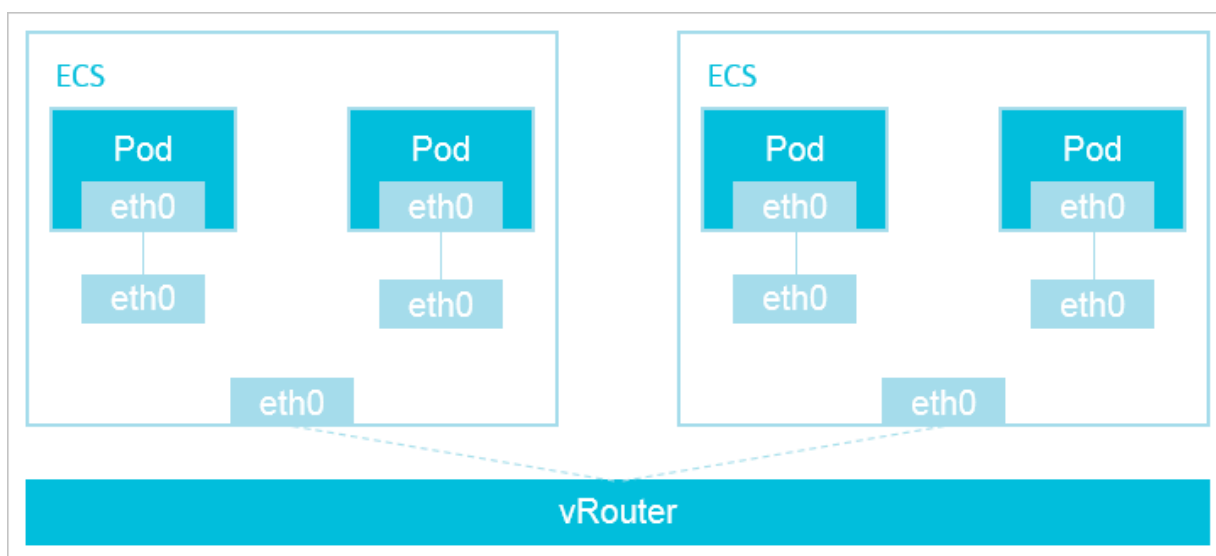
1.9.10 Terway network plugin

This topic describes how to use the Terway network plugin in a Kubernetes cluster that runs on Alibaba Cloud Container Service.

Terway network plugin

Terway, a network plugin developed by Alibaba Cloud Container Service, is fully compatible with Flannel, and provides the following features:

- Allocates Alibaba Cloud Elastic Network Interfaces (ENIs) to containers.
- Defines the access policies for containers according to the Kubernetes Network Policy. This network plugin is also compatible with the Calico Network Policy.

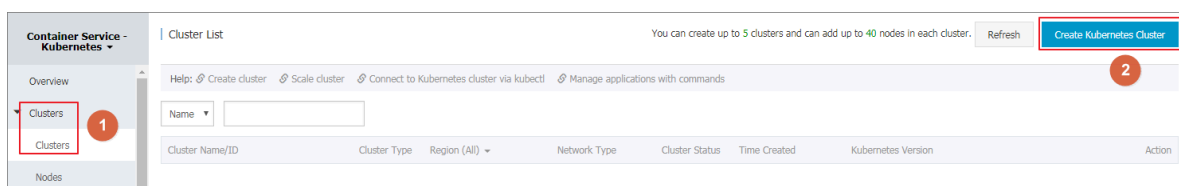


If you install the Terway network plugin in a Kubernetes cluster, each pod then has its own network stack and an IP address. Packets between pods on one ECS instance are forwarded directly by the instance. Packets between pods on different ECS instances are forwarded through the VRouter of a VPC. The Terway network plugin delivers high communication performance because it does not use tunneling technologies such as VXLAN to encapsulate packets.

Use the Terway network plugin

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Kubernetes, click Clusters.

3. In the upper-right corner, click Create Kubernetes Cluster.

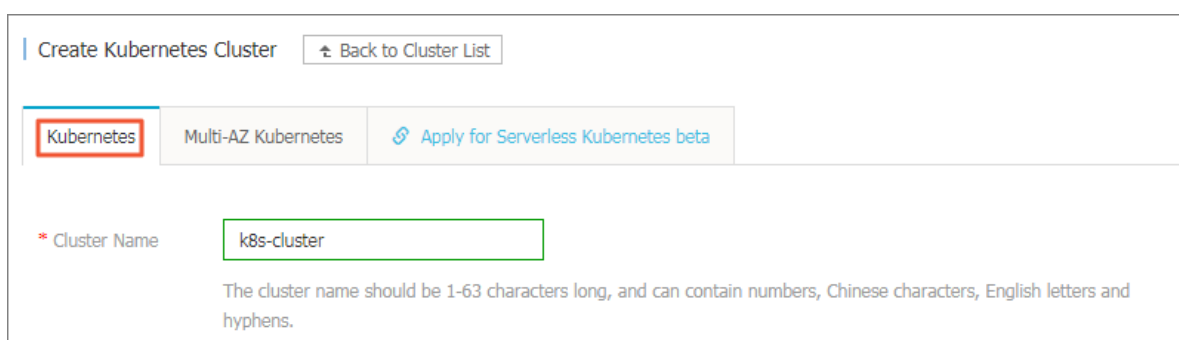


By default, the Create Kubernetes Cluster page is displayed.

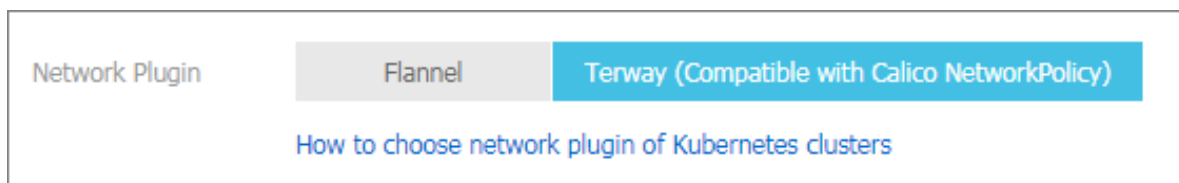


Note:

In this example, a dedicated Kubernetes cluster is created. For more information, see [Create a Kubernetes cluster](#).



4. Select the Terway network plugin.



Flannel and Terway

Alibaba Cloud Container Service for Kubernetes provides two types of network plugins for you to create a Kubernetes cluster: Terway and Flannel.



Note:

For how to select a network plugin, see [Do I select the Terway or Flannel plugin for my Kubernetes cluster network?](#)

- **Flannel:** a simple and stable community [Flannel](#) CNI plugin. Flannel can interoperate with the high-speed network of Alibaba Cloud VPC to provide a high-performance and stable container network for clusters. However, it provides a limited amount of 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, you can use this network plugin to limit the bandwidth traffic of a single container. If you do not need to use the Network Policy, we recommend that you select Flannel. In other cases, we recommend that you select Terway.

**Note:**

- Terway provides the same Network Policy as Calico because Terway is integrated with the Felix component of Calico. If you create a cluster to use Calico, you can use Terway to switch to Alibaba Cloud Container Service for Kubernetes.
- Terway is integrated with the Felix component V2.6.6.

1.9.11 Allocate an ENI to a pod

This topic describes how to allocate an Elastic Network Interface (ENI) to a pod.

Context

- When you create a Kubernetes cluster, you need to select Network Plugin as Terway. For more information, see [Create a Kubernetes cluster](#).
- If you use a Kubernetes cluster that is installed with the Terway network plugin, you must make sure that the Terway plugin is V1.0.0.1 or later.

**Note:**

1. Log on to the Container Service console, click Clusters under the Kubernetes menu.
2. In the action column of the target cluster, choose More > Addon Upgrade.
3. On the Addon Upgrade page, view your current version of Terway.

4. Determine whether to upgrade according to Current Version and Upgradeable Version. If you want to upgrade Terway, click Upgrade in the action column.

Addon Upgrade ×					
Component	Current Version	Upgradeable Version	Consistency Check	Action	Status
alicloud-application-controller	v0.1.0.1-f832bed-aliyun	v0.1.0.1-f832bed-aliyun	Success	Latest	
alicloud-disk-controller	v1.11.2.2-a390cfb-aliyun	v1.11.2.2-a390cfb-aliyun	Success	Latest	
Cloud Controller Manager 🔗 Readme 🔗 Version Information	v1.9.3.59-ge3bc999-aliyun	v1.9.3.59-ge3bc999-aliyun	Success	Latest	
flexvolume	v1.11.2.32-af2d48c-aliyun	v1.11.2.32-af2d48c-aliyun	Success	Latest	
Nginx Ingress Controller 🔗 Readme 🔗 Version Information	v0.20.0.1-4597ce2-aliyun	v0.20.0.1-4597ce2-aliyun	Success	Latest	
terway	v1.0.8.11-g323b1f3-aliyun	v1.0.8.11-g323b1f3-aliyun	Success	Latest	

Refresh
Close

Procedure

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Kubernetes, choose Application > Deployment.
3. In the upper-right corner, click Create by Template.

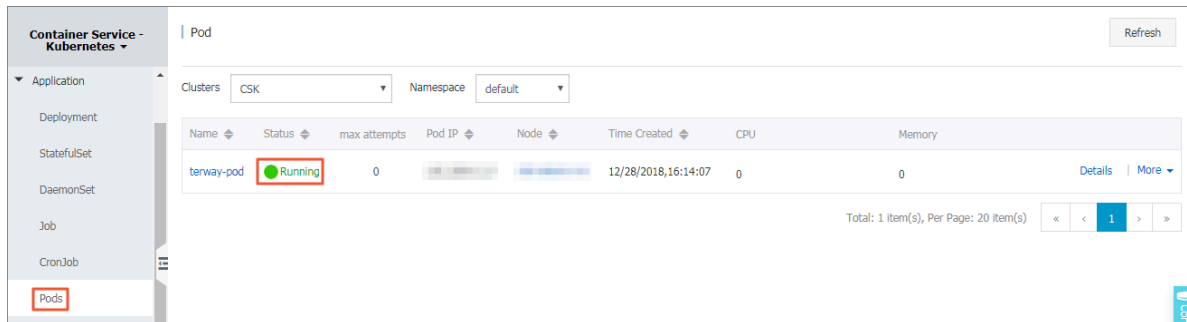
You can use the following YAML template to create a pod:

```
apiVersion : v1
kind : Pod
metadata :
  name : terway - pod
  labels :
    app : nginx
spec :
  containers :
    - name : nginx
      image : nginx
      ports :
        - containerPort : 80
      resources :
        limits :
```

```
aliyun / eni : 1
```

Result

1. In the left-side navigation pane under Kubernetes, choose Application > Pods. The pod named terway-pod is displayed.



2. In the left-side navigation pane under Kubernetes, click Clusters.
3. Click the name of the target cluster to view the cluster details.
4. In the Cluster Resource area, click VPC to view the VPC CIDR block of the cluster.
5. Run the following command to obtain the IP address of the deployed pod and verify that the IP address is within the VPC CIDR block of the cluster:

```
$ kubectl get pod -o wide
```

1.9.12 Use a network policy

Prerequisites

- You have created a Kubernetes cluster. For more information, see [Create a Kubernetes cluster](#).
- You have selected the Terway network plugin when creating the Kubernetes cluster. For more information, see [Create a Kubernetes cluster](#).
- You have connected to the Kubernetes cluster by using kubectl, see [Connect to a Kubernetes cluster by using kubectl](#).

Verify that an Nginx service is accessible to pods

1. Run the following command to create an Nginx application and expose it through a service named Nginx:

```
$ kubectl run nginx -- image = nginx
deployment.apps/nginx created
$ kubectl get pod
NAME          READY   STATUS    RESTARTS   AGE
nginx-64f497f8fd-znxbx 1 / 1   Running   0          45s
```

```
$ kubectl expose deployment nginx -- port = 80
service / nginx exposed
$ kubectl get service
NAME                TYPE                CLUSTER - IP          EXTERNAL - IP
PORT ( S )          AGE
kubernetes          ClusterIP           172 . 19 . 0 . 1      < none >
443 / TCP            3h
nginx               ClusterIP           172 . 19 . 8 . 48     < none >
80 / TCP             10s
```

2. Run the following command to create a pod named busybox and use the pod to access the Nginx service created in step 1:

```
$ kubectl run busybox -- rm - ti -- image = busybox / bin /
sh
kubectl run -- generator = deployment / apps . v1beta1 is
DEPRECATED and will be removed in a future version
. Use kubectl create instead .
If you don ' t see a command prompt , try pressing
enter .
/ # wget nginx
Connecting to nginx ( 172 . 19 . 8 . 48 : 80 )
index . html 100 % |
*****
612 0 : 00 : 00 ETA
/ #
```

Use a network policy to set the Nginx service to be accessible only to a specifically labeled application

1. Run the following command to create a *policy . yaml* file:

```
$ vim policy . yaml
kind : NetworkPol icy
apiVersion : networking . k8s . io / v1
metadata :
  name : access - nginx
spec :
  podSelecto r :
    matchLabel s :
      run : nginx
  ingress :
    - from :
      - podSelecto r :
        matchLabel s :
          access : " true "
```

2. Run the following command to create a network policy according to the *policy . yaml* file created in step 1:

```
$ kubectl apply - f policy . yaml
```

```
networkpol icy . networking . k8s . io / access - nginx created
```

3. Run the following command to verify that the Nginx service cannot be accessed if you do not define any access label in the command:

```
$ kubectl run busybox -- rm - ti -- image = busybox / bin / sh
If you don ' t see a command prompt , try pressing enter .
/ # wget nginx
Connecting to nginx ( 172 . 19 . 8 . 48 : 80 )
wget : can ' t connect to remote host ( 172 . 19 . 8 . 48 ) : Connection timed out
/ #
```

4. Run the following command to verify that the Nginx service can be accessed if an access label is defined in the command:

```
$ kubectl run busybox -- rm - ti -- labels =" access = true " -- image = busybox / bin / sh
If you don ' t see a command prompt , try pressing enter .
/ # wget nginx
Connecting to nginx ( 172 . 19 . 8 . 48 : 80 )
index . html 100 % |
*****
612 0 : 00 : 00 ETA
/ #
```

Use a network policy to specify a source IP CIDR block that can access a service exposed by an SLB service over the Internet

1. Run the following command to create an Alibaba Cloud SLB service for the preceding Nginx application, that is, specify `type = LoadBalancer` to expose the Nginx service to the Internet:

```
$ vim nginx - service . yaml
apiVersion : v1
kind : Service
metadata :
  labels :
    run : nginx
  name : nginx - slb
spec :
  externalTrafficPolicy : Local
  ports :
  - port : 80
    protocol : TCP
    targetPort : 80
  selector :
    run : nginx
  type : LoadBalancer

$ kubectl apply - f nginx - service . yaml
service / nginx - slb created

$ kubectl get service nginx - slb
```

NAME	TYPE	CLUSTER - IP	EXTERNAL - IP
nginx - slb	LoadBalancer	172 . 19 . 12 . 254	47 . 110
. 200 . 119	80 : 32240 / TCP	8m	

2. Run the following command to verify that the IP address of the created SLB service, that is, 47.110.200.119, cannot be accessed:

```
$ wget 47 . 110 . 200 . 119
-- 2018 - 11 - 21 11 : 46 : 05 -- http :// 47 . 110 . 200 . 119
/
Connecting to 47 . 110 . 200 . 119 : 80 ... failed :
Connection refused .
```



Note:

Access failure occurs due to the following reasons:

- You have configured access to the Nginx service only for the applications labeled with `access = true`.
- You have attempted to access the IP address of the SLB instance from outside the Kubernetes system. This is different from [Use a network policy to set the Nginx service to be accessible only to a specifically labeled application](#).

Solution: Modify the network policy and add a source IP CIDR block that is allowed to access the Nginx service.

3. Run the following command to view your local IP address:

```
$ curl myip . ipip . net

IP address : 10 . 0 . 0 . 1 from : China Beijing Beijing
# The local IP address varies by devices .
```

4. Run the following command to modify the created `policy . yaml` file:

```
$ vim policy . yaml
kind : NetworkPolicy
apiVersion : networking . k8s . io / v1
metadata :
  name : access - nginx
spec :
  podSelector :
    matchLabels :
      run : nginx
  ingress :
    - from :
      - podSelector :
          matchLabels :
            access : " true "
      - ipBlock :
          cidr : 100 . 64 . 0 . 0 / 10
      - ipBlock :
```

```

        cidr : 10 . 0 . 0 . 1 / 24      # Set the CIDR
block to which the local IP address belongs . This
is an example . Set the required parameters according
to your device .

$ kubectl apply -f policy . yaml
networkpolicy . networking . k8s . io / access - nginx
unchanged

```

**Note:**

- The outgoing interface of a network may have multiple IP addresses. We recommend that you specify an entire CIDR block.
- The SLB health check address belongs to the `100 . 64 . 0 . 0 / 10` CIDR block. Therefore, you must specify the `100 . 64 . 0 . 0 / 10` CIDR block.

5. Run the following command to verify that the Nginx service can be accessed:

```

$ kubectl run busybox -- rm - ti -- labels =" access = true
"-- image = busybox / bin / sh

If you don ' t see a command prompt , try pressing
enter .
/ # wget 47 . 110 . 200 . 119
Connecting to 47 . 110 . 200 . 119 ( 47 . 110 . 200 . 119 : 80
)
index . html          100 % |
*****|               612
0 : 00 : 00    ETA
/ #

```

Use a network policy to set a pod that can access only `www . aliyun . com`

1. Run the following command to obtain the IP address list resolved from the domain name of `www . aliyun . com` :

```

$ dig + short www . aliyun . com
www - jp - de - intl - adns . aliyun . com .
www - jp - de - intl - adns . aliyun . com . gds . alibabadns . com
.
v6wagbridg e . aliyun . com .
v6wagbridg e . aliyun . com . gds . alibabadns . com .
106 . 11 . 93 . 21
140 . 205 . 32 . 4
140 . 205 . 230 . 13
140 . 205 . 34 . 3

```

2. Run the following command to create a `busybox - policy` file:

```

$ vim busybox - policy . yaml
kind : NetworkPolicy
apiVersion : networking . k8s . io / v1
metadata :
  name : busybox - policy
spec :
  podSelector :

```

```

    matchLabels:
      s:
        run: busybox
  egress:
  - to:
    - ipBlock:
        cidr: 106.11.93.21/32
    - ipBlock:
        cidr: 140.205.32.4/32
    - ipBlock:
        cidr: 140.205.230.13/32
    - ipBlock:
        cidr: 140.205.34.3/32
  - to:
    - ipBlock:
        cidr: 0.0.0.0/0
  ports:
  - protocol: UDP
    port: 53

```

**Note:**

In the preceding `busybox - policy` file, an egress rule is set to specify the CIDR blocks that can be accessed by cluster applications. You need to set the condition that UDP requests are allowed. Otherwise, DNS resolution will fail.

3. Run the following command to create a network policy according to the `busybox - policy` file:

```

$ kubectl apply -f busybox - policy . yam
networkpol icy . networking . k8s . io / busybox - policy
created

```

4. Run the following command to verify that no website (for example, `www . google . com`) can be accessed except for `www . aliyun . com`:

```

$ kubectl run busybox -- rm - ti -- image = busybox / bin /
sh
If you don ' t see a command prompt , try pressing
enter .
/ # wget www . google . com
Connecting to www . google . com ( 64 . 13 . 192 . 74 : 80 )
wget : can ' t connect to remote host ( 64 . 13 . 192 .
74 ) : Connection timed out

```

5. Run the following command to verify that `www . aliyun . com` can be accessed:

```

/ # wget www . aliyun . com
Connecting to www . aliyun . com ( 140 . 205 . 34 . 3 : 80 )
Connecting to www . aliyun . com ( 140 . 205 . 34 . 3 : 443 )
wget : note : TLS certificat e validation not
implemente d
index . html 100 % |
*****| 462k
0 : 00 : 00 ETA

```

/ #

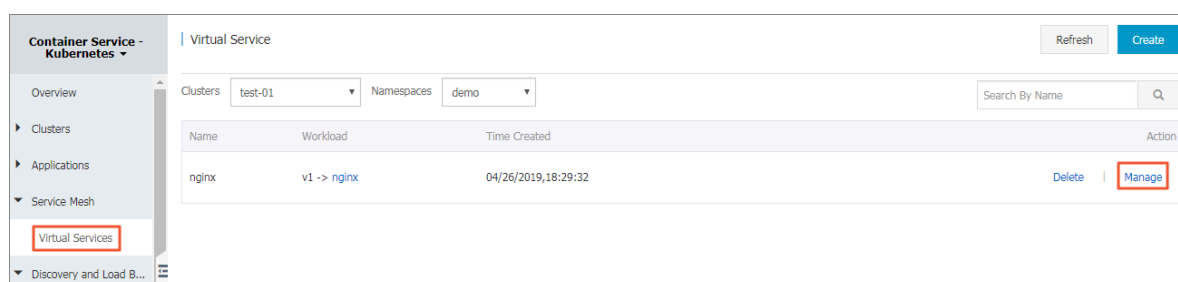
1.10 Service mesh

1.10.1 Manage traffic

This topic describes how to manage traffic by using Istio. Specifically, you can use Istio to set corresponding parameters for a load balancing algorithm, session affinity, connection pool, circuit breaker, or faulty injection.

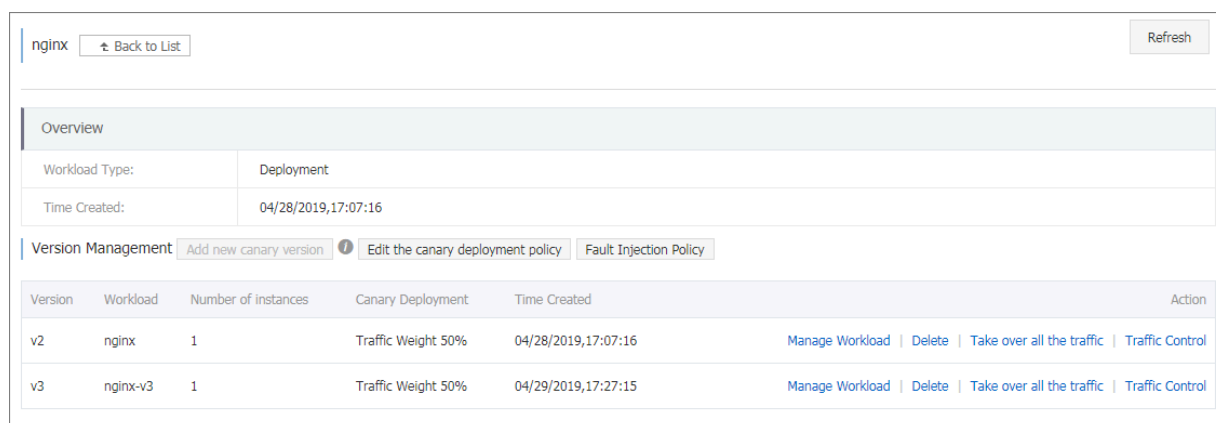
Before you begin

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Container Service-Kubernetes, choose Service Mesh > Virtual Services.
3. On the right of the target virtual service, click Manage. Then, you can manage traffic related to the service according to your needs.



Control traffic

Find the required version of the target virtual service, and then click Traffic Control.



- Set load balancing.

Istio provides two methods to set load balancing: Load Balancer Algorithm and Session Affinity.



Note:

These two methods are mutually exclusive.

Select a load balancing algorithm according to your needs.

- **ROUND_ROBIN:** Evenly distributes loads across the endpoints in the load balancing pool. This is the default algorithm used by the Istio proxy.
- **LEAST_CONN:** Selects two healthy hosts randomly and uses the host with fewer requests to provide service.
- **RANDOM:** Selects one healthy host randomly and evenly distributes loads on the endpoints in the load balancing pool. In the case that you have not set health checks, this is more efficient than ROUND_ROBIN.
- **PASSTHROUGH:** Forwards requests directly to the IP address specified by the client. For security purposes, we recommend that you exercise caution before implementing this algorithm.

Figure 1-1: Supported load balancing algorithms

The screenshot shows a 'Parameters Setting' dialog box with a close button (X) in the top right corner. It contains three tabs: 'Load Balancer Setting' (selected), 'Connection Pool Management', and 'Circuit Breaker Configuration'. Under the 'Load Balancer Setting' tab, there are two sub-tabs: 'Load Balancer Algorithm' (selected) and 'Session Affinity'. The 'Load Balancer Algorithm' sub-tab displays four radio button options: 'ROUND_ROBIN' (selected), 'LEAST_CONN', 'RANDOM', and 'PASSTHROUGH'. Below these options is a checkbox labeled 'Apply to all versions'. To the right of the radio buttons is a text area displaying the following YAML configuration:

```
trafficPolicy:
  loadBalancer:
    simple: ROUND_ROBIN
```

At the bottom right of the dialog box are 'OK' and 'Cancel' buttons.

Select a type of session affinity according to your needs.

- HTTP Headers: Obtains hashes based on a specific HTTP header.
- Cookie: Obtains hashes based on HTTP cookies.
- Source IP: Obtains hashes based on the source IP addresses.

Figure 1-2: Session affinity options

Parameters Setting

Load Balancer Setting

Connection Pool Management

Circuit Breaker Configuration

Load Balancer Alg

Session Affinity

☒ HTTP Headers

Name

User-Agent

☐ Cookie

Name

test

Life cycle

1s

☐ Source IP

☐ Apply to all versions

```
trafficPolicy:
  loadBalancer:
    consistentHash:
      httpHeaderName: ''
```

OK

Cancel

- Set a connection pool.
 - `maxConnections` : Indicates the maximum number of connections that Envoy will create for all the hosts in the upstream clusters. This parameter applies only to the connections that use the TCP or HTTP/1.1 protocol.
 - `connectTimeout` : Indicates the TCP connection timeout. The minimum value of this parameter must be greater than 1 ms. This parameter applies only to the connections that use the TCP or HTTP protocol.
 - `maxRequestsPerConnection` : Indicates the maximum number of requests that are destined for a backend through a connection. Setting this parameter to 1 disables the keep-alive feature. This parameter only applies to the connections that use the HTTP/1.1, HTTP/2, or gRPC protocol.
 - `maxRetries` : Indicates the maximum number of retries of an HTTP request that is sent to a destination host during a specified period of time. The default

value of this parameter is 3. This parameter only applies to the connections that use the HTTP/1.1, HTTP/2, or gRPC protocol.

Parameters Setting ✕

Load Balancer Setting

Connection Pool Management

Circuit Breaker Configuration

Maximum number of HTTP1 /TCP connections

100

Connection Timeout

100ms

Maximum number of requests per connection

100

Maximum number of retries

10

☐ Apply to all versions

```

trafficPolicy:
  connectionPool:
    tcp:
      maxConnections: 100
      connectTimeout: 100ms
    http:
      maxRequestsPerConnection: 100
      maxRetries: 10
  loadBalancer:
    simple: ROUND_ROBIN

```

OK

Cancel

- Set a circuit breaker.
- `consecutiveErrors` : Indicates the number of consecutive errors that occur to a host in a specified interval. If the value set is exceeded, the host where the consecutive errors occurred is removed from the connection pool. The default value of this parameter is 5.



Note:

If the upstream host is accessed through an HTTP connection, a status code of the `5xx` format is identified as an error. If the upstream host is accessed

through a TCP connection, a TCP connection timeout, a connection error, or a connection failure will be identified as an error.

- `maxEjectionPercent` : Indicates the maximum ratio of removable hosts to all the hosts in a load balancing pool of the upstream service. The default value of this parameter is 10%.
- `baseEjectionTime` : Indicates the minimum interval of a time an unhealthy host can be removed from the load balancing pool. The amount of time that an unhealthy host is removed from the load balancing pool is equal to this parameter multiplied by the number of times the host has already been removed. By using this parameter, the amount of time that an unhealthy host is removed from the load balancing pool can be increased automatically.
- `interval` : Indicates the period of time during which the system detects errors. The default value of this parameter is 10s. The minimum value of this parameter is 1 ms.

Parameters Setting ✕

Load Balancer Setting

Connection Pool Management

Circuit Breaker Configuration

Consecutive times with 5XX error code

10

Maximum % of hosts that can be ejected

80

Minimum ejection duration

3m

Time interval between ejection sweep analysis

1s

☐ Apply to all versions

```

trafficPolicy:
  loadBalancer:
    simple: ROUND_ROBIN
  outlierDetection:
    consecutiveErrors: 10
    interval: 1s
    baseEjectionTime: 3m
    maxEjectionPercent: 80

```

OK

Cancel

**Note:**

If you want the preceding settings to take effect for all the versions of the target virtual service, you need to select the Apply to all versions check box.

Inject faults to traffic

On the page that displays the details of the target virtual service, click **Fault Injection Policy**.

The screenshot shows the 'nginx' virtual service details page. At the top, there's a 'Back to List' button and a 'Refresh' button. The 'Overview' section displays 'Workload Type: Deployment' and 'Time Created: 04/28/2019,17:07:16'. The 'Version Management' section includes a table with two versions, v2 and v3, each with a 'Traffic Weight 50%'. A red box highlights the 'Fault Injection Policy' link in the 'Edit the canary deployment policy' dropdown menu.

Version	Workload	Number of instances	Canary Deployment	Time Created	Action
v2	nginx	1	Traffic Weight 50%	04/28/2019,17:07:16	Manage Workload Delete Take over all the traffic Traffic Control
v3	nginx-v3	1	Traffic Weight 50%	04/29/2019,17:27:15	Manage Workload Delete Take over all the traffic Traffic Control

You can inject two types of faults: delay faults and abort faults. The fault injection feature supports the HTTP protocol.

- Create a fault injection to delay traffic flow.

This type of fault injection is used to simulate faults, such as network faults and faults caused by upstream service overload.

Fault Injection

Delay Abort

Percentage 10 %

Delay(second) 10s

```
fault:
  delay:
    percent: 10
    fixedDelay: 10s
```

OK Cancel

Create a fault injection rule to delay traffic, set the following parameters.

- **percent** : Set the ratio of the requests to be delayed as all requests that are forwarded to the requested destination. The value range of this parameter is 0 to 100.
- **fixedDelay** : Set the delayed time before the specified ratio of requests are forwarded (in seconds by default). You can also set the delayed time in hours, minutes, or milliseconds. This is the required parameter for injecting a delay fault. The minimum value of this parameter is 1ms.

- Create an fault injection that terminates a request that comes from a downstream service and return a corresponding error to the downstream service.

This type of fault injection is used to simulate a condition where an error code occurs to the upstream service.

To create an fault injection that terminates a request that comes from a downstream service, set these parameters.

- `percent` : Set the ratio of the requests that you want to terminate to all the requests that are forwarded to the requested destination. The value range of this parameter is 0 to 100.
- `httpStatus` : Set the HTTP status code that will be returned to a client service. This is the required parameter for injecting an abort fault.

1.11 Server Load Balancer and Ingress management

1.11.1 Overview

Kubernetes clusters provide a diversity of approaches to access container applications, and support accessing internal services and realizing load balancing by means of Alibaba Cloud Server Load Balancer or Ingress.

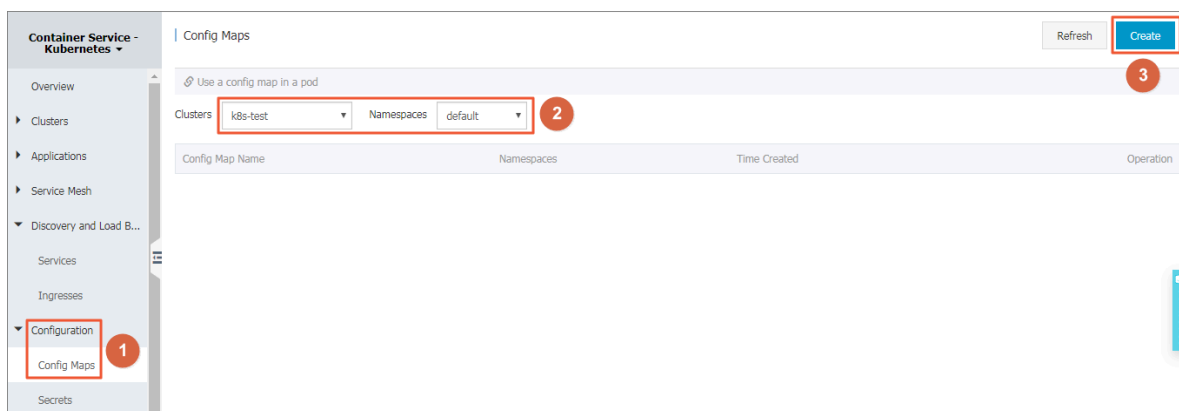
1.12 Config Map and Secret management

1.12.1 Create a Config Map

In the Container Service console, you can create a Config Map on the Config Maps page or by using a template.

Create a Config Map on Config Maps page

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Container Service-Kubernetes, choose **Configuration > Config Maps**.
3. Select the target cluster and namespace. Then, in the upper-right corner, click **Create**.



4. Complete the settings and then click OK.

- **Namespace:** Select the namespace to which the Config Map belongs. Config Map is a Kubernetes resource object that must be applied to the namespace.
- **Config Map Name:** Enter the Config Map name, which can contain lowercase letters, numbers, hyphens (-), and periods (.). The name cannot be empty.

Other resource objects must reference the Config Map name to obtain the configuration information.

- **Configuration:** Enter the Variable Name and the Variable Value. Then, click Add on the right. You can also click Edit, complete the configuration in the displayed dialog box, and click OK.

* Namespace:

default

* Config Map Name:

test-config

Name must consist of lowercase alphanumeric characters, '-' or '.'. Name cannot be empty.

Configuration:

Variable Name	Variable Value	Action
enemies	aliens	Edit Delete
lives	3	Edit Delete

Name

Value

Add

Variable key must be unique. Variable key and value cannot be empty.

Edit YAML file

OK

Cancel

In this example, configure the variables enemies and lives to pass the parameters aliens and 3 respectively.

YAML format

```
1 data:
2   enemies: aliens
3   lives: '3'
4 metadata:
5   name: test-config
6   namespace: default
7
```

* Configuration must be in YAML format.

OK

Cancel

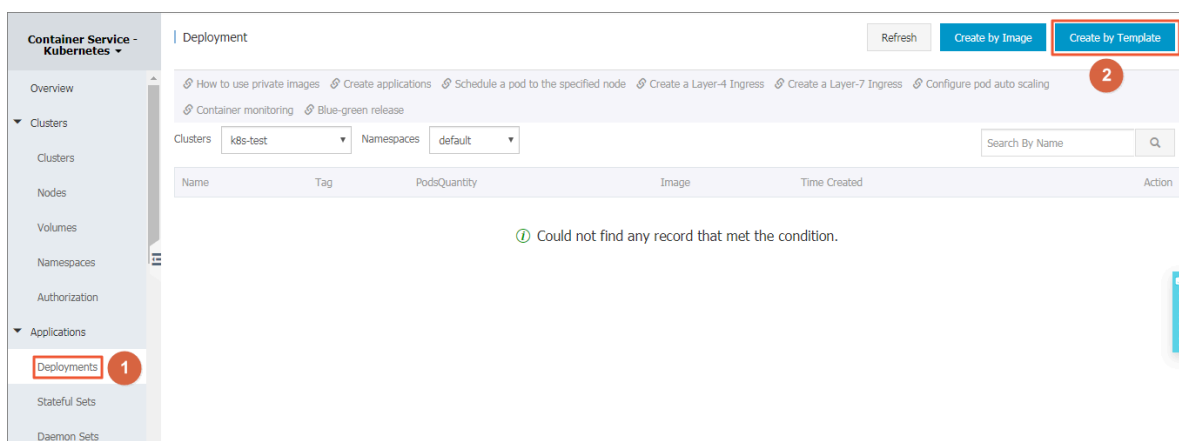
5. You can view the Config Map test-config on the Config Maps page after clicking OK.

Config Maps				Refresh	Create
Cluster: test					
Config Map Name	Namespace	Time Created		Operation	
test	default	2018-02-09 03:30:31		Delete	Modify
test-config	default	2018-02-09 05:56:47		Delete	Modify

Create a Config Map by using a template

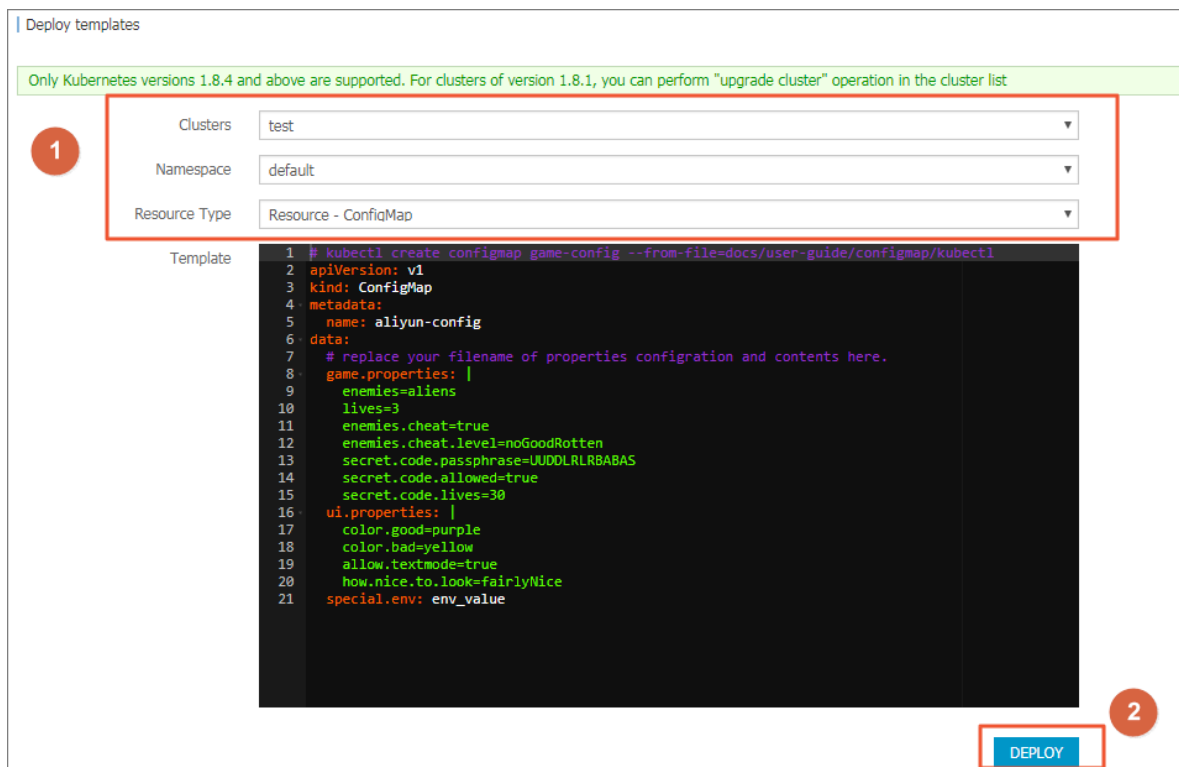
1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Container Service-Kubernetes, choose Applications > Deployments.

3. Click Create by template in the upper-right corner.



4. On the Deploy templates page, complete the settings and then click DEPLOY.

- **Clusters:** Select the cluster in which the Config Map is to be created.
- **Namespace:** Select the namespace to which the Config Map belongs. Config map is a Kubernetes resource object that must be applied to the namespace.
- **Resource Type:** You can write your own Config Map based on the Kubernetes YAML syntax rules, or select the sample template resource-ConfigMap. In the sample template, the Config Map is named as aliyun-config and includes two variable files `game . properties` and `ui . properties`. You can make modifications based on the sample template. Then, click DEPLOY.



5. After the deployment, you can view the Config Map `aliyun-config` on the Config Maps page.

Config Maps				Refresh	Create
Cluster: test					
Config Map Name	Namespace	Time Created		Operation	
aliyun-config	default	04/24/2018,15:41:32		Delete	Modify

1.12.2 Use a config map in a pod

You can use a config map in a pod in the following scenarios:

- Use a config map to define the pod environment variables.
- Use a config map to configure command line parameters.
- Use a config map in data volumes.

For more information, see [Configure a pod to use a ConfigMap](#).

Limits

To use a config map in a pod, make sure the config map and the pod are in the same cluster and namespace.

Create a config map

In this example, create a config map `special-config`, which includes two key-value pairs: `SPECIAL_LE VEL : very` and `SPECIAL_TY PE : charm`.

Create a config map by using an orchestration template


1. Log on to the [Container Service console](#).
2. Under Kubernetes, click **Application > Deployment**. Click **Create by template** in the upper-right corner.
3. Select the cluster and namespace from the **Clusters** and **Namespace** drop-down lists. Select a sample template or **Custom** from the **Resource Type** drop-down list. Click **DEPLOY**.

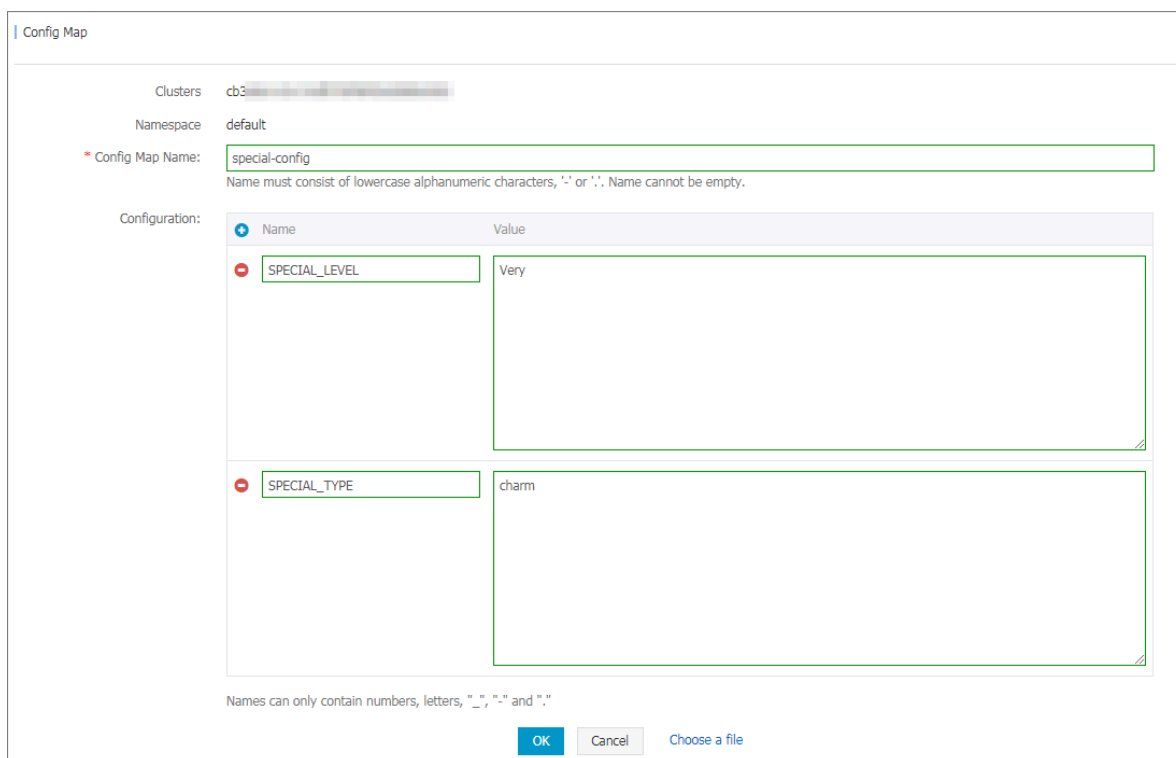
You can use the following YAML sample template to create a config map.

```
apiVersion : v1
kind : ConfigMap
metadata :
  name : special - config
  namespace : default
data :
  SPECIAL_LE VEL : very
```

SPECIAL_TY PE : charm

Create a config map on Config Maps page

1. Log on to the [Container Service console](#).
2. Under Kubernetes, choose Configuration > Config Maps in the left-side navigation pane.
3. Select the cluster and namespace from the Clusters and Namespace drop-down lists. Click Create in the upper-right corner.
4. Enter a config map name, click the plus icon , set the name and value for each entry of the config map, and then click OK.



Config Map

Clusters: cb3

Namespace: default

* Config Map Name: special-config
Name must consist of lowercase alphanumeric characters, '-' or '.'. Name cannot be empty.

Configuration:

Name	Value
SPECIAL_LEVEL	Very
SPECIAL_TYPE	charm

Names can only contain numbers, letters, "-", "." and "."

OK Cancel Choose a file

Use a config map to define pod environment variables

Use config map data to define pod environment variables

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Application > Deployment . Click Create by template in the upper-right corner.

3. Select the cluster and namespace from the Clusters and Namespace drop-down lists. Select a sample template or Custom from the Resource Type drop-down list. Click DEPLOY.

You can define the environment variables in a pod. Use `valueFrom` to reference the value of `SPECIAL_LEVEL` to define the pod environment variables.

See the following orchestration example:

```
apiVersion : v1
kind : Pod
metadata :
  name : config - pod - 1
spec :
  containers :
    - name : test - container
      image : busybox
      command : [ "/ bin / sh ", "- c ", " env " ]
      env :
        - name : SPECIAL_LEVEL_KEY
          valueFrom :
            configMapKeyRef :
              name : special - config
              key : SPECIAL_LEVEL
          restartPolicy : Never
```

valueFrom to specify env to reference the value of the config map .
 configMapKeyRef :
 name : special - config ## The
 referenced config map name .
 key : SPECIAL_LEVEL ## The
 referenced config map key .
 restartPolicy : Never

Similarly, to define the values of multiple config maps to the environment variable values of the pod, add multiple env parameters in the pod.

Configure all key-value pairs of a config map to pod environment variables

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Application > Deployment. Click Create by template in the upper-right corner.
3. Select the cluster and namespace from the Clusters and Namespace drop-down lists. Select a sample template or Custom from the Resource Type drop-down list. Click DEPLOY.

To configure all the key-value pairs of a config map to the environment variables of a pod, use the `envFrom` parameter. The key in a config map becomes the environment variable name in the pod.

See the following orchestration example:

```
apiVersion : v1
kind : Pod
```

```

metadata :
  name : config - pod - 2
spec :
  containers :
    - name : test - container
      image : busybox
      command : [ "/ bin / sh ", "- c ", " env " ]
      envFrom :
        - configMapRef :
            name : special - config
        - secretRef :
            name : special - config
      restartPolicy : Never

```

Use a config map to configure command line parameters

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click **Application > Deployment**. Click **Create by template** in the upper-right corner.
3. Select the cluster and namespace from the **Clusters and Namespace** drop-down lists. Select a sample template or **Custom** from the **Resource Type** drop-down list. Click **DEPLOY**.

You can use the config map to configure the commands or parameter values in the container by using the environment variable replacement syntax `$(VAR_NAME)`.

See the following orchestration example:

```

apiVersion : v1
kind : Pod
metadata :
  name : config - pod - 3
spec :
  containers :
    - name : test - container
      image : busybox
      command : [ "/ bin / sh ", "- c ", " echo $( SPECIAL_LE
VEL_KEY ) $( SPECIAL_TYPE PE_KEY )" ]
      env :
        - name : SPECIAL_LE VEL_KEY
          valueFrom :
            configMapKeyRef :
              name : special - config
              key : SPECIAL_LE VEL
        - name : SPECIAL_TYPE PE_KEY
          valueFrom :
            configMapKeyRef :
              name : special - config
              key : SPECIAL_TYPE PE

```

```
restartPolicy : Never
```

The output after running the pod is as follows:

```
very charm
```

Use a config map in data volumes

1. Log on to the [Container Service console](#).
2. Under the Kubernetes menu, click Application Deployment in the left-side navigation pane. Click Create by template in the upper-right corner.
3. Select the cluster and namespace from the Clusters and Namespace drop-down lists. Select a sample template or Custom from the Resource Type drop-down list. Click DEPLOY.

You can also use a config map in data volumes. Specifying the config map name under volumes stores the key-value pair data to the mountPath directory (`/ etc / config` in this example). It finally generates a configuration file with key as the file name and values as the contents of the file.

Then, the configuration file with key as the name and value as the contents is generated.

```
apiVersion : v1
kind : Pod
metadata :
  name : config - pod - 4
spec :
  containers :
    - name : test - container
      image : busybox
      command : [ "/ bin / sh ", "- c ", " ls / etc / config /" ]
  ## List the file names under this directory .
  volumeMounts :
    - name : config - volume
      mountPath : / etc / config
  volumes :
    - name : config - volume
      configMap :
        name : special - config
      restartPolicy : Never
```

Keys of the config map are output after running the pod.

```
SPECIAL_TYPE
```

SPECIAL_LEVEL

1.12.3 View a ConfigMap

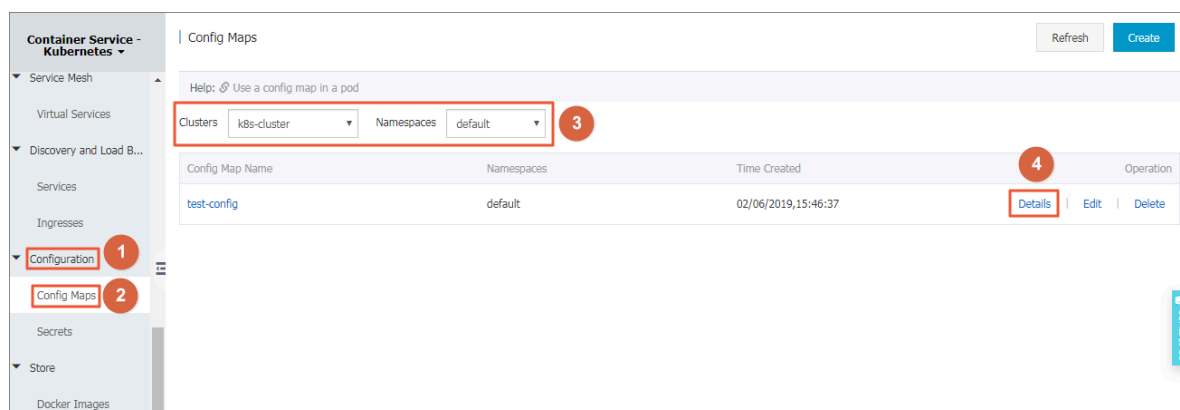
This topic describes how to view a created ConfigMap by using the Container Service console of Alibaba Cloud.

Prerequisites

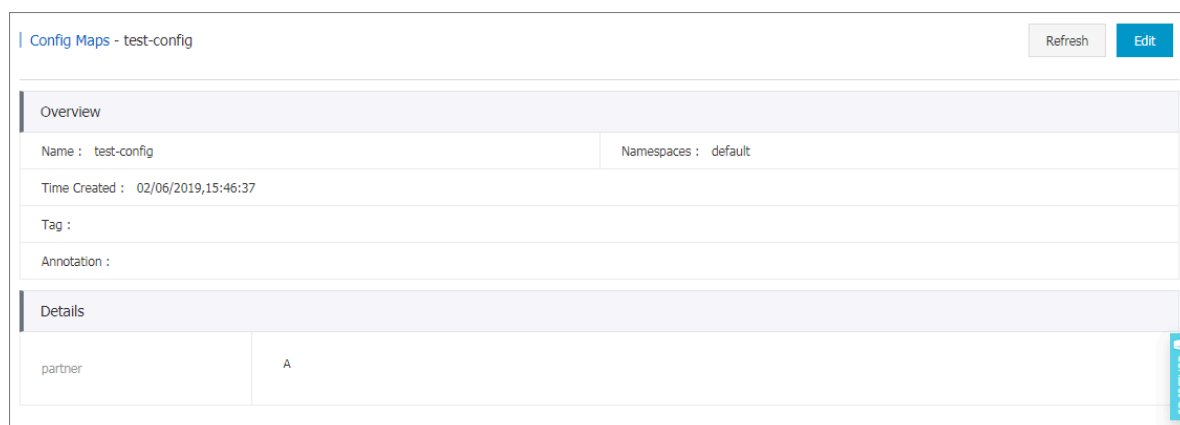
- A Kubernetes cluster is created. For more information, see [Create a Kubernetes cluster](#).
- A ConfigMap is created. For more information, see [Create a Config Map](#).

Procedure

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Container Service-Kubernetes, choose **Configuration > Config Maps**.
3. Select the target cluster and namespace, find the target ConfigMap, and then click **Details** in the Operation column.



Then, you can view the details of the ConfigMap.



1.12.4 Update a config map

You can modify the configurations of a config map.

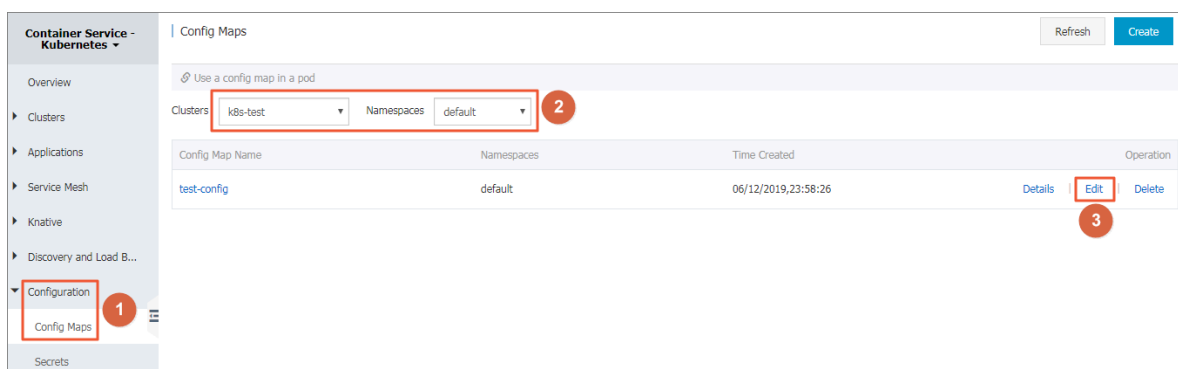


Note:

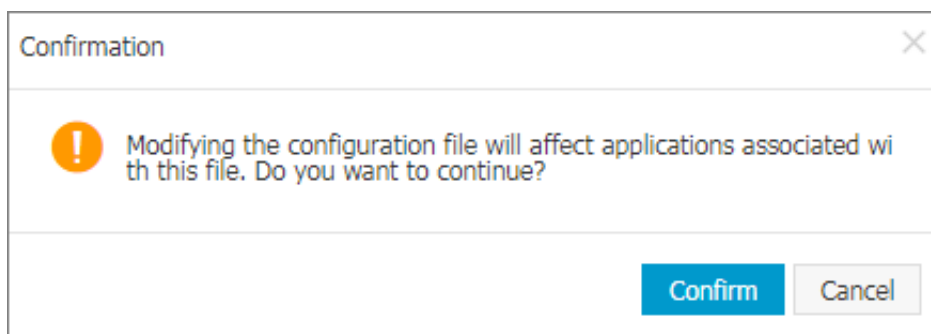
Updating a config map affects applications that use this config map.

Update a config map on Config Maps page

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click **Application > Config Maps** in the left-side navigation pane.
3. Select the cluster and namespace from the Clusters and Namespace drop-down lists. Click **Modify** at the right of the config map.



4. Click **Confirm** in the displayed dialog box.



5. Modify the configurations.

- Click Edit on the right of the configuration you want to modify. Update the configuration and then click Save.
- You can also click Edit YAML file. Click OK after making the modifications.

The screenshot shows a web-based configuration editor for a Kubernetes ConfigMap. On the left, there are filters for 'Clusters' (a dropdown), 'Namespaces' (set to 'default'), and 'Config Map Name' (set to 'test-config'). The main area is titled 'Configuration:' and contains a table with two columns: 'Name' and 'Value'. There are two rows in the table: one with 'enemies' in the Name column and 'aliens' in the Value column, and another with 'lives' in the Name column and '3' in the Value column. Each row has a red minus icon to its left. Below the table, there is a note: 'Names can only contain numbers, letters, "-", "." and "."'. At the bottom right, there are three buttons: 'OK' (blue), 'Cancel' (grey), and 'Choose a file' (blue).

Name	Value
enemies	aliens
lives	3

Names can only contain numbers, letters, "-", "." and "."

OK Cancel Choose a file

6. After modifying the configurations, click OK.

Update a config map in Kubernetes dashboard

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Clusters in the left-side navigation pane.

3. Click Dashboard at the right of the cluster.

Container Service - Kubernetes

Overview

Clusters

Nodes

Volumes

Namespaces

Authorization

Applications

Deployments

Stateful Sets

Daemon Sets

Jobs

Cron Jobs

Pods

Cluster List

You can create up to 100 clusters and add up to 1000 nodes in each cluster. To request a quota increase, [submit a ticket](#). [Refresh](#) [Create Kubernetes Cluster](#)

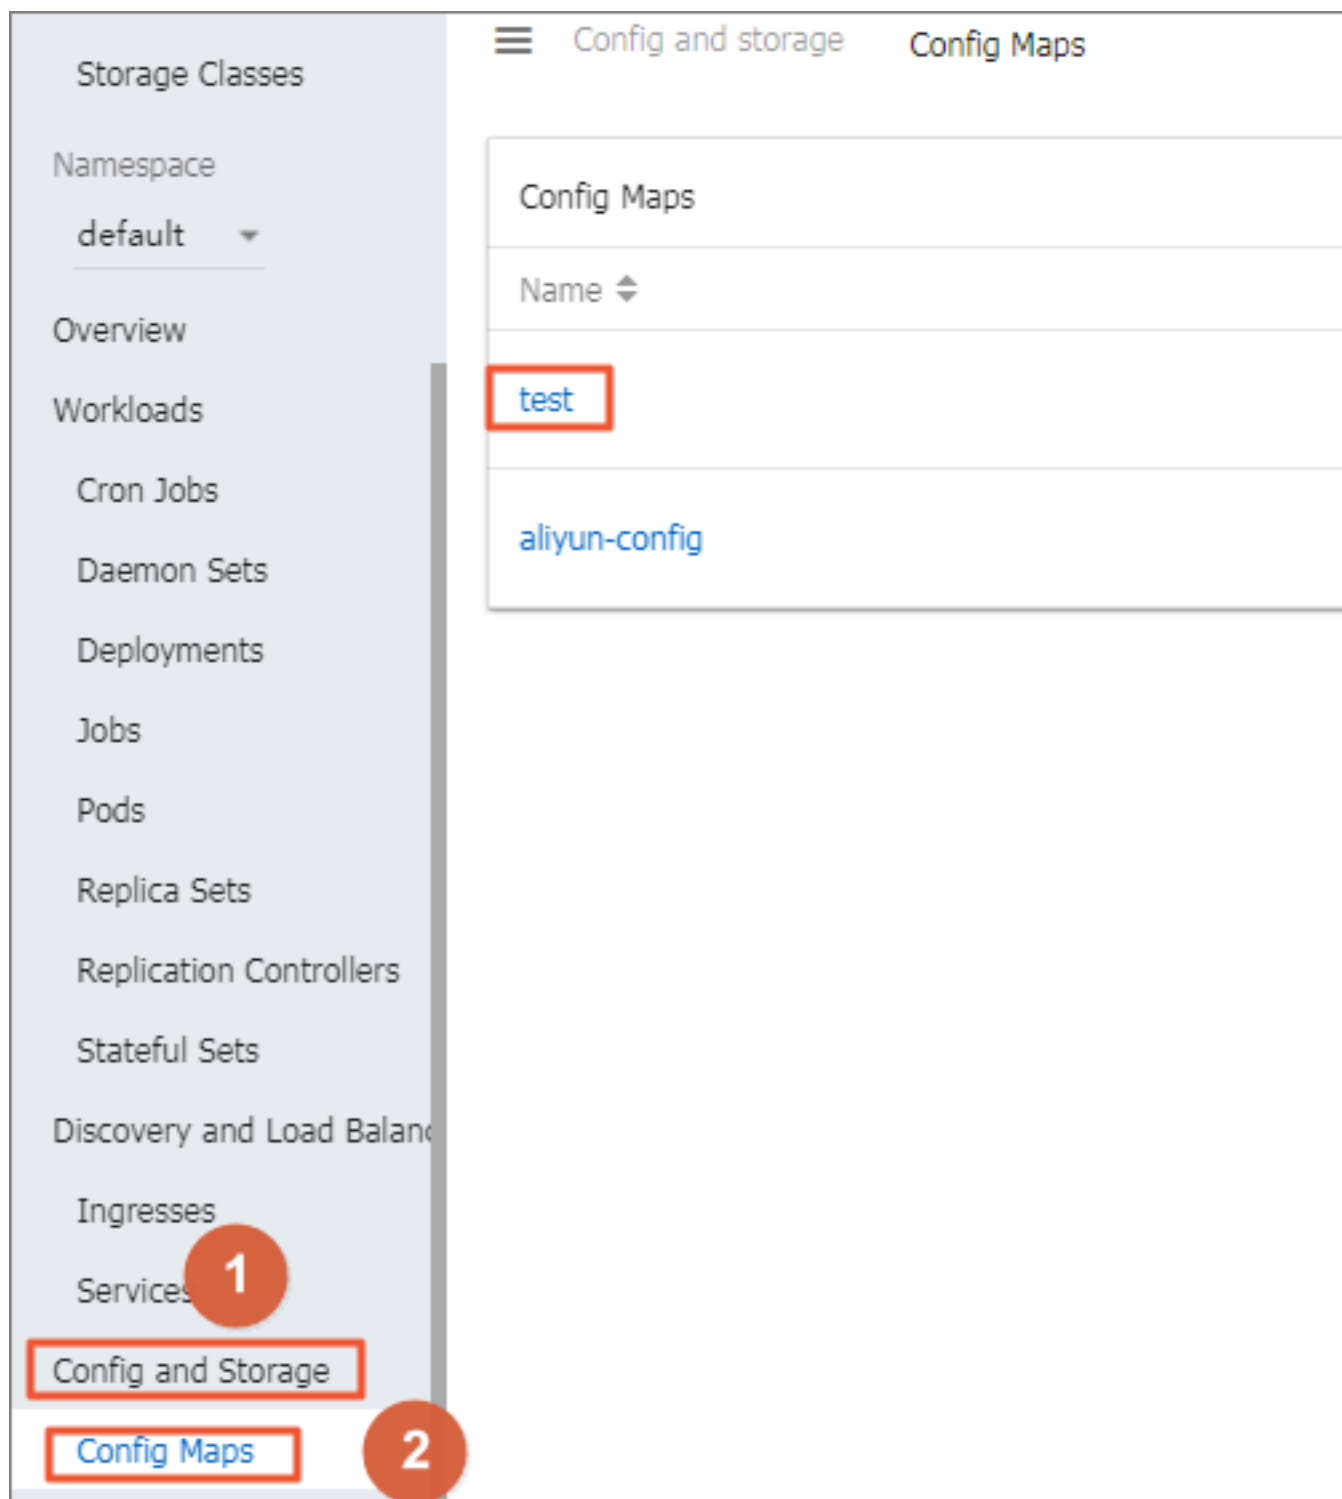
[Create cluster](#) [Create GPU clusters](#) [Scale cluster](#) [Connect to Kubernetes cluster via kubectl](#) [Manage applications with commands](#) [Plan Kubernetes CIDR Blocks in VPC Networks](#) [Troubleshoot cluster creation failures](#) [Authorization management](#) [Collect Kubernetes Diagnostics Information](#) [Submit ticket](#)

Name

Tags

Cluster Name/ID	Tags	Cluster Type	Region (All)	Network Type	Cluster Status	Number of Nodes	Time Created	Version	Action
managed-cluster		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		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		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		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. Under Kubernetes, select a namespace, click Config and Storage > Secrets in the left-side navigation pane. Select the target secret and click Actions > View/edit YAML.



5. The Edit a Secret dialog box appears. Modify the configurations and then click UPDATE.



Edit a Config Map

```
1 {
2   "kind": "ConfigMap",
3   "apiVersion": "v1",
4   "metadata": {
5     "name": "test",
6     "namespace": "default",
7     "selfLink": "/api/v1/namespaces/default/configmaps/test",
8     "uid": "0a826463-479e-11e8-a84c-00163e101791",
9     "resourceVersion": "52788",
10    "creationTimestamp": "2018-04-24T09:01:14Z"
11  },
12  "data": {
13    "enemies": "aliens",
14    "lives": "3"
15  }
16 }
```

CANCEL COPY UPDATE

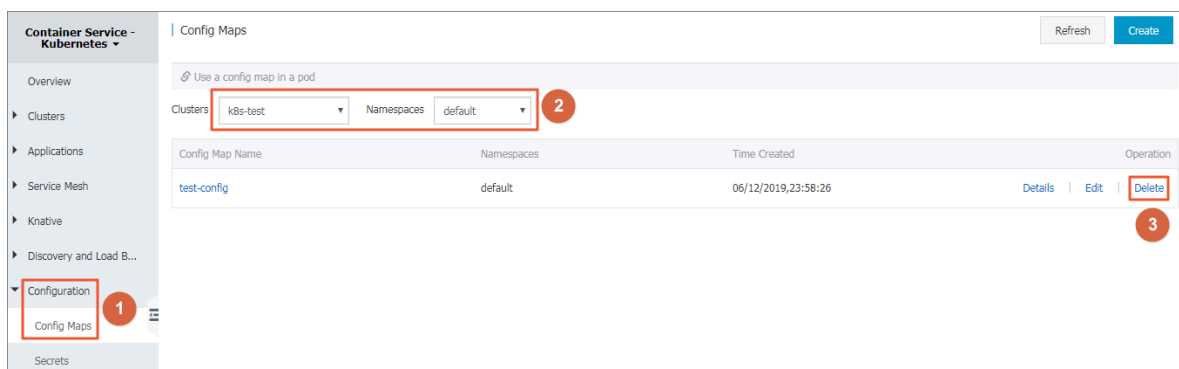
1.12.5 Delete a config map

You can delete a config map that is no longer in use.

Delete a config map on Config Maps page

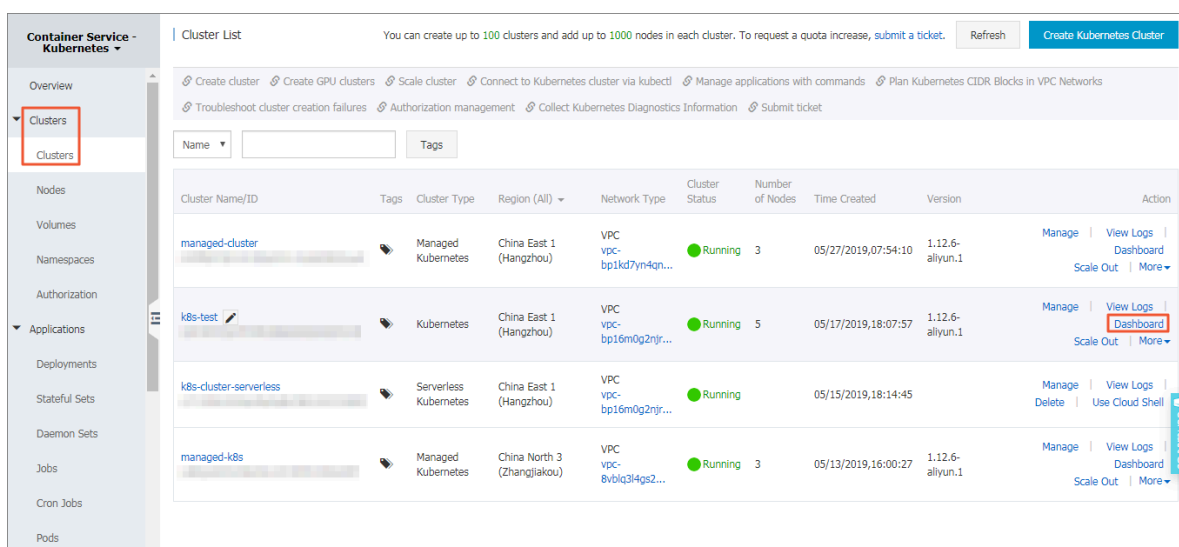
1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Application > Config Maps in the left-side navigation pane.

3. Select the target cluster from the Cluster drop-down list. Click Delete at the right of the config map.

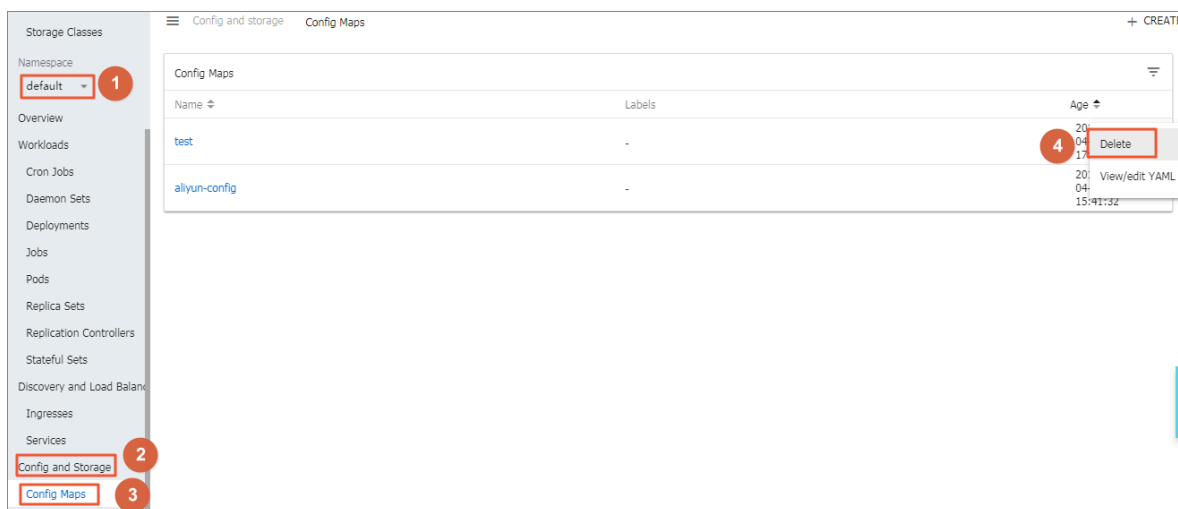


Delete a config map in Kubernetes dashboard

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Clusters in the left-side navigation pane.
3. Click Clusters in the left-side navigation pane, select the target cluster, and click Dashboard on the right.



4. Under Kubernetes, select a namespace, click Config and Storage > Secrets in the left-side navigation pane. Click the actions button on the right and click Delete in the drop-down list.



5. Click Delete in the displayed dialog box.

1.12.6 Create a Secret

This topic describes how to use the Container Service console to create a Secret.

Prerequisites

A Kubernetes cluster is created. For more information, see [Create a Kubernetes cluster](#).

Context

We recommend that you use Secrets for sensitive configurations in Kubernetes clusters, such as passwords and certificates.

Secrets have many types. For example:

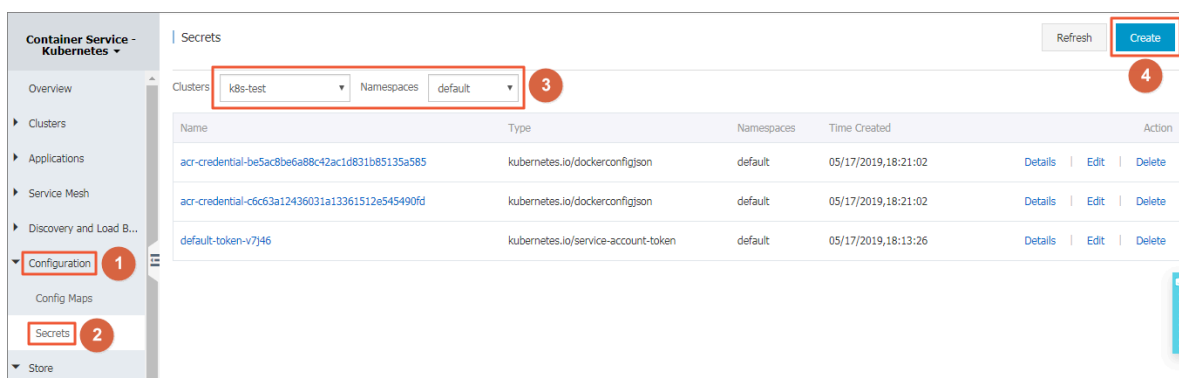
- **Service Account:** Automatically created by Kubernetes, which is used to access Kubernetes APIs and is automatically mounted to the pod directory `/run/secrets/kubernetes.io/serviceaccount`.
- **Opaque:** Secret in the base64 encoding format, which is used to store sensitive information such as passwords and certificates.

By default, you can only create secrets of the Opaque type in the Container Service console. Opaque data is of the map type, which requires the value to be in the base64 encoding format. Alibaba Cloud Container Service supports creating Secrets with one click and automatically encoding the clear data to base64 format.

You can also manually create Secrets by using command lines. For more information, see [Kubernetes Secrets](#) .

Procedure

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Container Service-Kubernetes, choose Configuration > Secrets.
3. Select the target cluster and namespace. Then, in the upper-right corner, click Create.



4. Complete the configurations to create a Secret.



Note:

To enter the clear data of the secret, select the Encode data values using Base64 check box.

* Name 1
Name must consist of lowercase alphanumeric characters, '-' or '.'. Name cannot be empty.

* Type ☒ Opaque ☐ Private Repository Logon Password ☐ TLS Certificate

* Data

Name	Value
password	1qaz2wsx09
username	admin

Names can only contain numbers, letters, "_", "-" and "."

☒ Encode data values using Base64 3

- Name: Enter the Secret name, which must be 1–253 characters long, and can only contain lowercase letters, numbers, hyphens (-), and dots (.).
- Configure the Secret data. Click the add icon next to Name and enter the name and value of the Secret, namely, the key-value pair. In this example, the Secret contains two values: `username : admin` and `password : 1f2d1e2e67 df`.
- Click OK.

5. The Secret page appears. You can view the created Secret in the Secret list.

Secrets					Refresh	Create
Clusters	k8s-test	Namespaces	default			
Name	Type	Namespaces	Time Created	Action		
account	Opaque	default	05/20/2019,10:52:53	Details	Edit	Delete
acr-credential-be5ac8be6a88c42ac1d831b85135a585	kubernetes.io/dockerconfigjson	default	05/17/2019,18:21:02	Details	Edit	Delete
acr-credential-c6c63a12436031a13361512e545490fd	kubernetes.io/dockerconfigjson	default	05/17/2019,18:21:02	Details	Edit	Delete
default-token-v7j46	kubernetes.io/service-account-token	default	05/17/2019,18:13:26	Details	Edit	Delete

1.12.7 View secret details

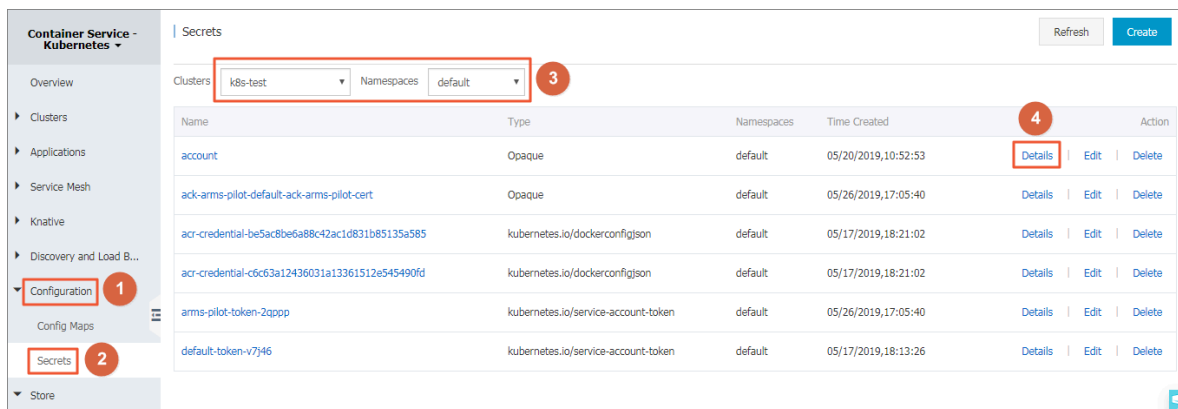
You can view the details of a created secret in the Container Service console.

Prerequisites

- You have created an Kubernetes cluster. For more information, see [Create a Kubernetes cluster](#).
- You have created a secret. For more information, see [Create a Secret](#).

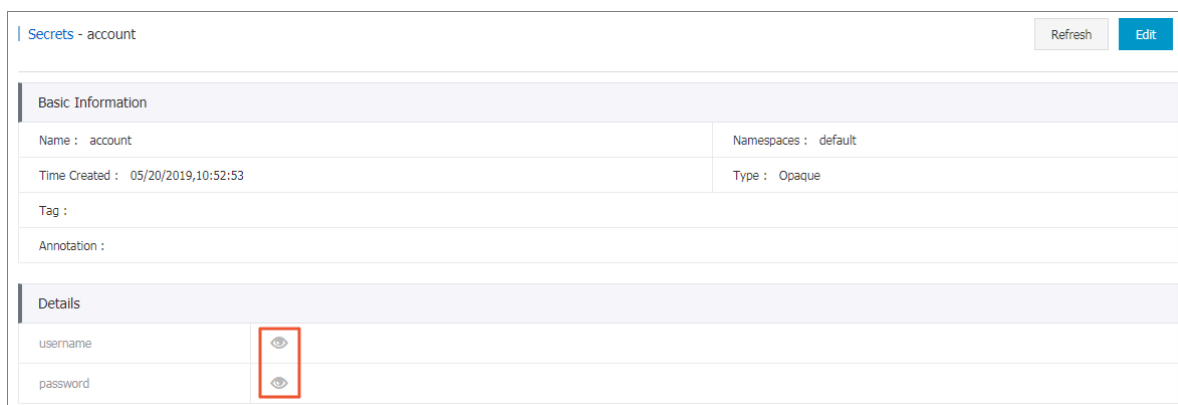
Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Application > Secrets in the left-side navigation pane.
3. Select the cluster and namespace from the Clusters and Namespace drop-down lists. Click Detail at the right of the secret.



4. You can view the basic information of the secret, and the data that the secret contains.

Click the icon at the right of the data name under Detail to view the clear data.



1.12.8 Update a secret

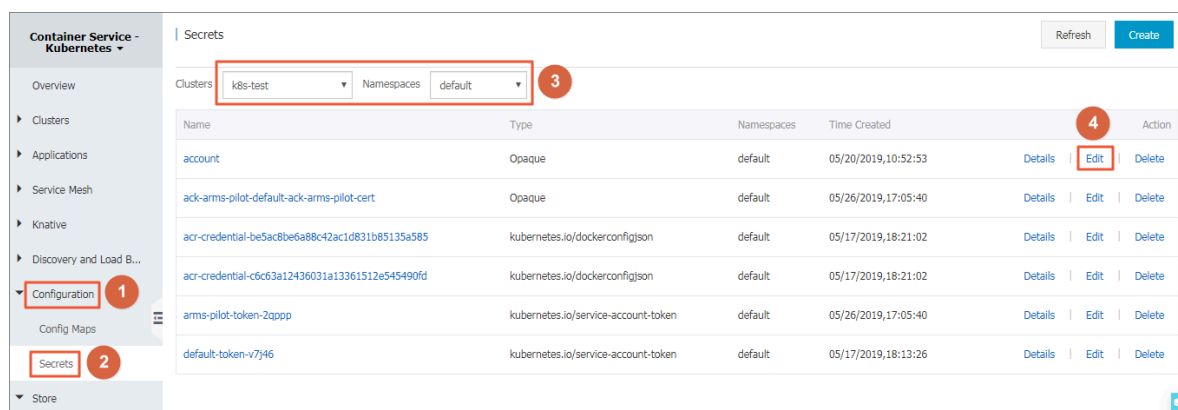
You can update an existing secret directly in the Container Service console.

Prerequisites

- You have created an Kubernetes cluster. For more information, see [Create a Kubernetes cluster](#).
- You have created a secret. For more information, see [Create a Secret](#).

Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Application > Secrets in the left-side navigation pane.
3. Select the cluster and namespace from the Clusters and Namespace drop-down lists. Click Edit at the right of the secret.



4. Update the secret data on the Edit Secret page.

Namespaces: default

* Name: account

* Type: ☒ Opaque ☐ Private Repository Logon Password ☐ TLS Certificate

* Data:

Name	Value
username	admin
password	1qaz2wsx09

Names can only contain numbers, letters, "_", "-" and "."

5. Click OK.

1.12.9 Delete a secret

You can delete an existing secret directly in the Container Service console.

Prerequisites

- You have created an Kubernetes cluster. For more information, see [Create a Kubernetes cluster](#).
- You have created a secret. For more information, see [Create a Secret](#).

Context



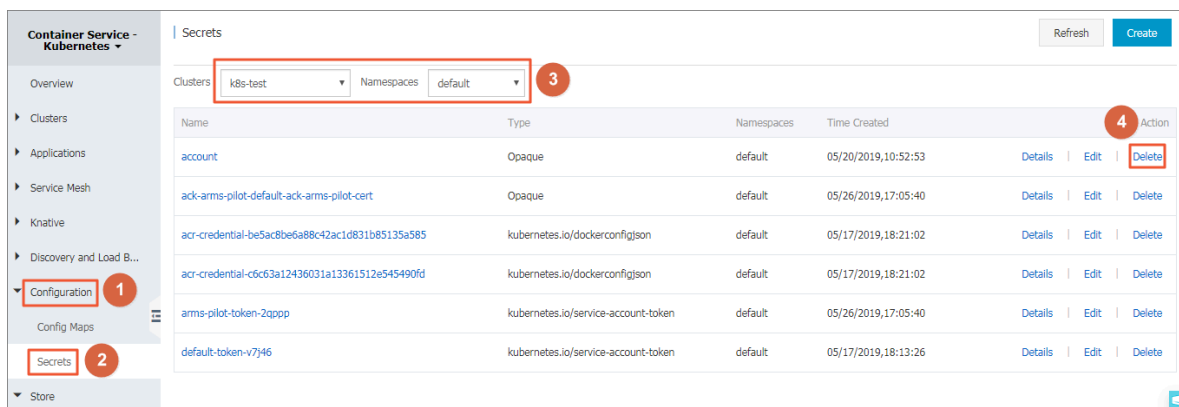
Note:

Do not delete the secret generated when the cluster is created.

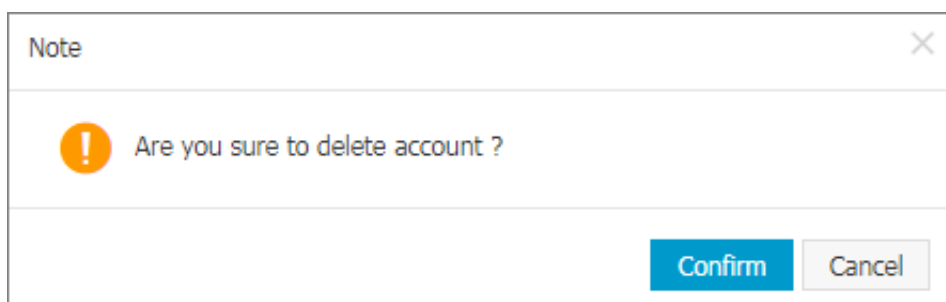
Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Application > Secrets in the left-side navigation pane.

3. Select the cluster and namespace from the Clusters and Namespace drop-down lists. Click Delete at the right of the secret.



4. Click Confirm in the displayed dialog box.



1.13 Storage management

1.13.1 Overview

Container Service supports automatically binding Kubernetes pods to Alibaba Cloud cloud disks, NAS, and Object Storage Service (OSS).

Currently, static storage volumes and dynamic storage volumes are supported.

See the following table for how each type of data volumes supports the static data volumes and dynamic data volumes.

Alibaba Cloud storage	Static data volume	Dynamic data volume
Alibaba Cloud cloud disk	<p>You can use the cloud disk static storage volumes by:</p> <ul style="list-style-type: none"> Using the volume method. Using PV/PVC. 	Supported.

Alibaba Cloud storage	Static data volume	Dynamic data volume
Alibaba Cloud NAS	You can use the NAS static storage volumes by: <ul style="list-style-type: none">· Using flexvolume plug-in.<ul style="list-style-type: none">- Using the volume method.- Using PV/PVC.· Using NFS drive of Kubernetes.	Supported.
Alibaba Cloud OSS	You can use the OSS static storage volumes by: <ul style="list-style-type: none">· Using the volume method.· Using PV/PVC.	Not supported.

1.13.2 Install the plug-in

Deploy the Alibaba Cloud Kubernetes storage plug-in by using the following yaml configurations.



Note:

If your Kubernetes cluster is created before February 6th, 2018, install the Alibaba Cloud Kubernetes storage plug-in before using the data volumes. If your Kubernetes cluster is created after February 6th, 2018, you can directly use the data volumes without installing the Alibaba Cloud Kubernetes storage plug-in.

Limits

Currently, CentOS 7 operating system is supported.

Instructions

- Disable the `-- enable - controller - attach - detach` option by using kubelet if you use the flexvolume. By default, Alibaba Cloud Kubernetes clusters have disabled this option.
- Deploy flexvolume in the kube-system user space.

Verify that the installation is complete

On the master node:

- Run the `kubectl get pod -n kube-system | grep flexvolume` command. Output is the list of running pods (number of nodes).
- Run the `kubectl get pod -n kube-system | grep alicloud-disk-controller` command. Output is the list of running pods.

Installation example

Install flexvolume

```

apiVersion : apps / v1 # for versions before 1.8.0 use
extensions / v1beta1
kind : DaemonSet
metadata :
  name : flexvolume
  namespace : kube-system
  labels :
    k8s-volume : flexvolume
spec :
  selector :
    matchLabels :
      name : acs-flexvolume
  template :
    metadata :
      labels :
        name : acs-flexvolume
    spec :
      hostPID : true
      hostNetwork : true
      tolerations :
        - key : node-role.kubernetes.io/master
          operator : Exists
          effect : NoSchedule
      containers :
        - name : acs-flexvolume
          image : registry.cn-hangzhou.aliyuncs.com/acs/flexvolume:v1.9.7-42e8198
          imagePullPolicy : Always
          securityContext :
            privileged : true
          env :
            - name : ACS_DISK
              value : "true"
            - name : ACS_NAS
              value : "true"
            - name : ACS_OSS
              value : "true"
          resources :
            limits :
              memory : 200Mi
            requests :
              cpu : 100m
              memory : 200Mi
          volumeMounts :
            - name : usrdir
              mountPath : /host/usr/
            - name : etcdir
              mountPath : /host/etc/
            - name : logdir
              mountPath : /var/log/alibabacloud/

```

```

volumes :
- name : usrdir
  hostPath :
    path : / usr /
- name : etcdire
  hostPath :
    path : / etc /
- name : logdire
  hostPath :
    path : / var / log / alicloud /

```

Install Disk provisioner

```

---
kind : StorageClass
apiVersion : storage . k8s . io / v1beta1
metadata :
  name : alicloud - disk - common
provisioner : alicloud / disk
parameters :
  type : cloud
---
kind : StorageClass
apiVersion : storage . k8s . io / v1beta1
metadata :
  name : alicloud - disk - efficiency
provisioner : alicloud / disk
parameters :
  type : cloud_efficiency
---
kind : StorageClass
apiVersion : storage . k8s . io / v1beta1
metadata :
  name : alicloud - disk - ssd
provisioner : alicloud / disk
parameters :
  type : cloud_ssd
---
kind : StorageClass
apiVersion : storage . k8s . io / v1beta1
metadata :
  name : alicloud - disk - available
provisioner : alicloud / disk
parameters :
  type : available
---
kind : ClusterRole
apiVersion : rbac . authorization . k8s . io / v1beta1
metadata :
  name : alicloud - disk - controller - runner
rules :
- apiGroups : ["" ]
  resources : [ " persistent volumes " ]
  verbs : [ " get ", " list ", " watch ", " create ", " delete " ]
- apiGroups : ["" ]
  resources : [ " persistent volumeclaims " ]
  verbs : [ " get ", " list ", " watch ", " update " ]
- apiGroups : [ " storage . k8s . io " ]
  resources : [ " storageclasses " ]
  verbs : [ " get ", " list ", " watch " ]
- apiGroups : ["" ]
  resources : [ " events " ]
  verbs : [ " list ", " watch ", " create ", " update ", " patch " ]

```

```

---
apiVersion : v1
kind : ServiceAccount
metadata :
  name : alicloud - disk - controller
  namespace : kube - system
---
kind : ClusterRoleBinding
apiVersion : rbac.authorization.k8s.io/v1beta1
metadata :
  name : run - alicloud - disk - controller
subjects :
- kind : ServiceAccount
  name : alicloud - disk - controller
  namespace : kube - system
roleRef :
  kind : ClusterRole
  name : alicloud - disk - controller - runner
  apiGroup : rbac.authorization.k8s.io
---
kind : Deployment
apiVersion : extensions/v1beta1
metadata :
  name : alicloud - disk - controller
  namespace : kube - system
spec :
  replicas : 1
  strategy :
    type : Recreate
  template :
    metadata :
      labels :
        app : alicloud - disk - controller
    spec :
      tolerations :
        - effect : NoSchedule
          operator : Exists
          key : node-role.kubernetes.io/master
        - effect : NoSchedule
          operator : Exists
          key : node.cloudprovider.kubernetes.io/uninitialized
      nodeSelector :
        node-role.kubernetes.io/master : ""
      serviceAccount : alicloud - disk - controller
      containers :
        - name : alicloud - disk - controller
          image : registry.cn-hangzhou.aliyuncs.com/acs/alibabacloud-disk-controller:v1.9.3-ed710ce
          volumeMounts :
            - name : cloud-config
              mountPath : /etc/kubernetes/
            - name : logdir
              mountPath : /var/log/alibabacloud/
      volumes :
        - name : cloud-config
          hostPath :
            path : /etc/kubernetes/
        - name : logdir
          hostPath :

```

```
path : / var / log / alicloud /
```

1.13.3 Use Alibaba Cloud cloud disk volumes

You can use Alibaba Cloud cloud disk volumes in a Kubernetes cluster of Alibaba Cloud Container Service.

You can mount an Alibaba Cloud cloud disk to a Kubernetes cluster by using the following two methods:

- [Static volumes](#)

You can use a static cloud disk volume in either of the following ways:

- [Use a cloud disk through a volume.](#)
- [Use a cloud disk through a PV and PVC.](#)

- [Dynamic volumes](#)



Note:

Depending on the type of cloud disk you create, the following requirements must be met:

- 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.

Static volumes

You can use an Alibaba Cloud cloud disk through a volume or through a PV and PVC.

Prerequisites

You have created a cloud disk in the ECS console. For more information, see [Create a cloud disk](#).

Limits

- A cloud disk is a non-shared storage device and can be mounted to only one pod.
- You must have created a cloud disk and obtained the disk ID before using the cloud disk volume. For more information, see [Create a cloud disk](#).
- The volumeId parameter indicates the ID of a mounted cloud disk. The volume name and PV name must be the same as the value of the volumeId parameter.

- In a Kubernetes cluster, a cloud disk can be mounted only to a node that resides in the same zone as the cloud disk.
- Only Pay-As-You-Go cloud disks can be mounted. In a Kubernetes cluster, the ECS instance billing method can be changed to Subscription, but the cloud disk billing method cannot be changed to Subscription. Otherwise, the cloud disks will fail to be mounted.

Use a cloud disk through a volume

Use the following `disk - deploy . yaml` file to create a pod:

```
apiVersion : extensions / v1beta1
kind : Deployment
metadata :
  name : nginx - disk - deploy
spec :
  replicas : 1
  template :
    metadata :
      labels :
        app : nginx
    spec :
      containers :
        - name : nginx - flexvolume - disk
          image : nginx
          volumeMounts :
            - name : " d - bp1j17ifxf asvts3tf40 "
              mountPath : "/" data "
          volumes :
            - name : " d - bp1j17ifxf asvts3tf40 "
              flexVolume :
                driver : " alicloud / disk "
                fsType : " ext4 "
                options :
                  volumeId : " d - bp1j17ifxf asvts3tf40 "
```

Use a cloud disk through a PV and PVC

Step 1: Create a cloud disk PV

You can create a cloud disk PV in the Container Service console or by using a YAML file.

Create a PV by using a YAML file

Use the following `disk - pv . yaml` file to create a PV:



Note:

The PV name must be the same as the cloud disk ID.

```
apiVersion : v1
kind : Persistent Volume
```

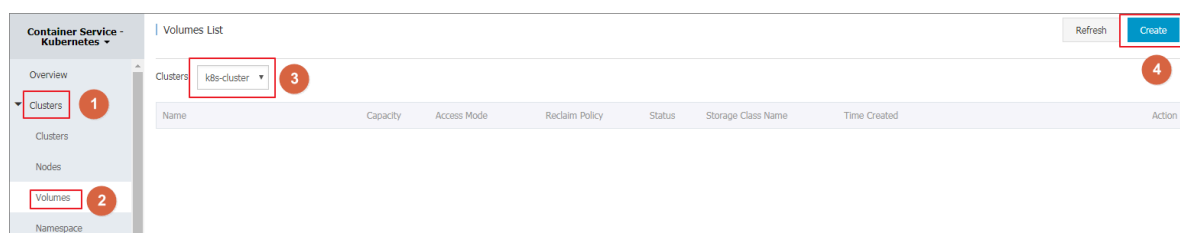
```

metadata :
  name : d - bp1j17ifxf asvts3tf40
  labels :
    failure - domain . beta . kubernetes . io / zone : cn - hangzhou
- b
  failure - domain . beta . kubernetes . io / region : cn -
hangzhou
spec :
  capacity :
    storage : 20Gi
  storageClassName : disk
  accessModes :
    - ReadWriteOnce
  flexVolume :
    driver : "alicloud / disk "
    fsType : "ext4 "
    options :
      volumeId : "d - bp1j17ifxf asvts3tf40 "

```

Create a cloud disk volume in the Container Service console

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Kubernetes, choose Clusters > Volumes.
3. Select the target cluster and then click Create in the upper-right corner.



4. In the displayed dialog box, set the volume parameters.

- **Storage type:** Cloud Disk is used in this example.
- **Access Mode:** By default, it is set to ReadWriteOnce.
- **Cloud Disk ID:** We recommend that you select a cloud disk that is in the same region and zone as the cluster.
- **File System Type:** Select a data type for the data to be stored. The available data types include ext4, ext3, xfs, and vfat. The default setting is ext4.
- **Tag:** Add tags to the volume.

5. Click Create.

Step 2: Create a PVC

Use the following `disk - pvc . yaml` file to create a PVC:

```
kind : Persistent VolumeClaim
apiVersion : v1
metadata :
  name : pvc - disk
spec :
  accessModes :
    - ReadWriteOnce
  storageClassName : disk
  resources :
    requests :
      storage : 20Gi
```

Step 3: Create a pod

Use the following `disk - pod . yml` file to create a pod:

```
apiVersion : v1
kind : Pod
metadata :
  name : " flexvolume - alicloud - example "
spec :
  containers :
    - name : " nginx "
      image : " nginx "
      volumeMounts :
        - name : pvc - disk
          mountPath : "/" data "
  volumes :
    - name : pvc - disk
      persistentVolumeClaim :
        claimName : pvc - disk
```

Dynamic volumes

To use a dynamic volume, you need to manually create a StorageClass, and specify a cloud disk type through `storageClassName` in a PVC.

Create a StorageClass

```
kind : StorageClass
apiVersion : storage.k8s.io / v1beta1
metadata :
  name : alicloud - disk - ssd - hangzhou - b
provisioner : alicloud / disk
parameters :
  type : cloud_ssd
  regionid : cn - hangzhou
  zoneid : cn - hangzhou - b
reclaimPolicy : Retain
```

Parameter setting:

- **provisioner:** Set this parameter to `alicloud/disk` to indicate that the StorageClass creates a cloud disk by using the provisioner plugin of Alibaba Cloud cloud disks.
- **type:** Specify the type of a cloud disk by using one of the following values: `cloud`, `cloud_efficiency`, `cloud_ssd`, and `available`. If you set this parameter to `available`, the system will cycle through `cloud_efficiency`, `cloud_ssd`, and `cloud` in order until one of them takes effect.
- **regionid:** Set the region in which you want to create a cloud disk.
- **reclaimPolicy:** Set the policy to reclaim a cloud disk. The default setting is `Delete`. You can also set this parameter to `Retain`.
- **zoneid:** Set the zone in which you want to create a cloud disk.



Note:

If you want to create cloud disks in multiple zones, you can set multiple values for the `zoneid` parameter, for example,

```
zoneid : cn - hangzhou - a , cn - hangzhou - b , cn - hangzhou - c
```

- **encrypted:** (optional) Set whether to encrypt a cloud disk. The default value is `false`. That is, a cloud disk will not be encrypted.

Create a service

```
kind : Persistent VolumeClaim
apiVersion : v1
metadata :
  name : disk - ssd
spec :
  accessModes :
    - ReadWriteOnce
  storageClassName : alicloud - disk - ssd - hangzhou - b
  resources :
    requests :
      storage : 20Gi
---
kind : Pod
apiVersion : v1
metadata :
  name : disk - pod - ssd
spec :
  containers :
    - name : disk - pod
      image : nginx
      volumeMounts :
        - name : disk - pvc
          mountPath : "/mnt"
      restartPolicy : "Never"
  volumes :
    - name : disk - pvc
      persistentVolumeClaim :
        claimName : disk - ssd
```

Default options

By default, Kubernetes clusters provide the following StorageClasses that can be used in the single-zone clusters:

- `alicloud-disk-common`, namely, a basic cloud disk.
- `alicloud-disk-efficiency`, namely, an Ultra disk.
- `alicloud-disk-ssd`, namely, an SSD disk.
- `alicloud-disk-available`: This StorageClass provides a systematic method of disk selection. Specifically, the system first attempts to create an Ultra disk. If the Ultra disks in the specified zone are sold out, the system tries to create an SSD disk. If the SSD disks are sold out, the system tries to create a basic cloud disk.

Create a multi-instance StatefulSet by using a cloud disk

We recommend that you create a multi-instance StatefulSet through volumeClaimTemplates so that you can dynamically create multiple PVCs and PVs, and connect the PVCs and PVs together.

```
apiVersion : v1
kind : Service
metadata :
  name : nginx
  labels :
    app : nginx
spec :
  ports :
    - port : 80
      name : web
  clusterIP : None
  selector :
    app : nginx
---
apiVersion : apps / v1beta2
kind : StatefulSet
metadata :
  name : web
spec :
  selector :
    matchLabels :
      app : nginx
  serviceName : "nginx"
  replicas : 2
  template :
    metadata :
      labels :
        app : nginx
    spec :
      containers :
        - name : nginx
          image : nginx
          ports :
            - containerPort : 80
              name : web
          volumeMounts :
            - name : disk-ssd
              mountPath : /data
  volumeClaimTemplates :
    - metadata :
        name : disk-ssd
      spec :
        accessModes : [ "ReadWriteOnce" ]
        storageClassName : "alicloud-disk-ssd"
        resources :
          requests :
            storage : 20Gi
```

1.13.4 Use NAS file systems of Alibaba Cloud

You can use Alibaba Cloud NAS volumes in a Kubernetes cluster of Container Service.

You can mount a NAS file system of Alibaba Cloud to a Kubernetes cluster as either of the following two types of volumes:

- [Static volumes](#)

You can use a static volume in either of the following two ways:

- Use a static volume through the flexvolume plugin.
 - Use a static volume directly.
 - Use a static volume through a Persistent Volume (PV) and a Persistent Volume Claim (PVC).
- Use a static volume through the NFS driver of Kubernetes.

- [Dynamic volumes](#)

Prerequisites

You have created a NAS file system in the NAS console and added a mount point for a Kubernetes cluster in the file system. You must make sure that the NAS file system and your cluster are in the same VPC.

Static volumes

You can use the Alibaba Cloud NAS file storage service by using the flexvolume plugin provided by Alibaba Cloud or the NFS driver of Kubernetes.

Use a static volume through the flexvolume plugin

With a flexvolume plugin, you can use an Alibaba Cloud NAS volume directly or through a PV and a PVC.



Note:

- **NAS:** a shared storage system that can provide storage services for multiple pods at the same time.
- **server:** defines the mount point of a NAS file system.
- **path:** defines the mount directory that connects to the NAS volume. You can specify a NAS sub-directory and mount it to your NAS volume. If the NAS sub-directory specified by you does not exist, the system automatically creates the NAS sub-directory and mounts it to your NAS volume.
- **vers:** defines the version number of the NFS mount protocol. NFS file system versions 3.0 and 4.0 are supported.

- **mode**: defines the access permission to a mount directory. When the mount directory is the root directory of a NAS file system, the access permission to the root directory cannot be set. If you set the mode parameter for a NAS file system that stores a large amount of data, the process of mounting the NAS file system to a cluster may take an excessive amount of time or even fail.

Use a static volume directly

Use a `nas - deploy . yaml` file to create a pod as follows:

```
apiVersion : v1
kind : Pod
metadata :
  name : " flexvolume - nas - example "
spec :
  containers :
    - name : " nginx "
      image : " nginx "
      volumeMounts :
        - name : " nas1 "
          mountPath : "/" data "
  volumes :
    - name : " nas1 "
      flexVolume :
        driver : " alicloud / nas "
        options :
          server : " 0cd8b4a576 - grs79 . cn - hangzhou . nas .
aliyuncs . com "
          path : "/" k8s "
          vers : " 3 "
```

Use a static volume through a PV and a PVC

Step 1: Create a PV

You can create a NAS volume by using a YAML file or create a NAS volume in the Alibaba Cloud Container Service console.

- Create a PV by using a YAML file.

Use a `nas - pv . yaml` file to create a PV as follows:

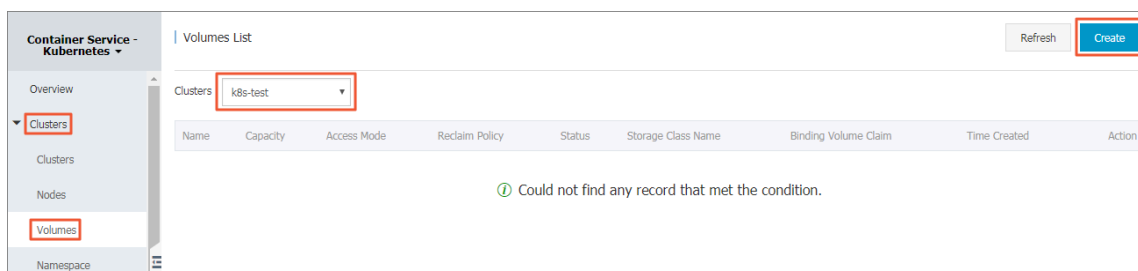
```
apiVersion : v1
kind : Persistent Volume
metadata :
  name : pv - nas
spec :
  capacity :
    storage : 5Gi
  storageClassName : nas
  accessModes :
    - ReadWriteMany
  flexVolume :
    driver : " alicloud / nas "
    options :
```

```
server : " 0cd8b4a576 - uih75 . cn - hangzhou . nas .  
aliyuncs . com "  
path : "/" k8s "
```

```
vers : " 3 "
```

- Create a NAS volume in the Container Service console.

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Kubernetes, choose Clusters > Volumes.
3. Select the target cluster from the cluster drop-down list and then click Create in the upper-right corner.

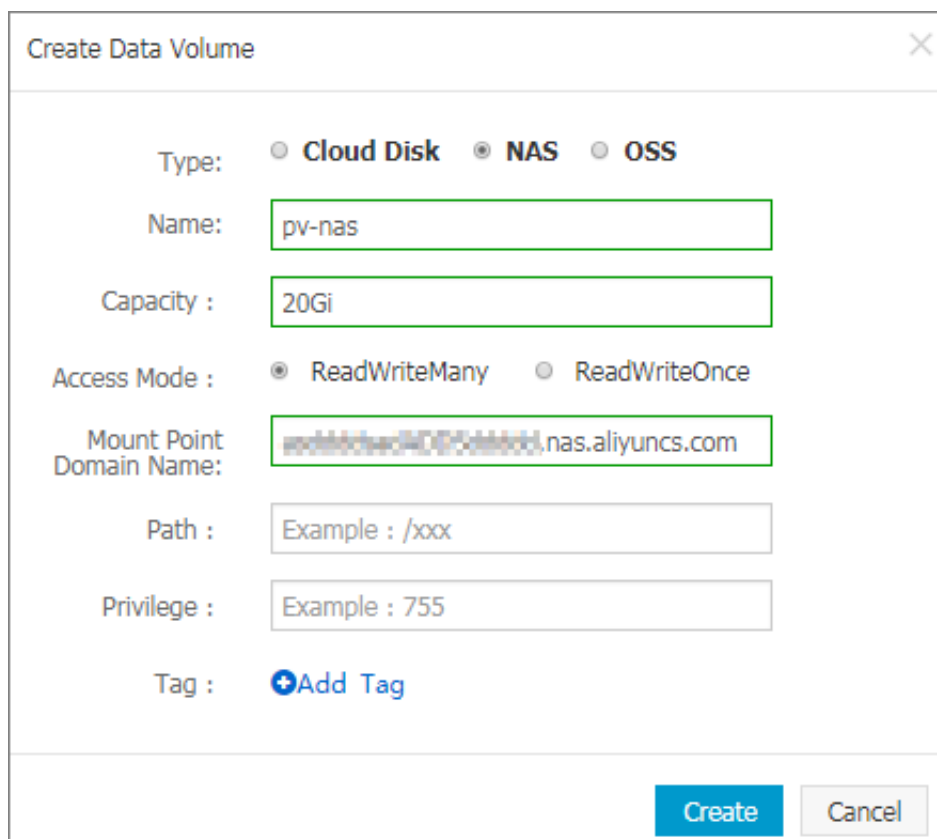


4. In the displayed dialog box, set the volume parameters.

- Storage type: NAS is selected in this example.
- Name: Customize a volume name. The volume name must be unique in the cluster. In this example, pv-nas is set as the volume name.
- Capacity: Set the volume capacity. Make sure that the volume capacity does not exceed the NAS file system capacity.
- Access Mode: By default, it is set to ReadWriteOnce.
- Mount Point Domain Name: Enter the mount address of the mount point that is used to mount the NAS file system to the Kubernetes cluster.
- Path: sub-directory under the NAS path, which starts with a forward slash (/). If you specify a sub-directory, your volume will be mounted to the sub-directory.
 - If no sub-directory exists in the root directory of a NAS file system, the system automatically creates a sub-directory by default.
 - This parameter is optional. A NAS volume is mounted to the root directory of a NAS file system by default.
- Privilege: Set the access permission to the mount directory. For example, you can set this parameter to 755, 644, or 777.
 - You can set this parameter only if you mount a NAS volume to the NAS sub-directory. This parameter cannot be set if you mount a NAS volume to the NAS root directory.

- This parameter is optional. By default, the original access permission to a NAS file system is used.

- Tag: Add tags to the volume.



The image shows a 'Create Data Volume' dialog box with the following fields and options:

- Type:** Radio buttons for ☐ Cloud Disk, ☒ NAS, and ☐ OSS.
- Name:** Text input field containing 'pv-nas'.
- Capacity :** Text input field containing '20Gi'.
- Access Mode :** Radio buttons for ☒ ReadWriteMany and ☐ ReadWriteOnce.
- Mount Point Domain Name:** Text input field containing 'xxxxxxxxxxxxx.nas.aliyuncs.com'.
- Path :** Text input field containing 'Example : /xxx'.
- Privilege :** Text input field containing 'Example : 755'.
- Tag :** A button labeled '+Add Tag'.
- Buttons:** 'Create' and 'Cancel' buttons at the bottom right.

5. Click Create.

Step 2: Create a PVC

Use a `nas - pvc . yaml` file to create a PVC as follows:

```
apiVersion : v1
kind : Persistent VolumeClaim
metadata :
  name : pvc - nas
spec :
  accessModes :
    - ReadWriteMany
  storageClassName : nas
  resources :
    requests :
      storage : 5Gi
```

Step 3: Create a pod

Use a `nas - pod . yaml` file to create a pod as follows:

```
apiVersion : v1
kind : Pod
metadata :
```

```

name : " flexvolume - nas - example "
spec :
  containers :
    - name : " nginx "
      image : " nginx "
      volumeMounts :
        - name : pvc - nas
          mountPath : "/" data "
  volumes :
    - name : pvc - nas
      persistent VolumeClaim :
        claimName : pvc - nas

```

Use the Kubernetes NFS driver



Note:

Alibaba Cloud NAS supports NFS 3.0 and NFS 4.0. You must specify a valid NFS version when you create a NAS volume.

Step 1: Create a NAS file system

Log on to the [NAS console](#) to create a NAS file system.



Note:

You must ensure that the NAS file system and your cluster are in the same region.

For example, assume that the mount point of your NAS file system is `055f84ad83 - ixxxx . cn - hangzhou . nas . aliyns . com .`

Step 2: Create a PV

You can create a NAS volume by using an orchestration template or the Alibaba Cloud Container Service console.

- Use an orchestration template to create a NAS volume

Use a `nas - pv . yaml` file to create a PV.

Run the following command to create a NAS PV:

```

root @ master # cat << EOF | kubectl apply -f -
apiVersion : v1
kind : Persistent Volume
metadata :
  name : nas
spec :
  capacity :
    storage : 8Gi
  accessMode s :
    - ReadWriteM any
  mountOptions :
    - noresvport
    - nfsvers = 3

```

```

persistent VolumeRecl aimPolicy : Retain
nfs :
  path : /
  server : 055f84ad83 - ixxxx . cn - hangzhou . nas . aliyuncs .
com
EOF

```

- Create a NAS volume in the Container Service console

For more information, see [Use a PV and a PVC](#).

Step 2: Create a PVC

Create a PVC to request to bind the PV.

```

root @ master # cat << EOF | kubectl apply -f -
apiVersion : v1
kind : Persistent VolumeClai m
metadata :
  name : nasclaim
spec :
  accessMode s :
    - ReadWriteM any
  resources :
    requests :
      storage : 8Gi
EOF

```

Step 3: Create a pod

Create an application to declare to mount and use the volume.

```

root @ master # cat << EOF | kubectl apply -f -
apiVersion : v1
kind : Pod
metadata :
  name : mypod
spec :
  containers :
    - name : myfrontend
      image : registry . aliyuncs . com / spacexnice / netdia :
latest
      volumeMoun ts :
        - mountPath : "/ var / www / html "
          name : mypd
  volumes :
    - name : mypd
      persistent VolumeClai m :
        claimName : nasclaim
EOF

```

The NAS file system is successfully mounted to the application that runs on the pod.

Dynamic volumes

To use a dynamic NAS volume, you need to manually install a driver plugin and configure a NAS mount point.

**Note:**

To dynamically generate a NAS volume is to automatically generate a directory in an existing NAS file system. This directory is defined as the target volume.

Install a plugin

```

apiVersion : storage . k8s . io / v1
kind : StorageClass
metadata :
  name : alicloud - nas
mountOptions :
- vers = 3
provisioner : alicloud / nas
reclaimPolicy : Retain
---
kind : Deployment
apiVersion : extensions / v1beta1
metadata :
  name : alicloud - nas - controller
  namespace : kube - system
spec :
  replicas : 1
  strategy :
    type : Recreate
  template :
    metadata :
      labels :
        app : alicloud - nas - controller
    spec :
      tolerations :
        - effect : NoSchedule
          operator : Exists
          key : node - role . kubernetes . io / master
        - effect : NoSchedule
          operator : Exists
          key : node . cloudprovider . kubernetes . io / uninitialized
      nodeSelector :
        node - role . kubernetes . io / master : ""
      serviceAccount : admin
      containers :
        - name : alicloud - nas - controller
          image : registry . cn - hangzhou . aliyuncs . com / acs /
alicloud - nas - controller : v3 . 1 . 0 - k8s1 . 11
          volumeMounts :
            - mountPath : / persistent volumes
              name : nfs - client - root
          env :
            - name : PROVISIONER_NAME
              value : alicloud / nas
            - name : NFS_SERVER
              value : 0cd8b4a576 - mmi32 . cn - hangzhou . nas .
alicuncs . com
            - name : NFS_PATH
              value : /
          volumes :
            - name : nfs - client - root
              flexVolume :
                driver : alicloud / nas

```

```

        options :
          path : /
          server : 0cd8b4a576 - mmi32 . cn - hangzhou . nas .
aliyuncs . com
          vers : " 3 "

```

Use the dynamic volume

```

apiVersion : apps / v1beta1
kind : StatefulSet
metadata :
  name : web
spec :
  serviceName : " nginx "
  replicas : 2
  volumeClaimTemplates :
  - metadata :
    name : html
    spec :
      accessModes :
        - ReadWriteOnce
      storageClassName : alicloud - nas
      resources :
        requests :
          storage : 2Gi
  template :
    metadata :
      labels :
        app : nginx
    spec :
      containers :
        - name : nginx
          image : nginx : alpine
          volumeMounts :
            - mountPath : "/usr/share/nginx/html/"
              name : html

```

1.13.5 Use Alibaba Cloud OSS volumes

You can use Alibaba Cloud OSS volumes in a Kubernetes cluster of Alibaba Cloud Container Service.

Specifically, you can only use static OSS volumes. Dynamic OSS volumes are not supported. You can use a static OSS volume in either of the following two ways:

- Use an OSS bucket through a volume.
- Use an OSS bucket through a Persistent Volume (PV) and a Persistent Volume Claim (PVC).

Prerequisites

You have created a bucket in the OSS console.

OSS parameter setting

- **OSS:** OSS is a shared storage system that can provide storage services to multiple pods at the same time.
- **bucket:** Only buckets can be mounted to a Kubernetes cluster. The sub-directories or files under a bucket cannot be mounted to a Kubernetes cluster.
- **url:** Specify an OSS endpoint, namely, the domain name used to mount an OSS bucket to a cluster.
- **akId:** Enter your Access Key ID.
- **akSecret:** Enter your Access Key Secret.
- **otherOpts:** Customize other parameters in the format of `- o *** - o ***`.

Notices

- If your Kubernetes cluster is created before February 6th, 2018, [Install the plug-in](#) before using a volume. Before you can use the OSS volume, you must first create a secret and then enter your Access Key information into the secret when you deploy the flexvolume service.
- If you use the flexvolume component of an earlier version, we recommend that you upgrade it to the latest version.



Note:

When you upgrade a Kubernetes cluster or restart a kubelet, an OSS volume mounted by the flexvolume component of an earlier version will cause the OSSFS driver to restart. To solve the exception caused by this event, you must recreate the pod that used the OSS volume to remount the OSS volume to the cluster. However, you can solve this issue more easily by upgrading the flexvolume component to the latest version.

Use a static OSS volume

Use an OSS bucket through a volume

Use a `oss - deploy . yaml` file to create a pod.

```
apiVersion : extensions / v1beta1
kind : Deployment
metadata :
  name : nginx - oss - deploy
spec :
  replicas : 1
  template :
    metadata :
```

```

labels :
  app : nginx
spec :
  containers :
    - name : nginx - flexvolume - oss
      image : nginx
      volumeMounts :
        - name : "oss1"
          mountPath : "/data"
      livenessProbe :
        exec :
          command :
            - sh
            - -c
            - cd /data
          initialDelaySeconds : 30
          periodSeconds : 30
      volumes :
        - name : "oss1"
          flexVolume :
            driver : "alicloud / oss"
            options :
              bucket : "docker"
              url : "oss-cn-hangzhou.aliyuncs.com"
              akId : ***
              akSecret : ***
              otherOpts : "-o max_stat_cache_size = 0 -o
allow_other "

```

Use a PV and a PVC

Step 1: Create a PV

You can create a PV by using a YAML file or the Container Service console.

Use a YAML file to create a PV

Use a `oss - pv . yaml` file to create a PV as follows:

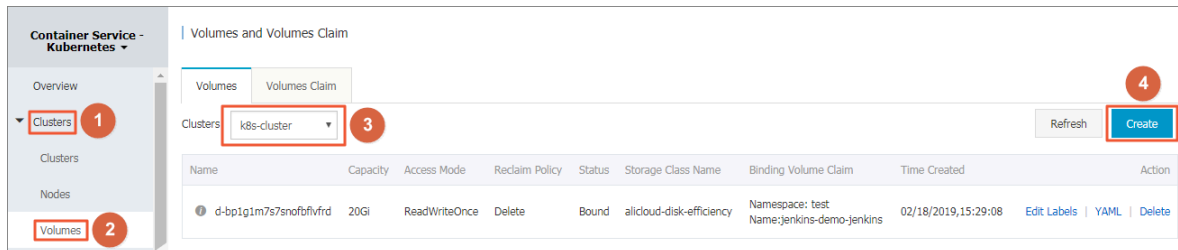
```

apiVersion : v1
kind : PersistentVolume
metadata :
  name : pv - oss
spec :
  capacity :
    storage : 5Gi
  accessModes :
    - ReadWriteMany
  storageClassName : oss
  flexVolume :
    driver : "alicloud / oss"
    options :
      bucket : "docker"
      url : "oss-cn-hangzhou.aliyuncs.com"
      akId : ***
      akSecret : ***
      otherOpts : "-o max_stat_cache_size = 0 -o allow_othe
r "

```

Create an OSS volume in the Container Service console

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Kubernetes, choose Clusters > Volumes.
3. Select the target cluster from the cluster drop-down list and then click Create in the upper-right corner.



4. In the displayed dialog box, set the volume parameters.
 - Storage type: OSS is selected in this example.
 - Name: Customize a volume name. The volume name must be unique in the cluster. In this example, pv-oss is set as the volume name.
 - Capacity: Set the volume capacity.
 - Access Mode: By default, it is set to ReadWriteMany.
 - AccessKey ID and AccessKey Secret: Use these two parameters to specify the Access Key used to access OSS.
 - Bucket ID: Select an OSS bucket name. Click Select Bucket. In the displayed dialog box, select the target bucket and click Select.
 - Access Domain Name. If the selected bucket and the cluster ECS instances are in different regions, you need to select Internet. If they are in the same region, your choice is dependent on your cluster network type. If your cluster uses a VPC,

you need to select VPC; if your cluster uses a classic network, you need to select Intranet.

- Tag: Add tags to the volume.

Create Data Volume

Type: ☐ Cloud Disk ☐ NAS ☒ OSS

Name:

Capacity :

Access Mode : ☒ ReadWriteMany

Access Key ID:

Access Key Secret:

Optional Parameters:

For the formats of other parameters, refer to this document. Example: -o allow_other -o default_permission=666 -onoxattr

Bucket ID: [Select Bucket](#)

Access Domain Name: ☐ Intranet ☐ Internet ☒ VPC ?

Tag : [+Add Tag](#)

Create

Cancel

5. Click Create.

Step 2: Create a PVC

Use a `oss - pvc . yaml` file to create a PVC as follows:

```
kind : Persistent VolumeClaim
apiVersion : v1
metadata :
  name : pvc - oss
spec :
  storageClassName : oss
  accessModes :
    - ReadWriteMany
  resources :
    requests :
```

```
storage : 5Gi
```

Step 3: Create a pod

Use a `oss - pod . yaml` file to create a pod.

```
apiVersion : v1
kind : Pod
metadata :
  name : " flexvolume - oss - example "
spec :
  containers :
    - name : " nginx "
      image : " nginx "
      volumeMounts :
        - name : pvc - oss
          mountPath : "/" data "
      livenessProbe :
        exec :
          command :
            - sh
            - - c
            - cd / data
          initialDelaySeconds : 30
          periodSeconds : 30
  volumes :
    - name : pvc - oss
      persistentVolumeClaim :
        claimName : pvc - oss
```

1.13.6 Create a Persistent Volume Claim

You can create a Persistent Volume Claim (PVC) by using the Container Service console.

Prerequisites

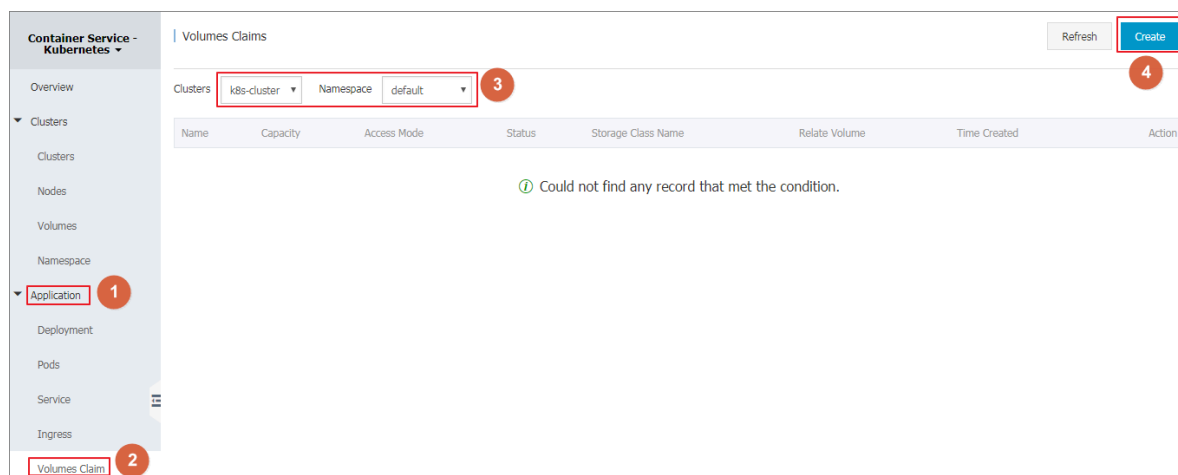
- You have created a Kubernetes cluster. For more information, see [Create a Kubernetes cluster](#).
- You have created a volume. In this example, use a cloud disk to create a cloud storage volume. For more information, see [Use Alibaba Cloud cloud disk volumes](#).

By default, the storage claim is bound to the storage volume depending on the label `alicloud - pvname`. When the data volume is created by using the Container Service console, the storage volume is labeled by default. If the storage volume label does not exist, you must add a label before you select to bound this storage volume.

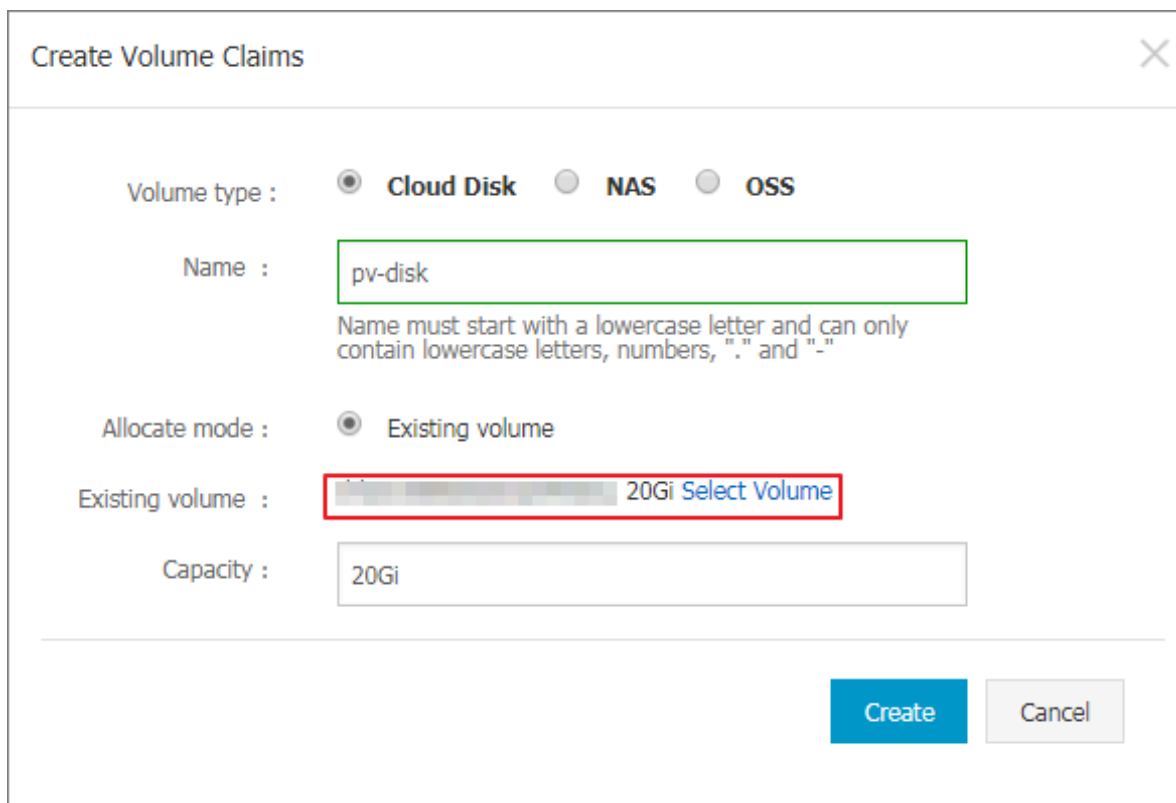
Context

Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Application > Volumes Claim in the left-side navigation pane to enter the Volumes Claims list page.
3. Select the target cluster and namespace, and click Create in the upper-right corner.



4. Complete the configurations in the Create Volume Claim dialog box, and click Create.



Create Volume Claims

Volume type : ☒ Cloud Disk ☐ NAS ☐ OSS

Name :
Name must start with a lowercase letter and can only contain lowercase letters, numbers, "." and "-"

Allocate mode : ☒ Existing volume ☐ New volume

Existing volume : [Select Volume](#)

Capacity :

[Create](#) [Cancel](#)

- Volume claim type: Consistent with volume, including cloud disk, NAS, and OSS types.
- Name: Enter the storage volume claim name.
- Distribution mode: Currently, only existing storage volumes are supported.
- Existing storage volume: Select to bind the storage volume of this type.
- Total: Claim usage, cannot be greater than the total amount of storage volumes.

**Note:**

If a storage volume already exists in your cluster and is not used, but cannot be found in Select Existing Storage Volume, maybe the `alicloud - pvname` label is not defined.

If you cannot find an available storage volume, you can click **Clusters > Volumes** in the left-side navigation pane. Find the target storage volume, click **Label Management** on the right. Add the corresponding label `alicloud - pvname`, the

value is the name of the storage volume. The cloud storage volume defaults to the cloud disk ID as the name of the storage volume.

Name	Value
alicloud-pvname	d-bp1-7350t00c7-emix3iv0e
failure-domain.beta.kubernetes.io/zone	cn-hangzhou-g
failure-domain.beta.kubernetes.io/region	cn-hangzhou

- Return to the Volumes Claims list, you can see that the newly created storage claim appears in the list.

1.13.7 Use a persistent volume claim

In the Container Service console, use an image or a template to deploy an application, so that you can use a persistent volume claim (PVC). In this example, an image is used to create an application. If you want to use a PVC through a template, see [Use Alibaba Cloud cloud disk volumes](#).

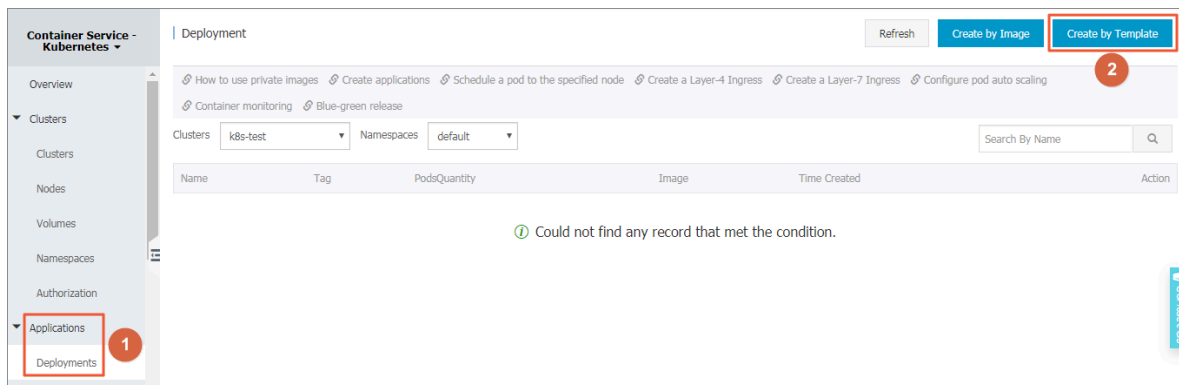
Prerequisites

- A Kubernetes cluster is created. For more information, see [Create a Kubernetes cluster](#).
- A PVC is created. In this topic, a pvc-disk PVC is created by using a disk. For more information, see [Create a persistent volume claim](#).

Procedure

- Log on to the [Container Service console](#).

2. In the left-side navigation pane under Container Service-Kubernetes, choose Applications > Deployments. Then, click Create by Image in the upper-right corner.



3. Set the application name and select the target cluster and the namespace. Then click Next.
4. Select an image, set the data volume of disk type, and then click Next. Disk, NAS, and OSS types are available. In this example, an existing disk PVC is used.
5. Set services for the test-nginx application, and then click Create.
6. In the left-side navigation pane under Container Service-Kubernetes, choose Applications > Pods. Then, click Details on the right of the pod to which the test-nginx application belongs.
7. Click Volumes to verify that the target pod is associated with the pvc-disk.

1.14 Log management

1.14.1 Application log management

A Kubernetes cluster that runs on Alibaba Cloud Container Service provides you with multiple methods to manage application logs.

- Following the instructions of [Use Log Service to collect Kubernetes cluster logs](#), you can make the best use of the functions provided by Alibaba Cloud Log Service, such as log statistics and analysis.

- With [Log-pilot](#), an open source project provided by Alibaba Cloud Container Service, and [A solution to log collection problems of Kubernetes clusters by using log-pilot, Elasticsearch, and Kibana](#), you can easily build your own application log clusters.

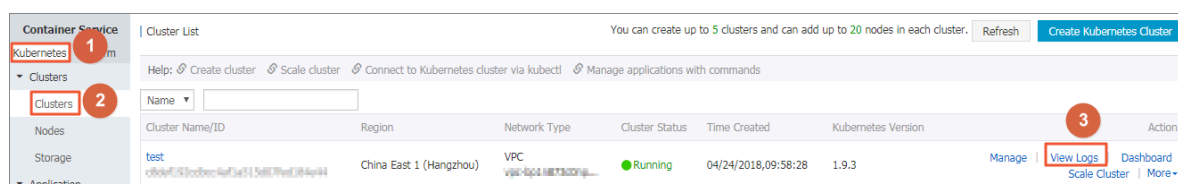
1.14.2 View cluster logs

Context

You can view the cluster operation logs by using the simple log service of Container Service.

Procedure

- Log on to the [Container Service console](#).
- Under Kubernetes, click Clusters in the left-side navigation pane.
- Click View Logs at the right of the cluster.



View the cluster operation information.

Cluster Logs: test Back to Cluster List		Refresh
Detailed resource deployment logs: Stack Events		
Time	Information	
04/24/2018,13:55:38	c8def192cdbcc4af1a515d07fed184e44 Start to client.DescribeTemplate	
04/24/2018,13:55:35	c8def192cdbcc4af1a515d07fed184e44 Start describeStackInfo	
04/24/2018,11:27:38	c8def192cdbcc4af1a515d07fed184e44 Start to client.DescribeTemplate	
04/24/2018,11:27:36	c8def192cdbcc4af1a515d07fed184e44 Start describeStackInfo	
04/24/2018,11:27:08	c8def192cdbcc4af1a515d07fed184e44 Start to client.DescribeTemplate	
04/24/2018,11:27:06	c8def192cdbcc4af1a515d07fed184e44 Start describeStackInfo	
04/24/2018,11:26:55	c8def192cdbcc4af1a515d07fed184e44 Start to client.DescribeTemplate	
04/24/2018,11:26:54	c8def192cdbcc4af1a515d07fed184e44 Start describeStackInfo	
04/24/2018,11:25:02	c8def192cdbcc4af1a515d07fed184e44 Start to client.DescribeTemplate	
04/24/2018,11:25:00	c8def192cdbcc4af1a515d07fed184e44 Start describeStackInfo	
04/24/2018,10:16:42	c8def192cdbcc4af1a515d07fed184e44 Set up k8s DNS configuration successfully	
04/24/2018,10:15:36	c8def192cdbcc4af1a515d07fed184e44 Start describeStackInfo	
04/24/2018,10:15:36	c8def192cdbcc4af1a515d07fed184e44 Stack CREATE completed successfully:	
04/24/2018,10:15:34	c8def192cdbcc4af1a515d07fed184e44 Start describeStackInfo	

1.14.3 Use Log Service to collect Kubernetes cluster logs

Log Service is integrated with Kubernetes clusters of Alibaba Cloud Container Service. You can enable Log Service when creating a cluster to quickly collect

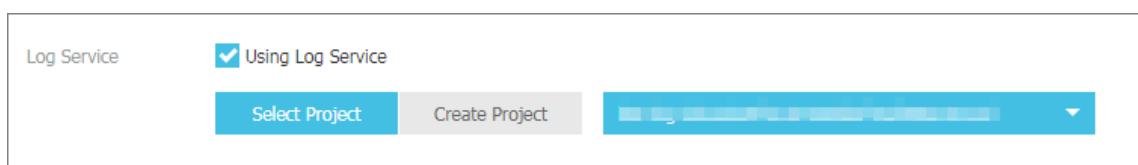
container logs for the Kubernetes cluster, such as the standard output of the container and text files of the container.

Enable Log Service when creating a Kubernetes cluster

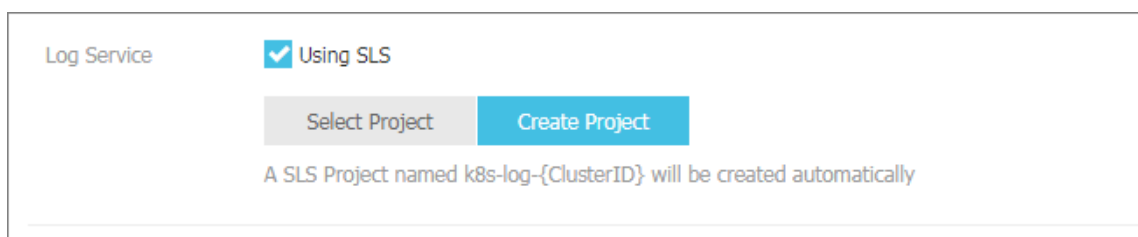
If you have not created any Kubernetes clusters, follow steps in this section to enable Log Service:

1. Log on to the [Container Service console](#).
2. Click Clusters in the left-side navigation pane and click Create Kubernetes Cluster in the upper-right corner.
3. For how to configure a cluster on the creation page, see [Create a Kubernetes cluster](#).
4. Drag to the bottom of the page and select the Using Log Service check box. The log plug-in will be installed in the newly created Kubernetes cluster.
5. When you select the Using Log Service check box, project options are displayed. A project is the unit in Log Service to manage logs. For more information about projects, see [Project](#) . Currently, two ways of using a project are available:

- Select an existing project to manage collected logs.

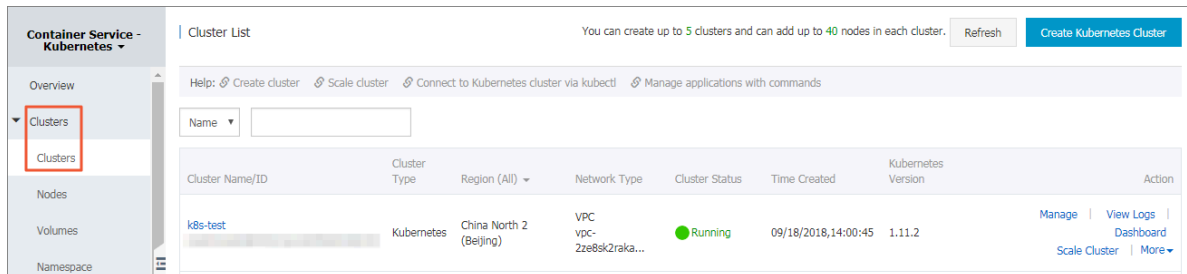


- The system automatically creates a new project to manage collected logs. The project is automatically named `k8s - log -{ ClusterID }`, where ClusterID represents the unique identifier of the created Kubernetes cluster.



- After you complete the configurations, click **Create** in the upper-right corner. In the displayed dialog box, click **OK**.

After the cluster creation is completed, the newly created Kubernetes cluster is displayed on the cluster list page.



Manually install Log Service components in a created Kubernetes cluster

If you have created a Kubernetes cluster, following instructions in this section to use Log Service:

- Log Service components are not installed. Manually install the components.
- Log Service components are installed but in an earlier version. Upgrade the components. If you do not upgrade the components, you can only use the Log Service console or custom resource definition (CRD) to configure log collection.

Check the Log Service component version

- Configure the local kubeconfig to connect to the Kubernetes cluster through `kubectl`.

For information about the configuration, see [Connect to a Kubernetes cluster by using kubectl](#).

- Run the following command to fast determine whether an upgrade or migration operation is required:

```
$ kubectl describe daemonsets -n kube-system logtail -
ds | grep ALICLOUD_LOG_DOCKER_ENV_CONFIG
```

- If `ALICLOUD_LOG_DOCKER_ENV_CONFIG : true` is output, the components can be used directly without requiring upgrade or migration.
 - If other results are output, check the components further.
- Run the following command to determine whether Helm is used to install the components.

```
$ helm get alibaba-log-controller | grep CHART
```

```
CHART : alibaba - cloud - log - 0 . 1 . 1
```

- 0.1.1 in the output indicates the version of the Log Service components. Please use the version of 0.1.1 and later. If the version is too early, please see [Upgrade Log Service components](#) to upgrade the components. If you have used Helm to install the components of a valid version, you can skip next steps.
- If no results are output, the components are installed by using Helm. But the DaemonSet installation method might be used. Follow the next step to check further.

4. DaemonSet can be an old one or a new one:

```
$ kubectl get daemonsets -n kube-system logtail
```

- If no result is output or `No resources found` is output, the Log Service components are not installed. For information about the installation method, see [Manually install Log Service components](#).
- If the correct result is output, an old DaemonSet is used to install the components which require upgrade. For information about upgrading the components, see [Upgrade Log Service components](#).

Manually install Log Service components

1. Configure the local kubeconfig to connect to the Kubernetes cluster through kubectl.

For information about the configuration, see [Connect to a Kubernetes cluster by using kubectl](#).

2. Replace a parameter and run the following command.

Replace `${your_k8s_cluster_id}` in the following command with your Kubernetes cluster ID, and then run the command.

```
wget https://acs-logging.oss-cn-hangzhou.aliyuncs.com/alicloud-k8s-log-installer.sh -O alicloud-k8s-log-installer.sh ; chmod 744 ./alicloud-k8s-log-installer.sh ; ./alicloud-k8s-log-installer.sh --cluster-id ${your_k8s_cluster_id} --ali-uid ${
```

```
your_ali_u id } -- region - id ${ your_k8s_c luster_reg ion_id
}
```

Parameter descriptions:

- **your_k8s_cluster_id**: You Kubernetes cluster ID.
- **your_ali_uid**: You account ID of Alibaba Cloud, which can be viewed in the user info.
- **your_k8s_cluster_region_id**: The region in which you Kubernetes cluster resides, which can be found in [Regions and zones](#). For example, if the cluster resides in Hangzhou, the value of this parameter cn-hangzhou.

Installation example

```
[ root @ iZbp ***** biaZ ~]# wget https://acs-logging.oss
-cn-hangzhou.aliyuncs.com/alicloud-k8s-log-installer
.sh -O alicloud-k8s-log-installer.sh; chmod 744 ./
alicloud-k8s-log-installer.sh; ./alicloud-k8s-log-
installer.sh --cluster-id c77a*****0106 --ali
-uid 19*****19 --region-id cn-hangzhou
--2018-09-28 15:25:33 -- https://acs-logging.oss-
cn-hangzhou.aliyuncs.com/alicloud-k8s-log-installer.
sh
Resolving acs-logging.oss-cn-hangzhou.aliyuncs.com ...
118.31.219.217, 118.31.219.206
Connecting to acs-logging.oss-cn-hangzhou.aliyuncs.
com | 118.31.219.217 | : 443 ... connected.
HTTP request sent, awaiting response ... 200 OK
Length: 2273 (2.2K) [text/x-sh]
Saving to: 'alicloud-k8s-log-installer.sh'

alicloud-k8s-log-installer.sh                                100
%[=====
  2.22K --. - KB / s      in 0s

2018-09-28 15:25:33 (13.5 MB / s) - 'alicloud-k8s
-log-installer.sh' saved [2273 / 2273]

--2018-09-28 15:25:33 -- http://logtail-release-
cn-hangzhou.oss-cn-hangzhou.aliyuncs.com/kubernetes/
alibaba-cloud-log.tgz
Resolving logtail-release-cn-hangzhou.oss-cn-hangzhou.
aliyuncs.com ... 118.31.219.49
Connecting to logtail-release-cn-hangzhou.oss-cn-
hangzhou.aliyuncs.com | 118.31.219.49 | : 80 ... connected
.
HTTP request sent, awaiting response ... 200 OK
Length: 2754 (2.7K) [application/x-gzip]
Saving to: 'alibaba-cloud-log.tgz'

alibaba-cloud-log.tgz                                        100
%[=====
  2.69K --. - KB / s      in 0s

2018-09-28 15:25:34 (79.6 MB / s) - 'alibaba -
cloud-log.tgz' saved [2754 / 2754]
```

```
[ INFO ] your k8s is using project : k8s - log -
c77a92ec5a 3ce4e64a1b f13bde1820 106
NAME : alibaba - log - controller
LAST DEPLOYED : Fri Sep 28 15 : 25 : 34 2018
NAMESPACE : default
STATUS : DEPLOYED

RESOURCES :
==> v1beta1 / CustomResourceDefinition
NAME AGE
aliyunlogconfig.alibabacloud.com 0s

==> v1beta1 / ClusterRole
alibaba - log - controller 0s

==> v1beta1 / ClusterRoleBinding
NAME AGE
alibaba - log - controller 0s

==> v1beta1 / DaemonSet
NAME DESIRED CURRENT READY UP - TO - DATE
AVAILABLE NODE SELECTOR AGE
logtail - ds 2 2 0 2 0
< none > 0s

==> v1beta1 / Deployment
NAME DESIRED CURRENT UP - TO - DATE
AVAILABLE AGE
alibaba - log - controller 1 1 1 0
0s

==> v1 / Pod ( related )
NAME READY STATUS
RESTARTS AGE
logtail - ds - 6v979 0 / 1 ContainerC
reating 0 0s
logtail - ds - 7ccqv 0 / 1 ContainerC
reating 0 0s
alibaba - log - controller - 84d8b6b8cf - nkrkx 0 / 1
ContainerC reating 0 0s

==> v1 / ServiceAccount
NAME SECRETS AGE
alibaba - log - controller 1 0s

[ SUCCESS ] install helm package : alibaba - log - controller
success .
```

Upgrade Log Service components

If you have installed Log Service components of an early version through Helm or DaemonSet, upgrade or migrate the components as follows.



Note:

To perform the following operations, first log on to the master node of your Kubernetes cluster of Alibaba Cloud Container Service. For information about the logon method, see [Connect to a Kubernetes cluster by using kubectl](#).

Use Helm to upgrade Log Service components (recommended)

1. Run the following command to download the latest Helm package of Log Service components:

```
wget http://logtail-release-cn-hangzhou.oss-cn-hangzhou.aliyuncs.com/kubernetes/alibaba-cloud-log.tgz -O alibaba-cloud-log.tgz
```

2. Upgrade the components by using `helm upgrade`. The command is as follows:

```
helm get values alibaba-log-controller --all > values.yaml && helm upgrade alibaba-log-controller alibaba-cloud-log.tgz --recreate-pods -f values.yaml
```

Use DaemonSet to upgrade Log Service components

You can upgrade Log Service components by modifying the DaemonSet template.

If your image account is `acs`, upgrade the image tag to the latest version that can be viewed in [Container Registry](#). If your image account is `acs`, upgrade the image tag to the latest version that can be viewed in [Container Registry](#).



Note:

- If upgrading the tag has not enabled a rolling update of Logtail, you must manually remove the Logtail pod to trigger a Logtail update.
- You need to check whether Logtail runs on all nodes, including Master nodes. If Logtail does not run on all nodes, you must set [tolerations](#) for Logtail.

```
toleration s :  
- operator : "Exists "
```

For more information, see [Latest Helm package configurations](#).

DaemonSet migrate

This upgrade method is applicable to the situation that you find the components are installed through the old DaemonSet when you check the Log Service component version. This method does not support configuring Log Service in Container Service. You can upgrade the components as follows:

1. At the end of the installation command, add a parameter which is the name of the project of Log Service used by your Kubernetes cluster.

For example, if the project name is k8s-log-demo and the cluster ID is c12ba2028cxxxxxxxxxx6939f0b, then the installation command is:

```
wget https://acs-logging.oss-cn-hangzhou.aliyuncs.com/alicloud-k8s-log-installer.sh -O alicloud-k8s-log-installer.sh; chmod 744 ./alicloud-k8s-log-installer.sh; ./alicloud-k8s-log-installer.sh --cluster-id c12ba2028cxxxxxxxxxx6939f0b --ali-uid 19***** 19 --region-id cn-hangzhou --log-project k8s-log-demo
```

2. After you complete the installation, log on the [Log Service console](#).
3. After you complete the installation, log on the [Log Service console](#).
4. In the Log service console, apply the history collection configuration of the project and Logstore to the new machine group `k8s - group - ${your_k8s_cluster_id}`.
5. After one minute, the history collection configuration is unbound from the history machine group.
6. When log collection is normal, you can delete the previously installed Logtail DaemonSet.



Note:

During the upgrade, some logs are duplicated. The CRD configuration management takes effect only for the configuration created by using CRD. The history configuration does not support the CRD management because the history configuration is created by using the non-CRD mode.

Configure Log Service when creating an application

Container Service allows you to configure Log Service to collect container logs when creating an application. Currently, you can use the console or a YAML template to create an application.

Create an application by using the console

1. Log on to the [Container Service console](#).
2. Under the Kubernetes menu, click Application > Deployment in the left-side navigation pane, and then click Create by Image in the upper-right corner.

3. Configure Name, Cluster, Namespace, Replicas, and Type, and then click Next.

Create Application

Basic Information Container Advanced Done

Name:

The name should be 1-64 characters long, and can contain numbers, lower case English letters and hyphens, but cannot start with a hyphen.

Cluster:

Namespace :

Replicas:

Type:

Back Next

4. On the Container page, select the Tomcat image and configure container log collection.

The following describes only configurations related to Log Service. For information about other application configurations, see [Create a deployment application by using an image](#).

5. Configure Log Service. Click the + sign to create a configuration which consists of a Logstore name and a log path.

- Logstore name: Specify a Logstore in which collected logs are stored. If your specified Logstore does not exist, the system automatically creates the Logstore in the project of Log Service with which the cluster is associated .



Note:

A Logstore name cannot contain underscores (_). You can use hyphens (-) instead.

- Log path: Specify the path where logs to be collected reside. For example, use `/usr / local / tomcat / logs / catalina . *. log` to collect text logs of tomcat.



Note:

If you specify the log path as stdout, the container standard output and standard error output will be collected.

Each configuration is automatically created as a configuration for the corresponding Logstore. By default, the simple mode (by row) is used to collect

logs. To use more collection modes, log on go to the Log Service console, and enter the corresponding project (prefixed with k8s-log by default) and Logstore to modify the configuration.

Log Service: Note: please ensure that cluster has deployed log plug-ins.

Configuration

Log Store	Log path in the container (can be set to stdout)
access	/usr/local/tomcat/logs/catalina.*.log
catalina	stdout

6. Custom tag. Click the + sign to create a new custom tag. Each custom tag is a key-value pair which will be added to collected logs. You can use a custom tag to mark container logs. For example, you can create a custom tag as a version number.

Custom Tag

Name Of Tag	Value Of Tag
release	1.0.0

7. When you complete all the configurations of the container, click Next in the upper-right corner to perform further configurations. For more information, see [Create a deployment application by using an image](#).

Create an application by using a YAML template

1. Log on to the [Container Service console](#).
2. Under the Kubernetes menu, click Application > Deployment in the left-side navigation pane, and then click Create by Template in the upper-right corner.
3. The syntax of the YAML template is the same as the Kubernetes syntax. To specify the collection configuration for the container, you need to use env to add collection configuration and custom tag for the container, and create corresponding volumeMounts and volumes. The following is a simple pod example:

```
apiVersion : v1
kind : Pod
metadata :
  name : my - demo
spec :
  containers :
    - name : my - demo - app
```

```

    image : ' registry . cn - hangzhou . aliyuncs . com / log -
service / docker - log - test : latest '
    env :
    ##### Configure environmen t variables #####
    - name : aliyun_log s_log - stdout
      value : stdout
    - name : aliyun_log s_log - varlog
      value : / var / log /*. log
    - name : aliyun_log s_mytag1_t ags
      value : tag1 = v1
    #####
    ##### Configure vulume mount #####
    volumeMoun ts :
    - name : volumn - sls - mydemo
      mountPath : / var / log
    volumes :
    - name : volumn - sls - mydemo
      emptyDir : {}
    #####

```

- Configure three parts in order based on your needs.
- In the first part, use environment variables to create your collection configurat ion and custom tag. All environment variables related to configuration are prefixed with `aliyun_log s_`.
- Rules for creating the collection configuration are as follows:

```

- name : aliyun_log s_ { Logstore name }
  value : { log path }

```

In the example, create two collection configurations. The `aliyun_log s_log - stdout` env creates a configuration that contains a Logstore named log-stdout and the log path of stdout. The standard output of the container is collected and stored to the Logstore named log-stdout.



Note:

A Logstore name cannot contain underscores (_). You can use hyphens (-) instead.

- Rules for creating a custom tag are as follows:

```

- name : aliyun_log s_ { a name without ' _ ' } _tags

```

```
value : { Tag   name }={ Tag   value }
```

After a tag is configured, when logs of the container are collected, fields corresponding to the tag are automatically attached to Log Service.

- If you specify a non-stdout log path in your collection configuration, create corresponding volumeMounts in this part.

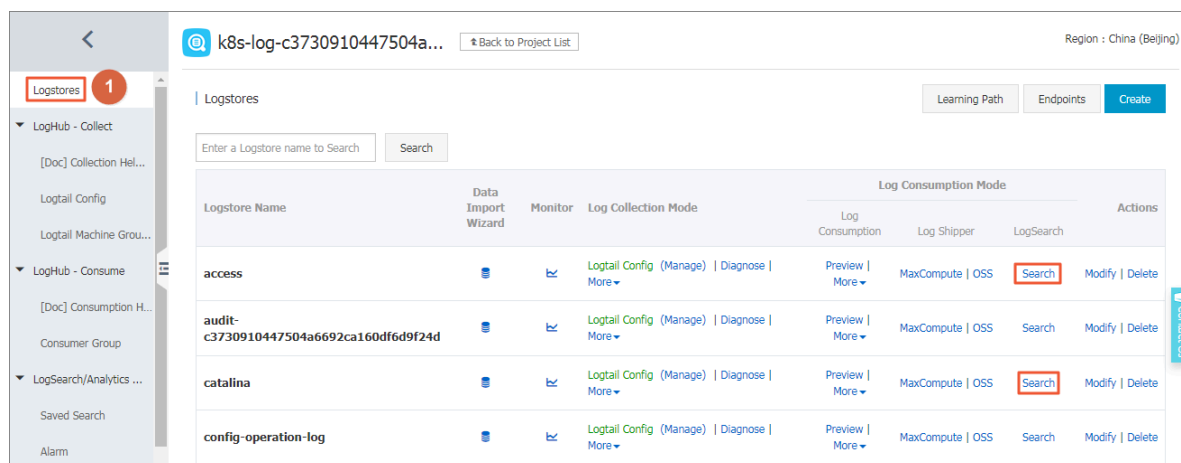
In the example, the `/ var / log /*. log` log path is added to the collection configuration, therefore, the `/ var / log` volumeMounts is added.

4. When you complete a YAML template, click **DEPLOY** to deliver the configurations in the template to the Kubernetes cluster to execute.

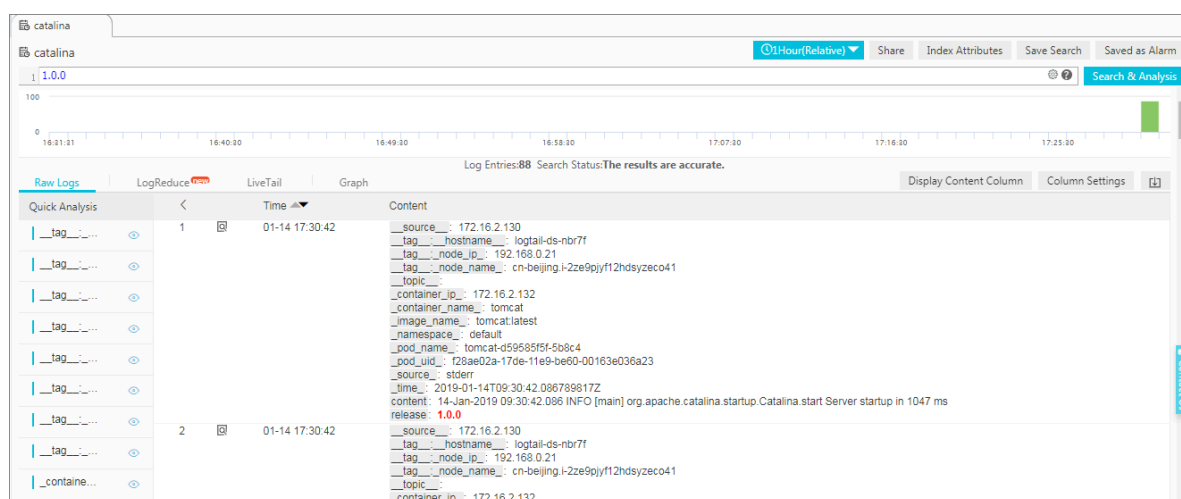
View logs

In this example, view logs of the tomcat application created in the console. After you complete the application configuration, logs of the tomcat application are collected and stored to Log Service. You can view your logs as follows:

1. Log on to the [Log Service console](#).
2. Log on to the [Log Service console](#).
3. In the console, select the project (k8s-log-{Kubernetes cluster ID} by default) corresponding to the Kubernetes cluster.
4. In the Logstore list, locate the Logstore specified in your configuration and click **Search**.



5. In this example, on the log search page, you can view the standard output logs of the tomcat application and text logs in the container, and you can find your custom tag is attached to log fields.



More information

1. By default, the system use the simple mode to collect your data, that is, to collect data by row without parsing. To perform more complex configurations, see the following Log Service documents and log on to the Log Service console to modify configurations.
 - [Container text logs](#)
 - [Container stdout](#)
2. Currently, Log Service uses plug-ins to collect the standard output logs of containers. You can configure more plug-ins to process collected logs further, such as to filter and extract fields.
3. In addition to configuring log collection through the console, you can also directly collect logs of the Kubernetes cluster through the CRD configuration. For more information, see [Configure Kubernetes log collection on CRD](#).
4. For troubleshooting exceptions, see [Troubleshoot collection errors](#).

1.14.4 A solution to log collection problems of Kubernetes clusters by using log-pilot, Elasticsearch, and Kibana

Requirements for logs of distributed Kubernetes clusters always bother developers. This is mainly because of the characteristics of containers and the defects of log collection tools.

- Characteristics of containers:
 - Many collection targets: The characteristics of containers cause the number of collection targets is large, which requires to collect the container logs and container stdout. Currently, no good tool can collect file logs from containers dynamically. Different data sources have different collection softwares. However, no one-stop collection tool exists.
 - Difficulty caused by auto scaling: Kubernetes clusters are in the distributed mode. The auto scaling of services and the environment brings great difficulty to log collection. You cannot configure the log collection path in advance, the same as what you do in the traditional virtual machine (VM) environment. The dynamic collection and data integrity are great challenges.
- Defects of current log collection tools:
 - Lack the capability to dynamically configure log collection: The current log collection tools require you to manually configure the log collection method and path in advance. These tools cannot dynamically configure the log collection because they cannot automatically detect the lifecycle changes or dynamic migration of containers.
 - Log collection problems such as logs are duplicate or lost: Some of the current log collection tools collect logs by using the tail method. Logs may be lost in this way. For example, the application is writing logs when the log collection tool is being restarted. Logs written during this period may be lost. Generally, the conservative solution is to collect logs of 1 MB or 2 MB previous to the current log by default. However, this may cause the duplicate log collection.
 - Log sources without clear marks: An application may have multiple containers that output the same application logs. After all the application logs are collected to a unified log storage backend, you cannot know a log is generated on which application container of which node when querying logs.

This document introduces log-pilot, a tool to collect Docker logs, and uses the tool together with Elasticsearch and Kibana to provide a one-stop solution to log collection problems in the Kubernetes environment.

Introduction on log-pilot

Log-pilot is an intelligent tool used to collect container logs, which not only collects container logs and outputs these logs to multiple types of log storage backends

efficiently and conveniently, but also dynamically discovers and collects log files from containers.

Log-pilot uses declarative configuration to manage container events strongly and obtain the stdout and file logs of containers, which solves the problem of auto scaling. Besides, log-pilot has the functions of automatic discovery, maintenance of checkpoint and handle, and automatic tagging for log data, which effectively deals with the problems such as dynamic configuration, duplicate logs, lost logs, and log source marking.

Currently, log-pilot is completely open-source in GitHub. The project address is <https://github.com/AliyunContainerService/log-pilot>. You can know more implementation principles about it.

Declarative configuration for container logs

Log-pilot supports managing container events, can dynamically listen to the event changes of containers, parse the changes according to the container labels, generate the configuration file of log collection, and then provide the file to collection plug-in to collect logs.

For Kubernetes clusters, log-pilot can dynamically generate the configuration file of log collection according to the environment variable `aliyun_log s_ $ name = $ path`. This environment variable contains the following two variables:

- One variable is `$name`, a custom string which indicates different meanings in different scenarios. In this scenario, `$name` indicates index when collecting logs to Elasticsearch.
- The other is `$path` which supports two input modes, stdout and paths of log files within containers, respectively corresponding to the standard output of logs and log files within containers.
 - Stdout indicates to collect standard output logs from containers. In this example, to collect Tomcat container logs, configure the label `aliyun . logs . catalina = stdout` to collect standard output logs of Tomcat.
 - The path of a log file within a container also supports wildcards. To collect logs within the Tomcat container, configure the environment variable `aliyun_log s_access = / usr / local / tomcat / logs /* . log`. To not use the keyword `aliyun`, you can use the environment variable `PILOT_LOG_PREFIX`,

which is also provided by log-pilot, to specify the prefix of your declarative log configuration. For example, `PILOT_LOG_PREFIX : " aliyun , custom "`.

Besides, log-pilot supports multiple log parsing formats, including none, JSON, CSV, Nginx, apache2, and regexp. You can use the `aliyun_log s_ $ name_format =<format >` label to tell log-pilot to use what format to parse logs when collecting logs.

Log-pilot also supports custom tags. If you configure `aliyun_log s_ $ name_tags =" K1 = V1 , K2 = V2 "` in the environment variable, K1=V1 and K2=V2 are collected to log output of the container during the log collection. Custom tags help you tag the log generation environment for convenient statistics, routing, and filter of logs.

Log collection mode

In this document, deploy a log-pilot on each machine and collect all the Docker application logs from the machines.

Compared with deploying a logging container on each pod, the most obvious advantage of this solution is less occupied resources. The larger the cluster scale is, the more obvious the advantage is. This solution is also recommended in the community.

Prerequisites

You have activated Container Service and created a Kubernetes cluster. In this example, create a Kubernetes cluster in China East 1 (Hangzhou).

Step 1 Deploy Elasticsearch

1. Connect to your Kubernetes cluster. For more information, see [#unique_63](#) or [#unique_170](#).
2. Deploy the resource object related to Elasticsearch first. Then, enter the following orchestration template. This orchestration template includes an elasticsearch-api

service, an elasticsearch-discovery service, and a status set of Elasticsearch. All of these objects are deployed under the namespace kube-system.

```
kubectl apply -f https://acs-logging.oss-cn-hangzhou.aliyuncs.com/elasticsearch.yml
```

3. After the successful deployment, corresponding objects are under the namespace kube-system. Run the following commands to check the running status:

```
$ kubectl get svc,StatefulSet -n=kube-system
NAME      TYPE      CLUSTER-IP    EXTERNAL-IP    PORT(S)    AGE
svc / elasticsearch-apiserver ClusterIP      172.17.0.134    <
none > 9200 / TCP    22h
svc / elasticsearch-discovery ClusterIP      172.17.0.91     <
none > 9300 / TCP    22h
...
NAME      DESIRED    CURRENT    AGE
statefulset / elasticsearch-apiserver 3          3          22h
```

Step 2 Deploy log-pilot and the Kibana service

1. Deploy the log-pilot log collection tool. The orchestration template is as follows:

```
kubectl apply -f https://acs-logging.oss-cn-hangzhou.aliyuncs.com/log-pilot.yml
```

2. Deploy the Kibana service. The sample orchestration template contains a service and a deployment.

```
kubectl apply -f https://acs-logging.oss-cn-hangzhou.aliyuncs.com/kibana.yml
```

Step 3 Deploy the test application Tomcat

After deploying the log tool set of Elasticsearch + log-pilot + Kibana, deploy a test application Tomcat to test whether or not logs can be successfully collected, indexed, and displayed.

The orchestration template is as follows:

```
apiVersion: v1
kind: Pod
metadata:
  name: tomcat
  namespace: default
  labels:
    name: tomcat
spec:
  containers:
    - image: tomcat
      name: tomcat-test
      volumeMounts:
        - mountPath: /usr/local/tomcat/logs
          name: accesslogs
      env:
```

```

- name : aliyun_log s_catalina
  value : " stdout " ## Collect standard output logs .
- name : aliyun_log s_access
  value : "/ usr / local / tomcat / logs / catalina . *. log "
## Collect log files within the container .
volumes :
- name : accesslogs
  emptyDir : {}

```

The Tomcat image is a Docker image that both uses stdout and file logs. In the preceding orchestration, the log collection configuration file is dynamically generated by defining the environment variable in the pod. See the following descriptions for the environment variable:

- `aliyun_log s_catalina = stdout` indicates to collect stdout logs from the container.
- `aliyun_log s_access = / usr / local / tomcat / logs / catalina . *. log` indicates to collect all the log files whose name matches `catalina . *. log` under the directory `/ usr / local / tomcat / logs /` from the container.

In the Elasticsearch scenario of this solution, the `$ name` in the environment variable indicates index. In this example, `$ name` is `catalina` and `access`.

Step 4 Expose the Kibana service to Internet

The Kibana service deployed in the preceding section is of the NodePort type, which cannot be accessed from the Internet by default. Therefore, create an Ingress in this document to access the Kibana service from Internet and test whether or not logs are successfully indexed and displayed.

1. Create an Ingress to access the Kibana service from Internet. In this example, use the simple routing service to create an Ingress. For more information, see [#unique_171](#). The orchestration template of the Ingress is as follows:

```

apiVersion : extensions / v1beta1
kind : Ingress
metadata :
  name : kibana - ingress
  namespace : kube - system # Make sure the namespace is
the same as that of the Kibana service .
spec :
  rules :
  - http :
    paths :
    - path : /
      backend :

```

```

serviceName : kibana # Enter the name of the
Kibana service .
servicePort : 80 # Enter the port exposed by
the Kibana service .

```

2. After the Ingress is successfully created, run the following commands to obtain the access address of the Ingress:

```

$ kubectl get ingress -n=kube-system
NAME HOSTS ADDRESS PORTS AGE
shared-dns * 120.55.150.30 80 5m

```

3. Access the address in the browser as follows.

4. Click Management in the left-side navigation pane. Then, click Index Patterns > Create Index Pattern. The detailed index name is the \$name variable suffixed with a time string. You can create an index pattern by using the wildcard *. In this example, use \$name* to create an index pattern.

You can also run the following commands to enter the corresponding pod of Elasticsearch and list all the indexes of Elasticsearch:

```

$ kubectl get pods -n=kube-system # Find the
corresponding pod of Elasticsearch .
...
$ kubectl exec -it elasticsearch-1 bash # Enter a
pod of Elasticsearch .
...
$ curl 'localhost:9200/_cat/indices?v' ## List all
the indexes .
health status index uuid pri rep docs . count docs .
deleted store . size pri . store . size
green open .kibana x06jj19PS4 CIm6Ajo51P Wg 1 1 4
0 53 . 6kb 26 . 8kb
green open access-2018.03.19 txd3tG-NR6-
guqmMEKKzE w 5 1 143 0 823 . 5kb 411 . 7kb
green open catalina-2018.03.19 ZgtWd16FQ7 qqJNNWXxFP
cQ 5 1 143 0 915 . 5kb 457 . 5kb

```

5. After successfully creating the indexes, click Discover in the left-side navigation pane, select the created index and the corresponding time range, and then enter the related field in the search box to query logs.

Then, you have successfully tested the solution to log collection problems of Alibaba Cloud Kubernetes clusters based on log-pilot, Elasticsearch, and Kibana. By using this solution, you can deal with requirements for logs of distributed Kubernetes clusters effectively, improve the Operation and Maintenance and operational efficiencies, and guarantee the continuous and stable running of the system.

1.14.5 Configure Log4jAppender for Kubernetes and Log Service

Log4j is an open-source project of Apache, which consists of three important components: log level, log output destination, and log output format. By configuring Log4jAppender, you can set the log output destination to console, file, GUI component, socket server, NT event recorder, or UNIX Syslog daemon.

This document introduces how to configure a YAML file to output Alibaba Cloud Container Service Kubernetes cluster logs to Alibaba Cloud Log Service, without modifying the application codes. In this document, deploy a sample API application in the Kubernetes cluster for demonstration.

Prerequisites

- You have activated Container Service and created a Kubernetes cluster.

In this example, create a Kubernetes cluster in the region of China East 1 (Hangzhou).

- Enable AccessKey or Resource Access Management (RAM). Make sure you have sufficient access permissions. Use the AccessKey in this example.

Step 1 Configure Log4jAppender in Alibaba Cloud Log Service

1. Log on to the [Log Service console](#).
2. On the Project List page, click Create Project in the upper-right corner. Complete the configurations and then click Confirm to create the project.

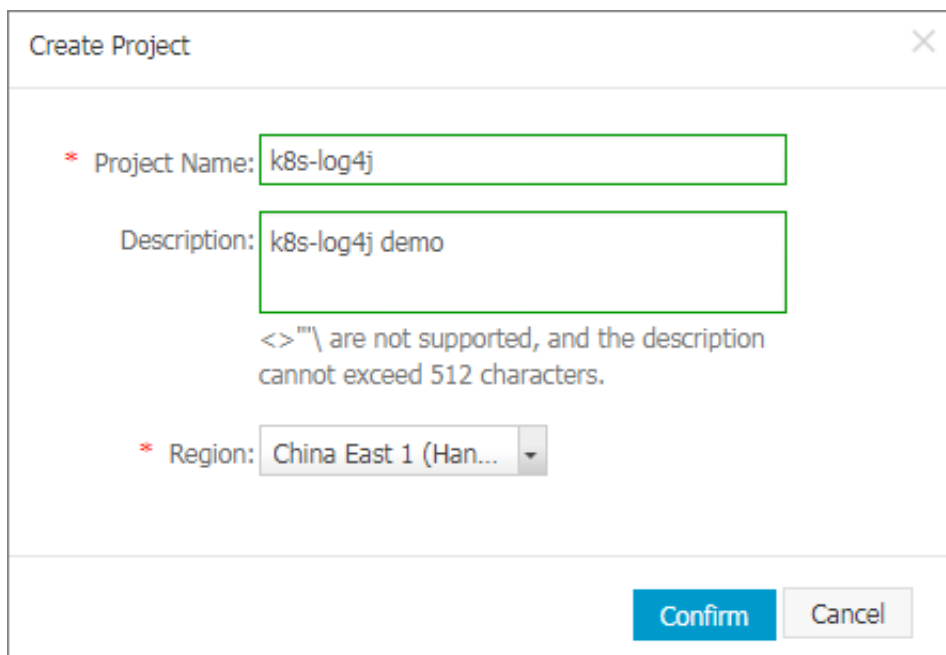
In this example, create a project named k8s-log4j and select the same region (China East 1 (Hangzhou)) as the Kubernetes cluster.



Note:

Generally, create a Log Service project in the same region as the Kubernetes cluster. When the Kubernetes cluster and Log Service project are in the same region, log data is transmitted by using the intranet, which saves the Internet

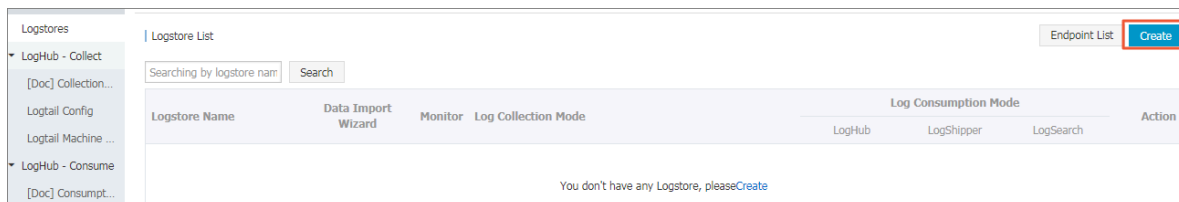
bandwidth cost and time of data transmission because of different regions, and implements the best practice of real-time collection and quick query.



The 'Create Project' dialog box contains the following fields and controls:

- Project Name:** A text input field containing 'k8s-log4j'.
- Description:** A text input field containing 'k8s-log4j demo'. Below this field is a warning message: '<>""\ are not supported, and the description cannot exceed 512 characters.'
- Region:** A dropdown menu showing 'China East 1 (Han...'.
- Buttons:** 'Confirm' (blue) and 'Cancel' (grey) buttons at the bottom right.

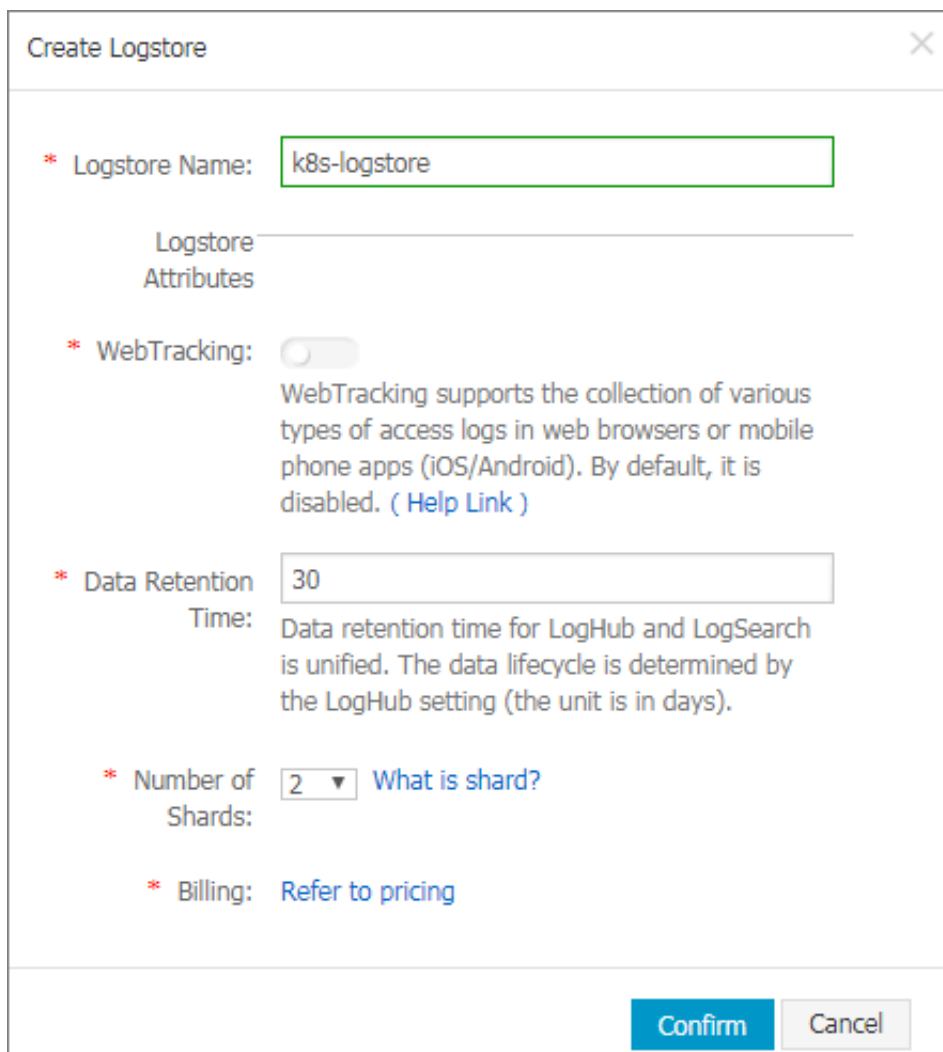
3. After being created, the project k8s-log4j is displayed on the Project List page. Click the project name.
4. The Logstore List page appears. Click Create in the upper-right corner.



The 'Logstore List' page features a sidebar on the left with navigation links: Logstores, LogHub - Collect, [Doc] Collection..., Logtail Config, Logtail Machine ..., LogHub - Consume, and [Doc] Consumpt... The main content area has a search bar 'Searching by logstore name' and a 'Search' button. Below is a table with columns: Logstore Name, Data Import Wizard, Monitor, Log Collection Mode, Log Consumption Mode (with sub-columns LogHub, LogShipper, LogSearch), and Action. The table is currently empty, displaying the message 'You don't have any Logstore, please Create'. In the top right corner, there are links for 'Endpoint List' and a 'Create' button (highlighted with a red box).

5. Complete the configurations and then click Confirm.

In this example, create a Logstore named k8s-logstore.

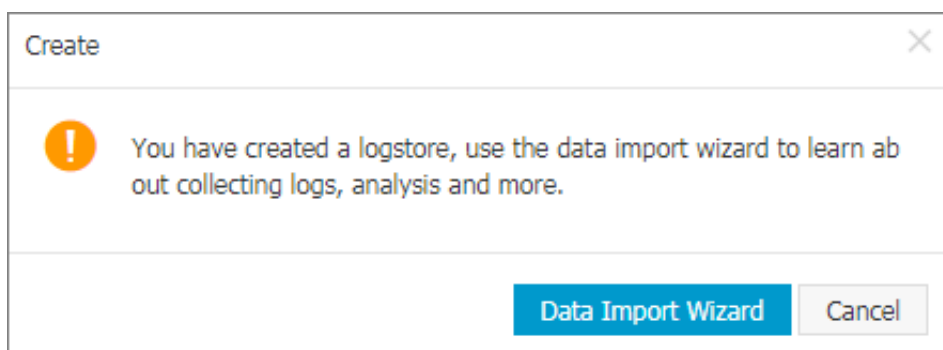


The 'Create Logstore' dialog box contains the following fields and options:

- Logstore Name:** A text input field containing 'k8s-logstore'.
- Logstore Attributes:** A section header.
- WebTracking:** A toggle switch that is currently disabled. Below it, text states: 'WebTracking supports the collection of various types of access logs in web browsers or mobile phone apps (iOS/Android). By default, it is disabled. ([Help Link](#))'.
- Data Retention Time:** A text input field containing '30'. Below it, text states: 'Data retention time for LogHub and LogSearch is unified. The data lifecycle is determined by the LogHub setting (the unit is in days).'.
- Number of Shards:** A dropdown menu showing '2' and a link 'What is shard?'.
- Billing:** A link 'Refer to pricing'.

At the bottom right, there are two buttons: 'Confirm' (highlighted in blue) and 'Cancel'.

6. Then, a dialog box asking you to use the data import wizard appears.

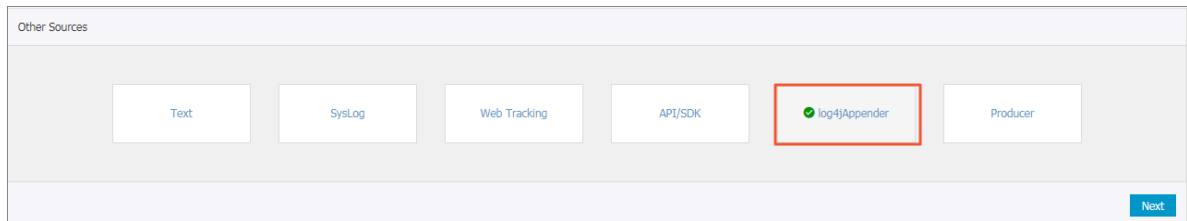


The 'Create' dialog box displays a warning icon (an orange circle with an exclamation mark) and the following text: 'You have created a logstore, use the data import wizard to learn about collecting logs, analysis and more.'

At the bottom right, there are two buttons: 'Data Import Wizard' (highlighted in blue) and 'Cancel'.

7. Click Data Import Wizard. In the Select Data Source step, select log4jAppender under Other Sources and then complete the configurations as instructed on the page.

Use the default configurations in this example. Configure the settings according to the specific scenarios of log data.



Step 2 Configure Log4jAppender in the Kubernetes cluster

In this example, use the sample YAML files [demo-deployment](#) and [demo-service](#) for demonstration.

1. Connect to your Kubernetes cluster.

For more information, see [Access Kubernetes clusters by using SSH](#) or [Connect to a Kubernetes cluster by using kubectl](#).

2. Obtain the `demo - deployment . yaml` file and configure the environment variable `JAVA_OPTS` to collect logs from the Kubernetes cluster.

The sample orchestration of the `demo - deployment . yaml` file is as follows:

```
apiVersion : apps / v1beta2
kind : Deployment
metadata :
  name : log4j - appender - demo - spring - boot
  labels :
    app : log4j - appender
spec :
  replicas : 1
  selector :
    matchLabels :
      app : log4j - appender
  template :
    metadata :
      labels :
        app : log4j - appender
    spec :
      containers :
        - name : log4j - appender - demo - spring - boot
          image : registry . cn - hangzhou . aliyuncs . com /
jaegertracing / log4j - appender - demo - spring - boot : 0 . 0 .
2
          env :
            - name : JAVA_OPTS ## Note
```

```

        value : "- Dproject={ your_project } - Dlogstore={
your_logstore } - Dendpoint={ your_endpoint } - Daccess_key_id={ your_access_key_id } - Daccess_key_secret={ your_access_key_secret }"
        ports :
        - containerPort : 8080

```

Wherein:

- `- Dproject` : The name of the used Alibaba Cloud Log Service project. In this example, it is `k8s-log4j`.
- `- Dlogstore` : The name of the used Alibaba Cloud Log Service Logstore. In this example, it is `k8s-logstore`.
- `- Dendpoint` : The service endpoint of Log Service. You must configure your service endpoint according to the region where the Log Service project resides. For more information, see [Service endpoint](#). In this example, it is `cn-hangzhou.log.aliyuncs.com`.
- `- Daccess_key_id` : Your AccessKey ID.
- `- Daccess_key_secret` : Your AccessKey Secret.

3. Run the following command in the command line to create the deployment:

```
kubectl create -f demo - deployment . yaml
```

4. Obtain the `demo - service . yaml` file and run the following command to create the service.

No need to modify the configurations in the `demo - service . yaml` file.

```
kubectl create -f demo - service . yaml
```

Step 3 Test to generate Kubernetes cluster logs

You can run the `kubectl get` command to view the deployment status of the resource object. Wait until the deployment and the service are successfully deployed. Then, run the `kubectl get svc` command to view the external access IP of the service, that is, the EXTERNAL-IP.

```

$ kubectl get svc
NAME      TYPE      CLUSTER - IP      EXTERNAL - IP      PORT ( S )      AGE
log4j - appender - demo - spring - boot - svc      LoadBalancer      172
. 21 . XX . XX      120 . 55 . XXX . XXX      8080 : 30398 / TCP      1h

```

In this example, test to generate Kubernetes cluster logs by running the `login` command, wherein, `K8S_SERVICE_IP` is the `EXTERNAL - IP`.



Note:

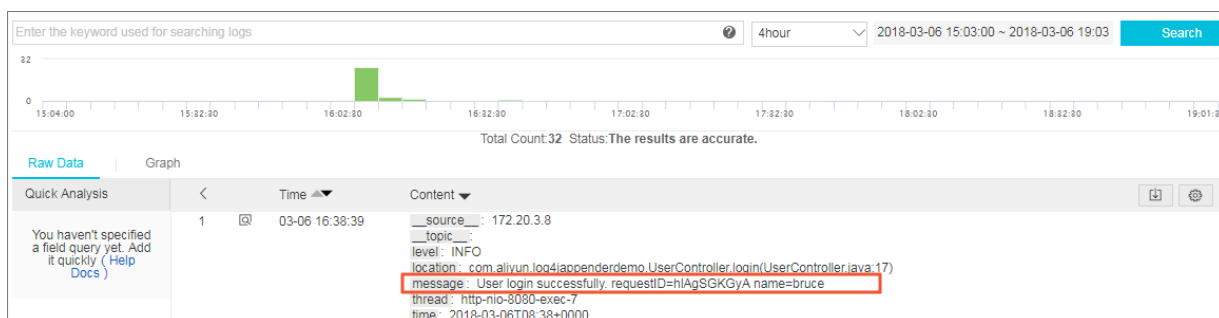
See [GitHub log4j-appender-demo](#) to view the complete collection of APIs.

```
curl http ://${ K8S_SERVIC E_IP }: 8080 / login ? name = bruce
```

Step 4 View logs in Alibaba Cloud Log Service

Log on to the [Log Service console](#).

Click the project name and click Search at the right of the Logstore k8s-logstore to view the output logs of the Kubernetes cluster.



The output content of the log corresponds to the preceding command. This example demonstrates how to output the logs of the sample application to Alibaba Cloud Log Service. By completing the preceding steps, you can configure Log4JAppender in Alibaba Cloud and implement advanced functions such as collecting logs in real time , filtering data, and querying logs by using Alibaba Cloud Log Service.

1.15 Monitoring management

1.15.1 Deploy the Prometheus monitoring system

Prometheus is an open source monitoring tool for cloud native applications. This topic describes how to deploy the Prometheus monitoring system by using Alibaba Cloud Container Service for Kubernetes.

Background information

A monitoring system monitors the following two types of objects:

- Resource, namely, the resource usage of a node or application. The monitoring system of Container Service for Kubernetes monitors node resource usage, cluster resource usage, and pod resource usage.

- Application, namely, internal metrics of an application. For example, The monitoring system collects statistics regarding the number of online users that use an application in real time, and performs service-level monitoring and alarming for the application by exposing ports.

The following are the objects monitored in a Kubernetes cluster:

- System components, which are built-in components of the Kubernetes cluster, such as apiserver, controller-manager, and etcd.
- Static resource entities, which include node resource status and kernel events.
- Dynamic resource entities, which are abstract workload entities of Kubernetes, such as deployment, DaemonSet, and pods.
- Customized application objects, which includes the data and metrics that require customization within an application.

To monitor system components and static resource entities, you need to specify monitoring methods for them in the configuration file.

To monitor dynamic resource entities, we recommend that you deploy the Prometheus monitoring system.

Prerequisites

- You have created a Kubernetes cluster. For more information, see [Create a Kubernetes cluster](#).
- You have connected to the Master node so that you can view node labels and other information. For more information, see [Connect to a Kubernetes cluster by using kubectl](#).

Deploy the Prometheus monitoring system

1. Run the following command to download the prometheus-operator code:

```
git clone https://github.com/AliyunContainerService/prometheus-operator
```

2. Run the following command to deploy the Prometheus monitoring system:



Note:

Some Prometheus components may fail to be deployed when you run this command for the first time because Prometheus components require a specific

sequence to be deployed. If any exceptions occur during your first deployment, you need to run the command again.

```
cd prometheus-operator/contrib/kube-prometheus
kubectl apply -f manifests
```

3. Run the following command to set the access method for Prometheus:

```
kubectl --namespace monitoring port-forward svc/prometheus-k8s 9090
```

4. View the deployment result

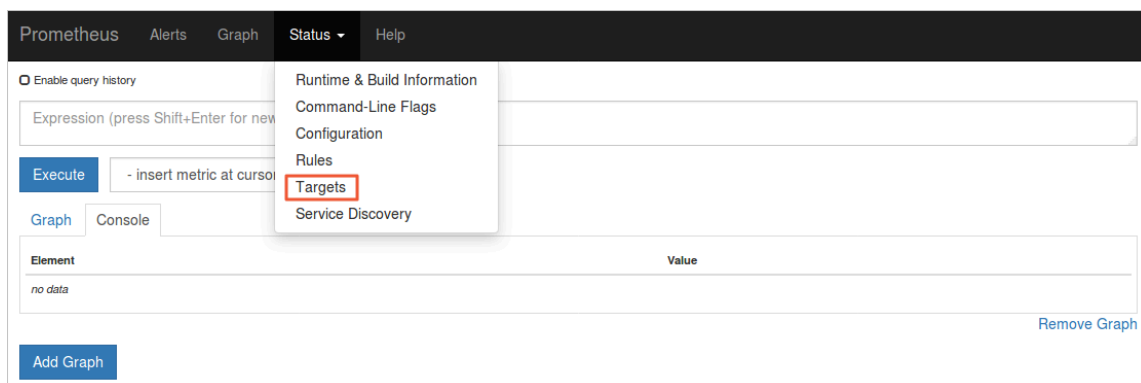
- a. To view Prometheus, access `localhost : 9090` in a browser.



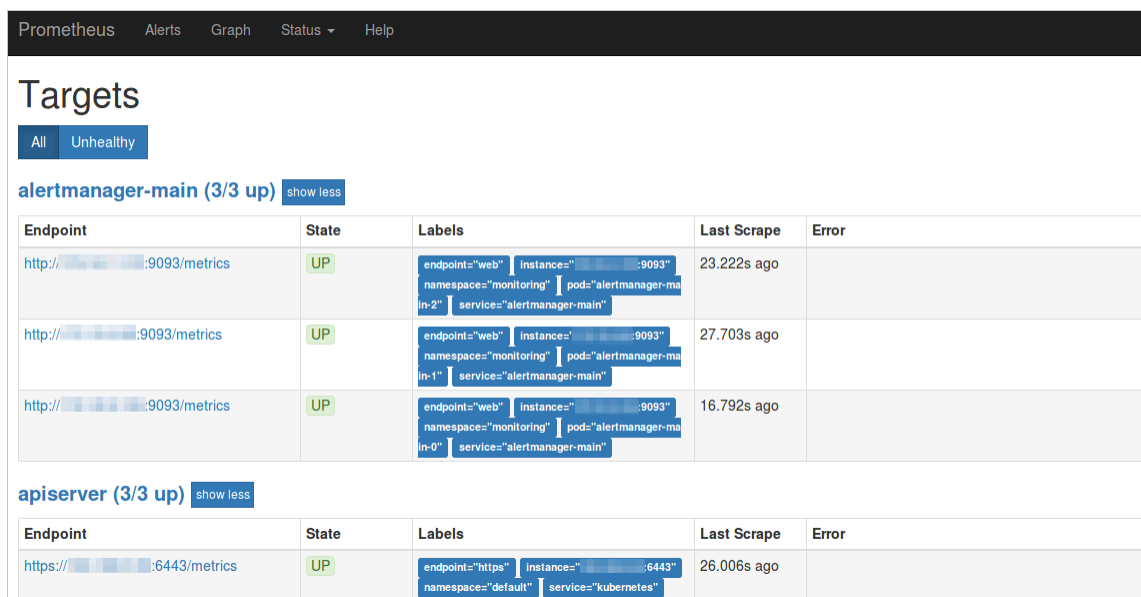
Note:

By default, Prometheus cannot be accessed through the Internet. You must use your local proxy to access it.

- b. Select Targets under the Status menu to view all collection tasks.



If the status of all tasks is UP, all collection tasks are running properly.



View and display data aggregation

1. Run the following command to access Grafana:

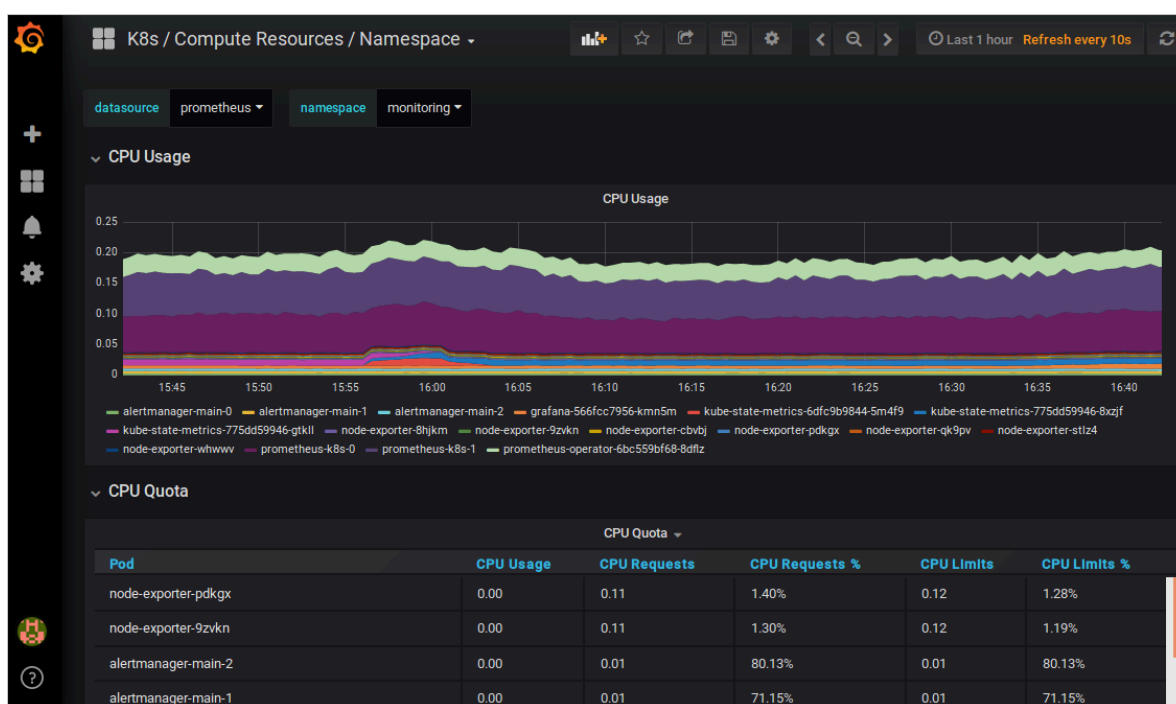
```
kubectl -- namespace monitoring port - forward svc / grafana 3000
```

2. Access `localhost : 3000` in your browser and then select a dashboard to view data aggregation.



Note:

The default user name and password are both admin.

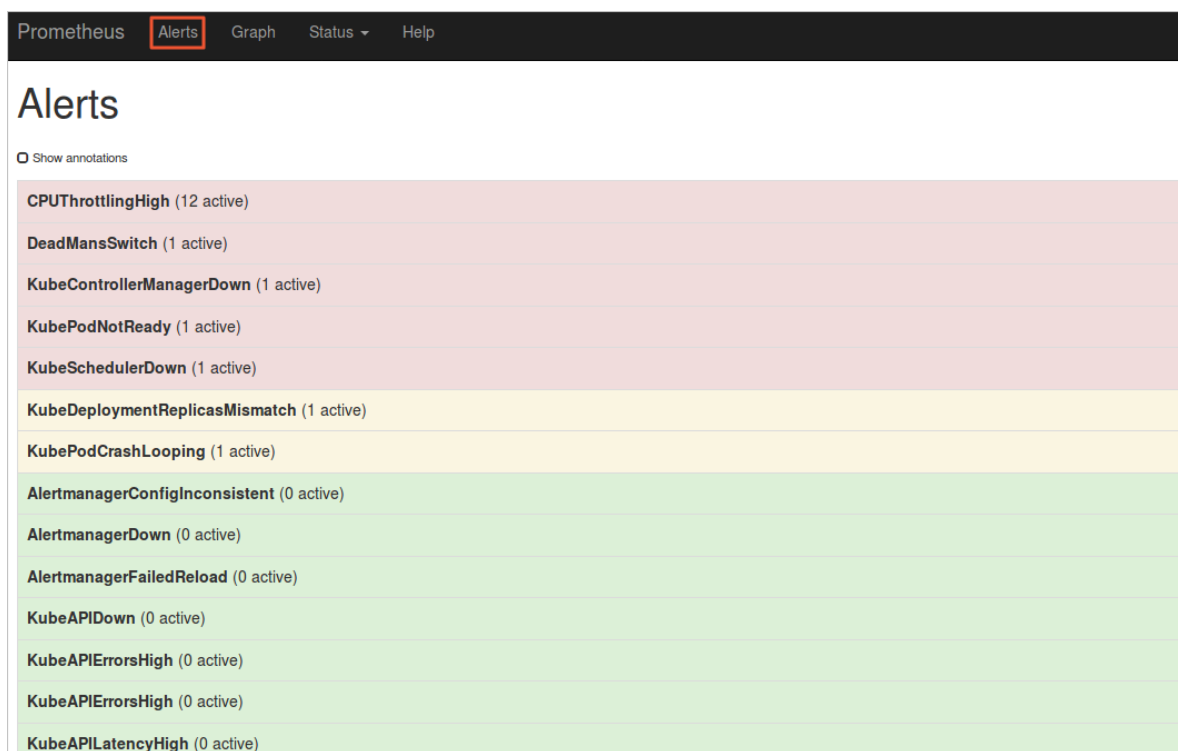


View alerting rules and set alert silencing

- View alerting rules

Access `localhost : 9090` in your browser and click the Alerts menu to view the current alerting rules.

- Red: indicates that an alert is triggered.
- Green: indicates the normal status.



The screenshot shows the Prometheus web interface with the 'Alerts' tab selected. The page title is 'Alerts'. Below the title is a link 'Show annotations'. The main content is a table of alerting rules, each with a name and an active count. The rules are color-coded: red for triggered alerts and green for normal status.

Alert Name	Active Count	Status
CPUThrottlingHigh	12 active	Red
DeadMansSwitch	1 active	Red
KubeControllerManagerDown	1 active	Red
KubePodNotReady	1 active	Red
KubeSchedulerDown	1 active	Red
KubeDeploymentReplicasMismatch	1 active	Yellow
KubePodCrashLooping	1 active	Yellow
AlertmanagerConfigInconsistent	0 active	Green
AlertmanagerDown	0 active	Green
AlertmanagerFailedReload	0 active	Green
KubeAPIDown	0 active	Green
KubeAPIErrorsHigh	0 active	Green
KubeAPIErrorsHigh	0 active	Green
KubeAPILatencyHigh	0 active	Green

- Set alert silencing

Run the following command, open `localhost : 9093` in your browser, and select **Silenced** to set alert silencing:

```
kubectl --namespace monitoring port -f svc /
alertmanager - main 9093
```

The screenshot shows the Alertmanager web interface. At the top, there are tabs for 'Filter' and 'Group', and a 'Receiver: All' dropdown with radio buttons for 'Silenced' and 'Inhibited'. Below this is a search bar with a placeholder 'Custom matcher, e.g. env="production"' and a blue '+' button. Underneath, there's a section for 'alertname="CPUThrottlingHigh"' with a '+' button. Below that, the date '08:14:17, 2018-10-29' is shown, followed by links for '+ Info', 'Source', and 'Silence' (the 'Silence' link is highlighted with a red box). At the bottom, there are several filter tags: 'severity="warning"', 'prometheus="monitoring/k8s"', 'pod_name="kube-state-metrics-775dd59946-8xzj"', 'namespace="monitoring"', and 'container_name="addon-resizer"', each with a '+' button.

1.15.2 Group-based monitoring and alarms

Alibaba Cloud Container Service is interoperable with CloudMonitor to enable group-based monitoring and alarms.

Prerequisites

- [Create a Kubernetes cluster](#) if you do not have one.
- The Kubernetes version must be 1.8.4 or later. Otherwise, you must first upgrade the cluster.

Context

In the Operation & Maintenance (O&M) of IT infrastructure, monitoring and alarms facilitate daily O&M, system monitoring, troubleshooting, and debugging, and guarantee the reliability and security of O&M.

The traditional container monitoring solution that uses a statically configured monitoring agent or a centralized server for monitoring and alarms may not be suitable for the Kubernetes scenario because it can cause some problems. For example, the information required to identify the monitoring objects is missing because containers are mostly scheduled in the resource pool whereas the monitoring agent is deployed on the host. Also, containers have shorter lives than applications. The monitoring and alarm rules, and such monitoring data as

ReplicaSet and Deployment for a single container cannot be used for the corresponding application.

Alibaba Cloud Container Service for Kubernetes is deeply integrated with CloudMonitor or to use application groups to unify the monitoring objects and metrics. In addition, CloudMonitor of Alibaba Cloud is equipped with many functions and custom tools, which provide you with the best practice to monitor your Kubernetes resources and manage the alarms.

Procedure

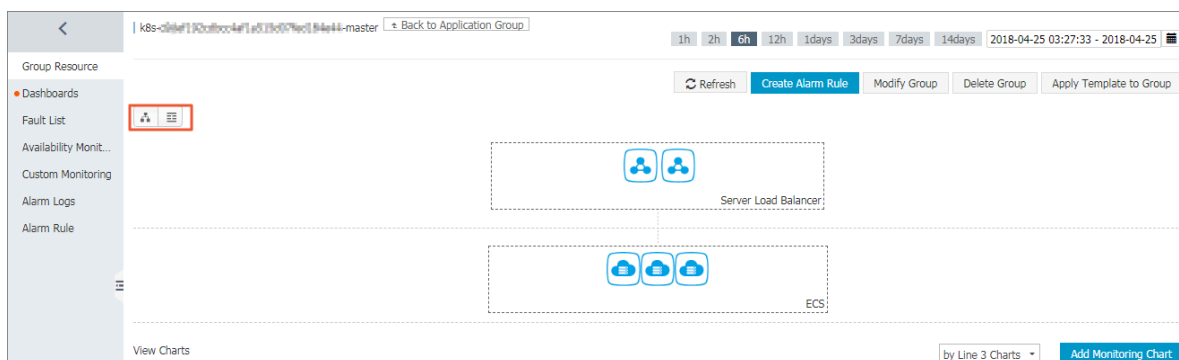
1. Log on to the [CloudMonitor console](#).
2. In the left-side navigation pane, click Application Groups. The Kubernetes groups with cluster IDs are displayed.

Group Name / Group ID	Health Status	Type	Total Server Number	Resource Types	Unhealthy Instances	Created At	Actions
g20 / 23939	✓	Custom	2	1	0	2017-09-08 13:45:22	Manage More
k8s-elasticsearch-64553-master	✓	kubernetes	3	2	0	2018-04-24 10:16:42	Manage More
k8s-elasticsearch-64554-worker	✓	kubernetes	1	1	0	2018-04-24 10:16:42	Manage More

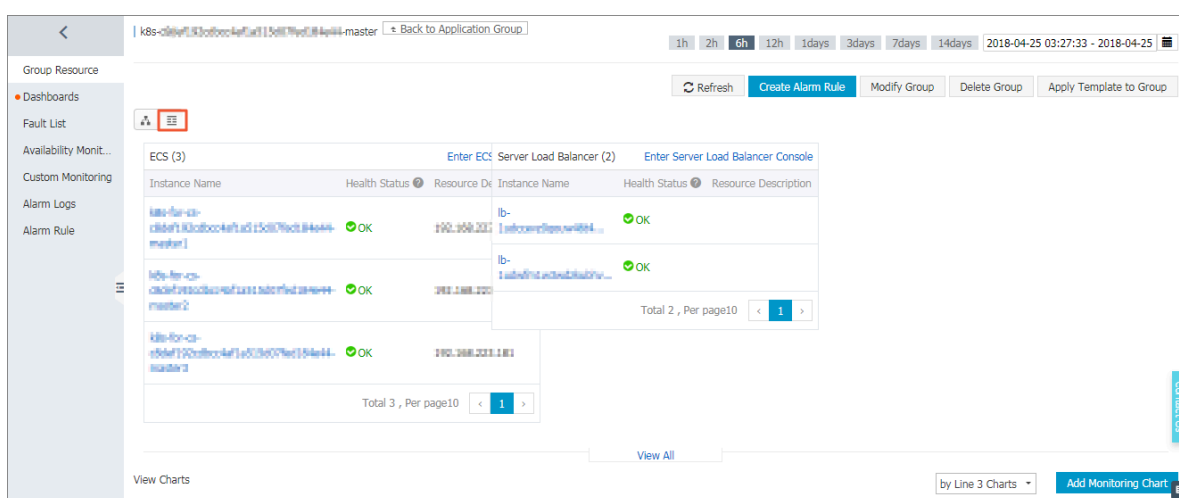
3. Click the Group Name to go to the group details page. You can view the resources contained in the group. For example, in a Master group of Kubernetes, you can see such resources as Elastic Compute Service (ECS) instances and Server Load Balancer (SLB) instances.

Kubernetes has two types of nodes: Worker nodes and Master nodes. Master nodes generally contain management and control applications and the resources are required to be highly robust. Worker nodes are generally responsible for scheduling pods and the overall requirement on the resources focuses on scheduling capability. When you create a group, Container Service automatically creates two resource groups, a Master group and a Worker group. The Master

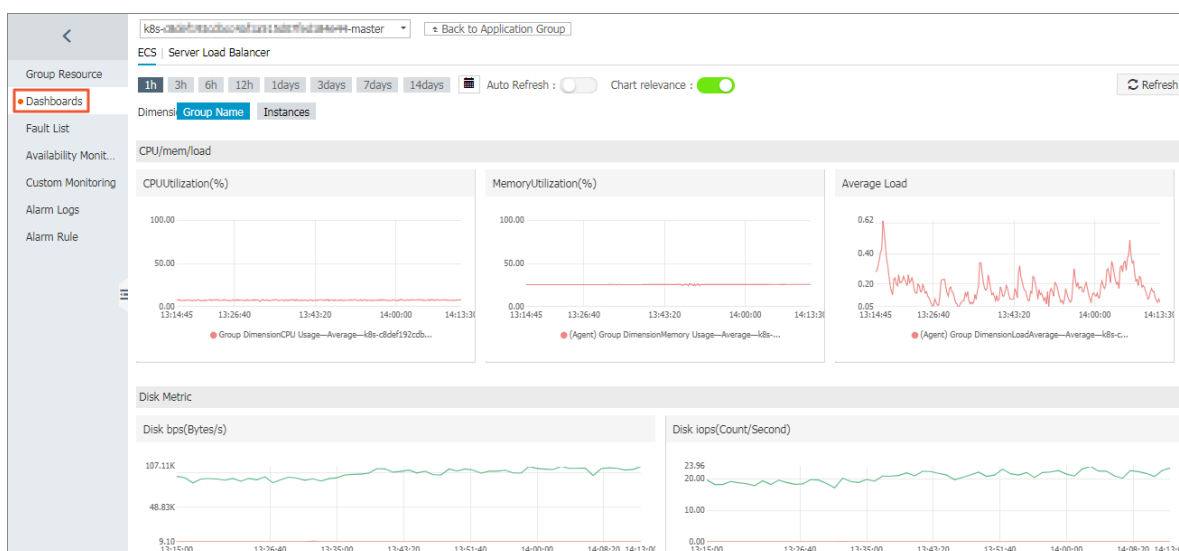
group includes the Master nodes and the related SLB instances. The Worker group includes all the Worker nodes.



4. You can view the details of other cloud products, such as SLB, in the group.

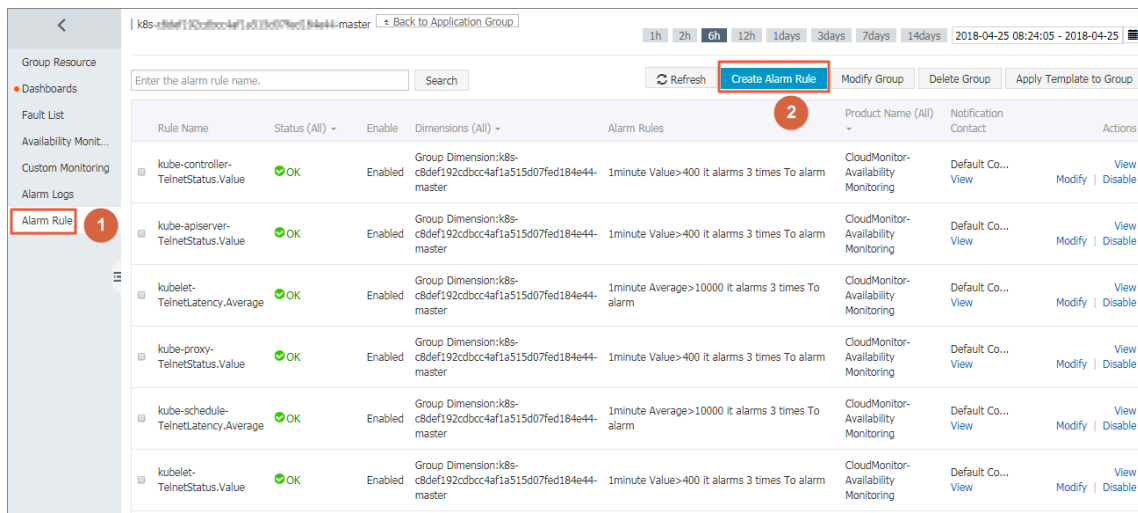


5. In the left-side navigation pane, click Dashboards to view the detailed monitoring metrics of each cloud product in the group.



6. In the left-side navigation pane, click Alarm Rule. A list of existing alarm rules in the group is displayed. By default, the health of the core components of all nodes in the Mater group is checked.

a. Click Create Alarm Rule to create an alarm rule for the group according to your business requirements.



The screenshot displays the 'Alarm Rule' management page. On the left, the navigation pane has 'Alarm Rule' selected, indicated by a red box and the number 1. The main content area shows a list of alarm rules for the 'k8s-c8def192cdccc4af1a515d07fed184e44-master' group. A red box and the number 2 highlight the 'Create Alarm Rule' button. The table below lists the existing alarm rules:

Rule Name	Status (All)	Enable	Dimensions (All)	Alarm Rules	Product Name (All)	Notification Contact	Actions
kube-controller- TelnStatus.Value	OK	Enabled	Group Dimension:k8s- c8def192cdccc4af1a515d07fed184e44- master	1minute Value>400 it alarms 3 times To alarm	CloudMonitor- Availability Monitoring	Default Co... View	Modify View Disable
kube-apiserver- TelnStatus.Value	OK	Enabled	Group Dimension:k8s- c8def192cdccc4af1a515d07fed184e44- master	1minute Value>400 it alarms 3 times To alarm	CloudMonitor- Availability Monitoring	Default Co... View	Modify View Disable
kubelet- TelnLatency.Average	OK	Enabled	Group Dimension:k8s- c8def192cdccc4af1a515d07fed184e44- master	1minute Average>10000 it alarms 3 times To alarm	CloudMonitor- Availability Monitoring	Default Co... View	Modify View Disable
kube-proxy- TelnStatus.Value	OK	Enabled	Group Dimension:k8s- c8def192cdccc4af1a515d07fed184e44- master	1minute Value>400 it alarms 3 times To alarm	CloudMonitor- Availability Monitoring	Default Co... View	Modify View Disable
kube-schedule- TelnLatency.Average	OK	Enabled	Group Dimension:k8s- c8def192cdccc4af1a515d07fed184e44- master	1minute Average>10000 it alarms 3 times To alarm	CloudMonitor- Availability Monitoring	Default Co... View	Modify View Disable
kubelet- TelnStatus.Value	OK	Enabled	Group Dimension:k8s- c8def192cdccc4af1a515d07fed184e44- master	1minute Value>400 it alarms 3 times To alarm	CloudMonitor- Availability Monitoring	Default Co... View	Modify View Disable

b. On the displayed Create Alarm Rule page, set the alarm rules.

- Select the related resource, such as ECS.
- Select whether to use a template to create the alarm rule. If yes, select an alarm template from the Select Template drop-down list. You can also click

Create Alarm Template to create a new custom alarm template. For more information, see .

- Set the notification method. For example, you can know the Kubernetes cluster status through DingTalk, email, and SMS.

Create Alarm Rule
Back to

1
Related Resource

Products : k8s

Resource Range : Application Group
Create Application Group Improving the efficiency of maintenance Best Practice

Group Name : k8s-c2174629b4ea049d887da6e71d...

2
Set Alarm Rules

Use Template : Yes No

Select Template : 常用基础模板 x Create Alarm Template

常用基础模板_cpu_total	Host.cpu.total	1m	it alar...	>	>	90	%
常用基础模板_diskusage_utili	Host.disk.utilization	1m	it alar...	>	>	90	%
常用基础模板_memory_usedt	Host.mem.usedutilization l...	1m	it alar...	>	>	90	%
常用基础模板_InternetOutRat	Internet Outbound Bandwi...	1m	it alar...	>	>	90	%
常用基础模板_agent_heartbe	Heartbeat not Detected	1m					

Mute for : 24h

Effective Period : 00:00 To: 23:59

3
Notification Method

Notification Contact : Contact Group All

Search

Selected Groups 1 count All

云账号报警联系人

Quickly create a contact group

c. Click Confirm. The created alarm rule is displayed on the Alarm Rule page.

Back
k8s-c2174629b4ea049d887da6e71d...-master
Back to Application Group

1h 2h 6h 12h 1days 3days 7days 14days 2018-04-25 19:26:17 - 2018-04-2

Group Resource

Dashboards

Fault List

Availability Monit...

Custom Monitoring

Alarm Logs

Alarm Rule

Enter the alarm rule name. Search

Refresh Create Alarm Rule Modify Group Delete Group Apply Template to Group

Rule Name	Status (All)	Enable	Dimensions (All)	Alarm Rules	Product Name (All)	Notification Contact	Actions
test_CPUUtilization	OK	Enabled	Group Dimension:k8s-c8def192cdbc4af1a515d07fed184e44-master	1minute CPU Usage Average>=90 % it alarms 1 times To alarm	ECS	Default Co... View	Modify Disable Delete
kube-controller-TelnetStatus.Value	OK	Enabled	Group Dimension:k8s-c8def192cdbc4af1a515d07fed184e44-master	1minute Value>400 it alarms 3 times To alarm	CloudMonitor-Availability Monitoring	Default Co... View	Modify Disable

What's next

More features are provided to meet your resource monitoring requirements, such as fault list, event monitoring, availability monitoring, and log monitoring. You can find them in the left-side navigation pane.

1.15.3 Integration and usage with CloudMonitor

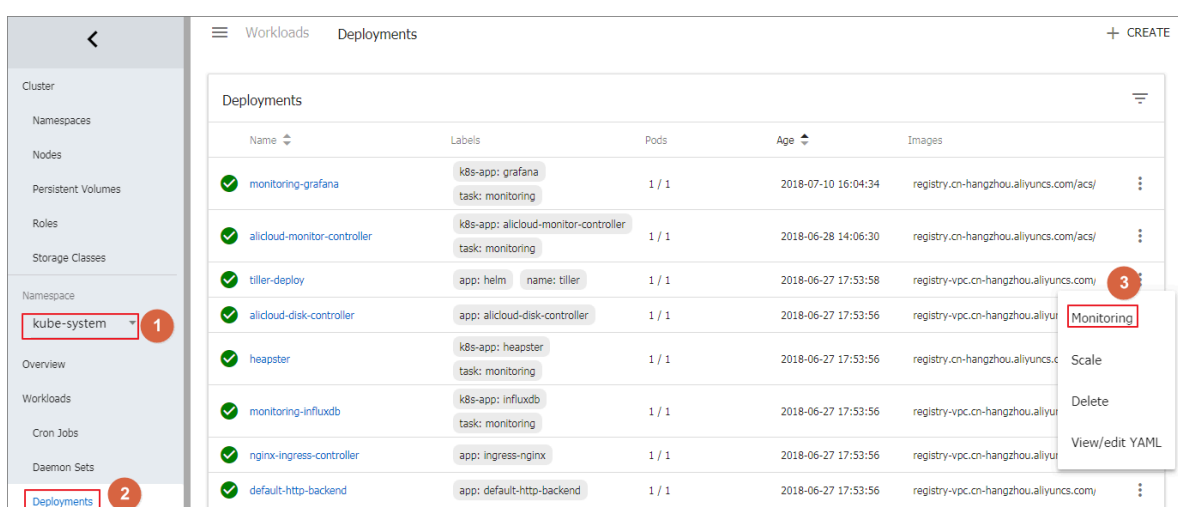
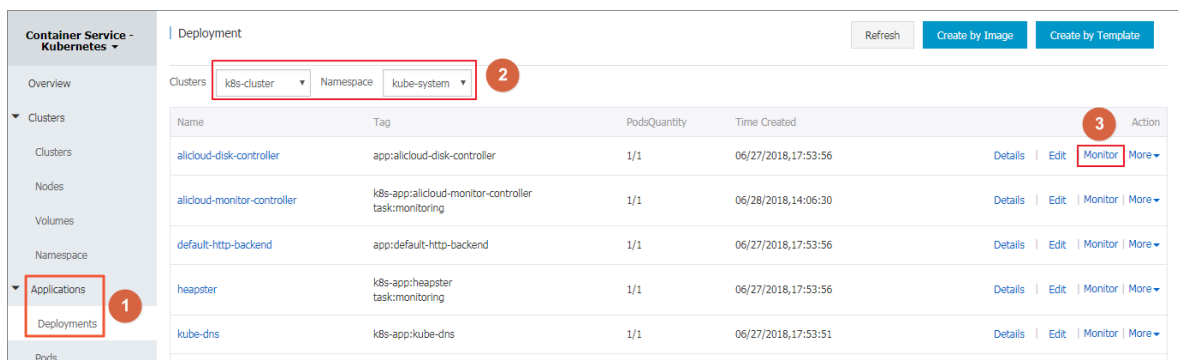
Prerequisites

Check whether `alicloud - monitor - controller` has been deployed in the `kube - system` namespace. If not, upgrade the version of the cluster.

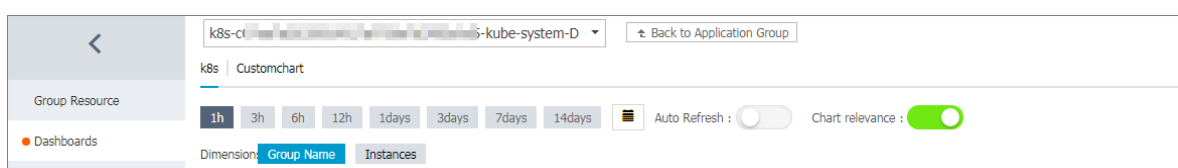
Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Deployment in the left-side navigation pane.

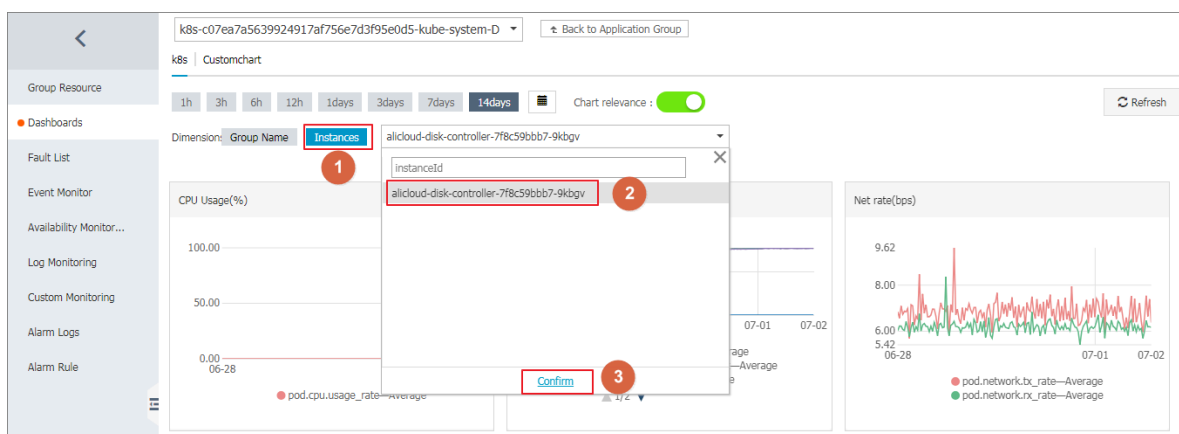
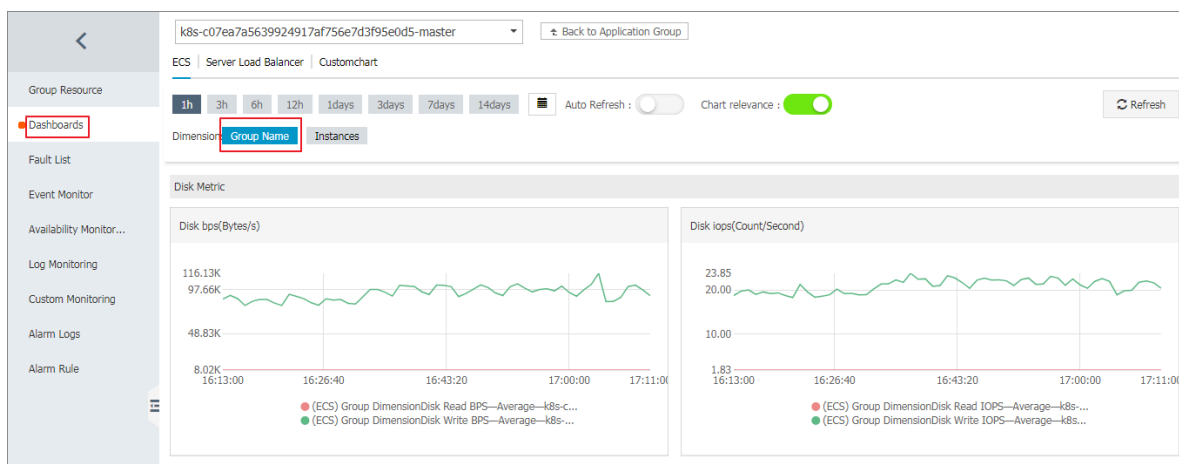
3. Select the target deployment, click Monitor on the right. You can also click Monitor on the Deployment page of the built-in kubernetes dashboard.



In this case, you jump to the corresponding Application group details page of CloudMonitor.



4. Application group supports monitoring in two dimensions: group and instance.



5. For alarm settings, the index of group level starts with group , and the instance level index starts with pod .

1
Related Resource

Products : k8s
Resource Range : Application Group
Group Name : k8s-c07ea7a5639924917af756e7d3f95e0d5-master

2
Set Alarm Rules

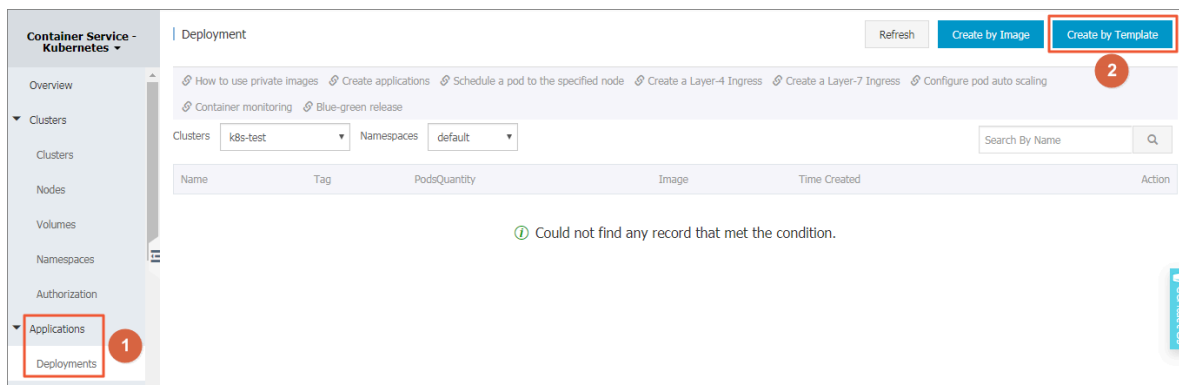
Use Template : Yes No
Alarm Rule : usage of Deployment
Rule Describe : group.cpu.usage_rate 5mins Total >= 80
Mute for : group.cpu.limit group.cpu.request group.cpu.usage
Triggered when threshold is exceeded for : group.cpu.usage_rate group.diskio.read_bytes group.diskio.read_bytes_rate
Effective Period : 00:00 To: 23:59

No Data

The line chart indicates the average aggregate value trend of instances under the application group, making your reference to set the alarm rule threshold.

Upgrade cluster version

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Application > Deployment in the left-side navigation pane to enter the Deployment List page. Click Create by template in the upper-right corner.



3. Select the target cluster, kube-system namespace, and use the following sample template. Then click Create.



Note:

Replace `REGION` and `CLUSTER_ID` with your actual cluster information, and redeploy heapster yaml template.

Clusters

k8s-cluster

Namespace

kube-system

Resource Type

Resource - basic Deployment

Template

```

1  apiVersion: extensions/v1beta1
2  kind: Deployment
3  metadata:
4    name: heapster
5    namespace: kube-system
6  spec:
7    replicas: 1
8    template:
9      metadata:
10     labels:
11       task: monitoring
12       k8s-app: heapster
13     annotations:
14       scheduler.alpha.kubernetes.io/critical-pod: ''
15     spec:
16       serviceAccount: admin
17       containers:
18       - name: heapster
19         image: registry.##REGION##.aliyuncs.com/acs/heapster-amd64:v1.5.1.1
20         imagePullPolicy: IfNotPresent
21         command:
22         - /heapster
23         - --source=kubernetes:https://kubernetes.default
24         - --historical-source=influxdb:http://monitoring-influxdb:8086
25         - --sink=influxdb:http://monitoring-influxdb:8086
26         - --sink=socket:tcp://monitor.csk.##REGION##.aliyuncs.com:8093?clusterId=##CLUSTER_ID
27         D##&public=true

```

DEPLOY

An example of heapster template is as follows. If you have an earlier version of the heapster in the cluster, you can log on to the Kubernetes cluster and run the

`kubectl apply -f xxx . yaml` command to upgrade it.

```

apiVersion : extensions / v1beta1
kind : Deployment
metadata :
  name : heapster
  namespace : kube - system
spec :
  replicas : 1
  template :
    metadata :
      labels :
        task : monitoring
        k8s - app : heapster
      annotation s :
        scheduler . alpha . kubernetes . io / critical - pod : ''
    spec :

```

```

    serviceAccount: admin
  containers:
  - name: heapster
    image: registry.## REGION ##.aliyuncs.com/acs/heapster-amd64:v1.5.1.1
    imagePullPolicy: IfNotPresent
    command:
    - /heapster
    - --source=kubernetes:https://kubernetes.default
    - --historical-source=influxdb:http://monitoring-influxdb:8086
    - --sink=influxdb:http://monitoring-influxdb:8086
    - --sink=socket:tcp://monitor.csk.## REGION ##.aliyuncs.com:8093?clusterId=## CLUSTER_ID ##&public=true

```

The example layout of alicloud-monitor-controller is as follows. Run the `kubectl create -f xxx.yaml` command to deploy alicloud-monitor-controller.

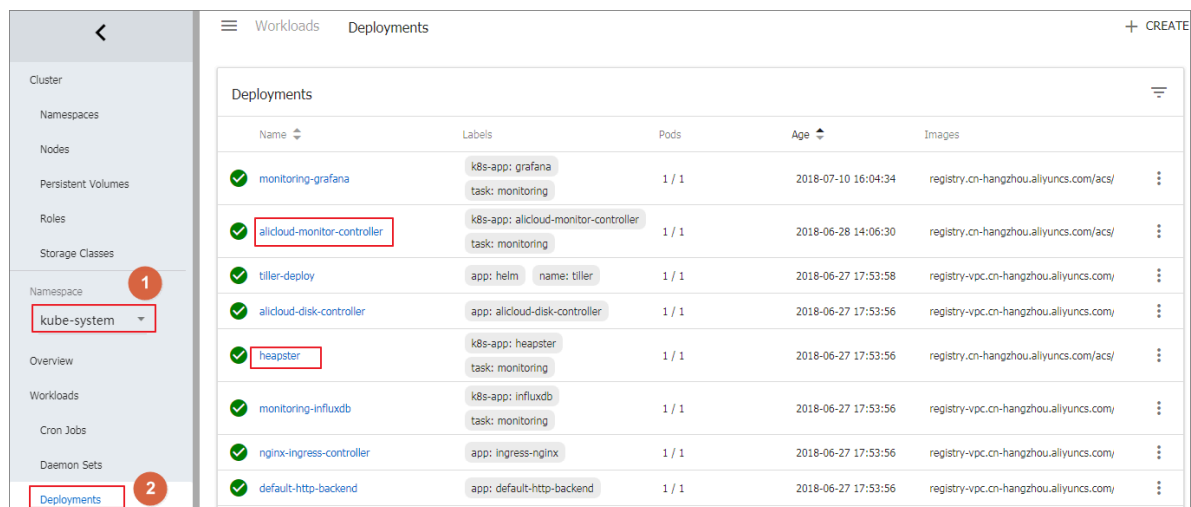
```

apiVersion: extensions/v1beta1
kind: Deployment
metadata:
  name: alicloud-monitor-controller
  namespace: kube-system
spec:
  replicas: 1
  template:
    metadata:
      labels:
        task: monitoring
        k8s-app: alicloud-monitor-controller
      annotations:
        scheduler.alpha.kubernetes.io/critical-pod: ''
    spec:
      hostNetwork: true
      tolerations:
      - effect: NoSchedule
        operator: Exists
        key: node-role.kubernetes.io/master
      - effect: NoSchedule
        operator: Exists
        key: node.cloudprovider.kubernetes.io/uninitialized
      serviceAccount: admin
      containers:
      - name: alicloud-monitor-controller
        image: registry.## REGION ##.aliyuncs.com/acs/alibabacloud-monitor-controller:v1.0.0
        imagePullPolicy: IfNotPresent
        command:
        - /alicloud-monitor-controller
        - agent
        - --regionId=## REGION ##
        - --clusterId=## CLUSTER_ID ##
        - --logstdout

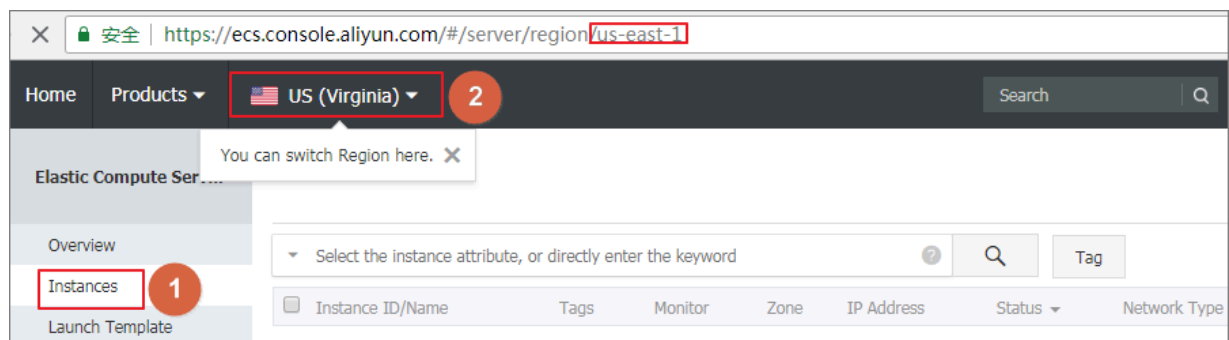
```

- -- v = 4

4. Go to the Kubernetes console. In the kube-system namespace, you can see that the two deployments are running, and the upgrade is complete.



If you do not know the REGION information, you can go to the ECS console and select the region where your cluster resides. The last segment of the page URL address is REGION.



1.15.4 Use Grafana to display monitoring data

Prerequisites

- You have successfully created a Kubernetes cluster. For more information, see [Create a Kubernetes cluster](#).
- In this example, use the Grafana with built-in monitoring templates and the image address is `registry.cn-hangzhou.aliyuncs.com/acs/grafana:5.0.4`.

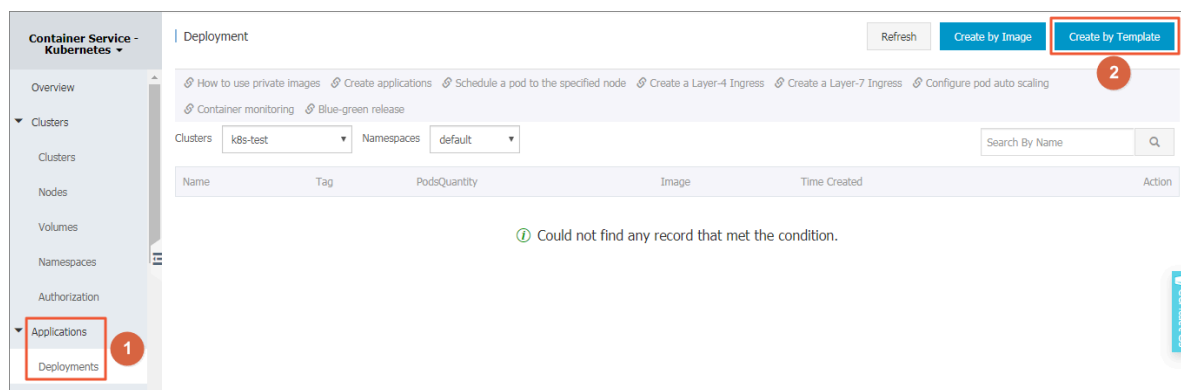
Context

Among Kubernetes monitoring solutions, compared with open-source solutions such as Prometheus, the combination of Heapster + InfluxDB + Grafana is more simple

and direct. Heapster not only collects monitoring data in Kubernetes, but also is relied on by the monitoring interface of the console and the POD auto scaling of HPA . Therefore, Heapster is an essential component of Kubernetes. An Alibaba Cloud Kubernetes cluster has the built-in Heapster + InfluxDB combination. To display the monitoring data, you must configure an available Grafana and the corresponding dashboard.

Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Application > Deployment in the left-side navigation pane.
3. Click Create by template in the upper-right corner.



4. Configure the template to create the deployment and service of Grafana. After completing the configurations, click **DEPLOY**.

- **Clusters:** Select a cluster.
- **Namespace:** Select the namespace to which the resource object belongs, which must be `kube-system`.
- **Resource Type:** Select `Custom` in this example. The template must contain a deployment and a service.

Deploy templates

Only Kubernetes versions 1.8.4 and above are supported. For clusters of version 1.8.1, you can perform "upgrade cluster" operation in the cluster list

Clusters: test
Namespace: kube-system
Resource Type: Custom

1

Template

```

1 apiVersion: extensions/v1beta1
2 kind: Deployment
3 metadata:
4   name: monitoring-grafana
5   namespace: kube-system
6 spec:
7   replicas: 1
8   template:
9     metadata:
10      labels:
11        task: monitoring
12        k8s-app: grafana
13     spec:
14       containers:
15         - name: grafana
16           image: registry.cn-hangzhou.aliyuncs.com/acs/grafana:5.0.4
17           ports:
18             - containerPort: 3000
19               protocol: TCP
20           volumeMounts:
21             - mountPath: /var
22               name: grafana-storage
23           env:
24             - name: INFLUXDB_HOST
25               value: monitoring-influxdb
26       volumes:
27         - name: grafana-storage

```

2

DEPLOY

The orchestration template in this example is as follows:

```

apiVersion : extensions / v1beta1
kind : Deployment
metadata :
  name : monitoring - grafana
  namespace : kube - system
spec :
  replicas : 1
  template :
    metadata :
      labels :
        task : monitoring
        k8s - app : grafana
    spec :
      containers :
        - name : grafana
          image : registry . cn - hangzhou . aliyuncs . com / acs /
grafana : 5 . 0 . 4
          ports :
            - containerP ort : 3000

```

```

      protocol : TCP
      volumeMounts :
      - mountPath : / var
        name : grafana - storage
      env :
      - name : INFLUXDB_HOST
        value : monitoring - influxdb
    volumes :
    - name : grafana - storage
      emptyDir : {}
---
apiVersion : v1
kind : Service
metadata :
  name : monitoring - grafana
  namespace : kube - system
spec :
  ports :
  - port : 80
    targetPort : 3000
  type : LoadBalancer
  selector :
    k8s - app : grafana

```

- Go back to the Deployment page after the successful deployment. Select the cluster from the Clusters drop-down list and then select kube-system from the Namespace drop-down list to view the deployed applications.

The screenshot shows the 'Deployment' page in the Container Service - Kubernetes console. The left sidebar has a menu with 'Applications' selected, highlighted by a red box and a red circle '1'. The main content area shows a table of deployments for the 'k8s-cluster' namespace 'kube-system'. The 'monitoring-grafana' deployment is highlighted with a red box. A red circle '2' is next to the 'Namespace' dropdown menu.

Name	Tag	PodsQuantity	Time Created	Action
alicloud-disk-controller	app:alicloud-disk-controller	1/1	06/27/2018,17:53:56	Details Edit Monitor More
alicloud-monitor-controller	k8s-app:alicloud-monitor-controller task:monitoring	1/1	06/28/2018,14:06:30	Details Edit Monitor More
default-http-backend	app:default-http-backend	1/1	06/27/2018,17:53:56	Details Edit Monitor More
heapster	k8s-app:heapster task:monitoring	1/1	06/27/2018,17:53:56	Details Edit Monitor More
kube-dns	k8s-app:kube-dns	1/1	06/27/2018,17:53:51	Details Edit Monitor More
monitoring-grafana	k8s-app:grafana task:monitoring	1/1	07/10/2018,16:04:34	Details Edit Monitor More

- Click the name monitoring-grafana to view the deployment status. Wait until the running status changes to Running.

Deploymentmonitoring-grafana

[← Back to List](#)

Refresh

Overview

Name:

monitoring-grafana

Namespace:

kube-system

Time Created:

2018-04-27 17:54:26

Label:

k8s-app:grafana
task:monitoring

annotation:

deployment.kubernetes.io/revision:1

Selector:

k8s-app:grafana
task:monitoring

Strategy:

RollingUpdate

Status:

Updated:1 , Unavailable:0 , Replica:1

Pods

RelatedService

Name

Status

Image

Events

monitoring-grafana-675dc8448c-drftq

Running

registry.cn-hangzhou.aliyuncs.com/acs/grafana:5.0.4

- Click Application > Service in the left-side navigation pane. Select the cluster from the Clusters drop-down list and kube-system from the Namespace drop-down list to view the external endpoint.

The external endpoint is automatically created by using the LoadBalancer type service. For developers who require more secure access policies, we recommend that you increase the security by adding the external endpoint to the IP whitelist or configuring the certificate.

Container Service - Kubernetes

Stateful Sets

Daemon Sets

Jobs

Cron Jobs

Pods

Volume Claims

Releases

Workflow

Service Mesh

Virtual Services

Discovery and Load Balancing Services

Service

Canary release

Clusters

k8s-test

Namespaces

kube-system

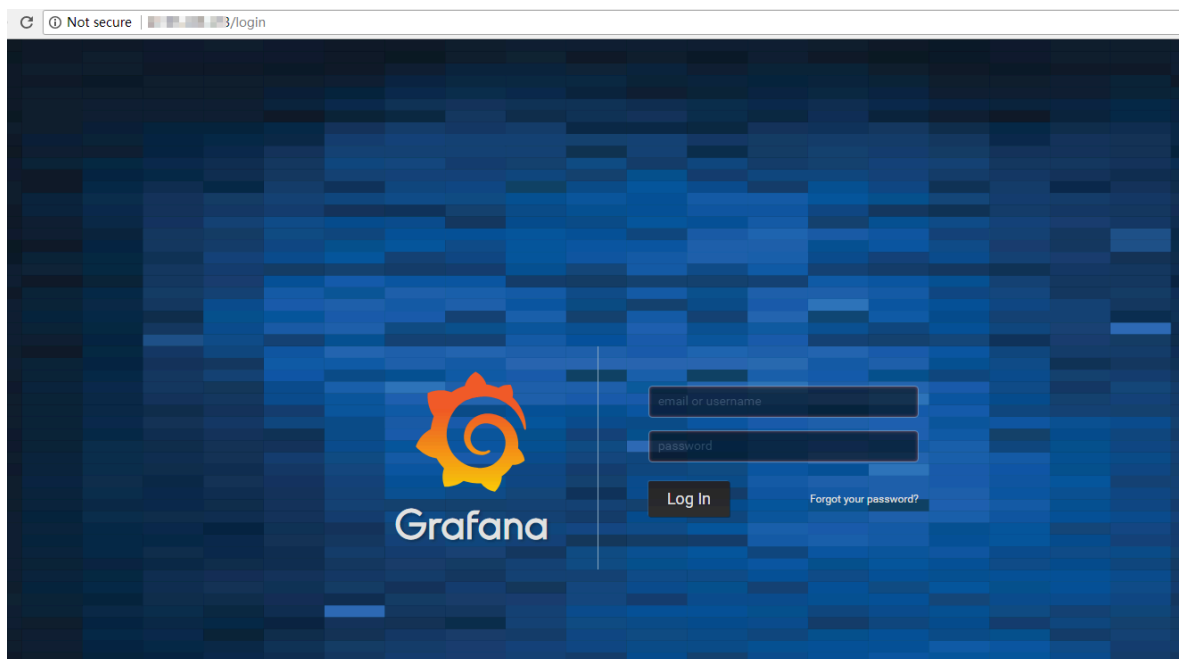
2

Search By Name

Name	Label	Type	Time Created	ClustersIP	InternalEndpoint	ExternalEndpoint	Action
heapster	kubernetes.io/cluster-service:true task:monitoring kubernetes.io/name:metrics-server	ClusterIP	05/17/2019,18:19:37		heapster:80 TCP	-	Details Update View YAML Delete
kube-dns	kubernetes.io/cluster-service:true k8s-app:kube-dns kubernetes.io/name:KubeDNS	ClusterIP	05/17/2019,18:19:36		kube-dns:53 UDP kube-dns:53 TCP	-	Details Update View YAML Delete
metrics-server	kubernetes.io/name:metrics-server	ClusterIP	05/17/2019,18:19:37		metrics-server:443 TCP	-	Details Update View YAML Delete
monitoring-grafana	-	LoadBalancer	05/30/2019,10:35:08		monitoring-grafana:80 TCP monitoring-grafana:30301 TCP		Details Update View YAML Delete

8. Click the external endpoint at the right of the monitoring-grafana service to log on to the Grafana monitoring page.

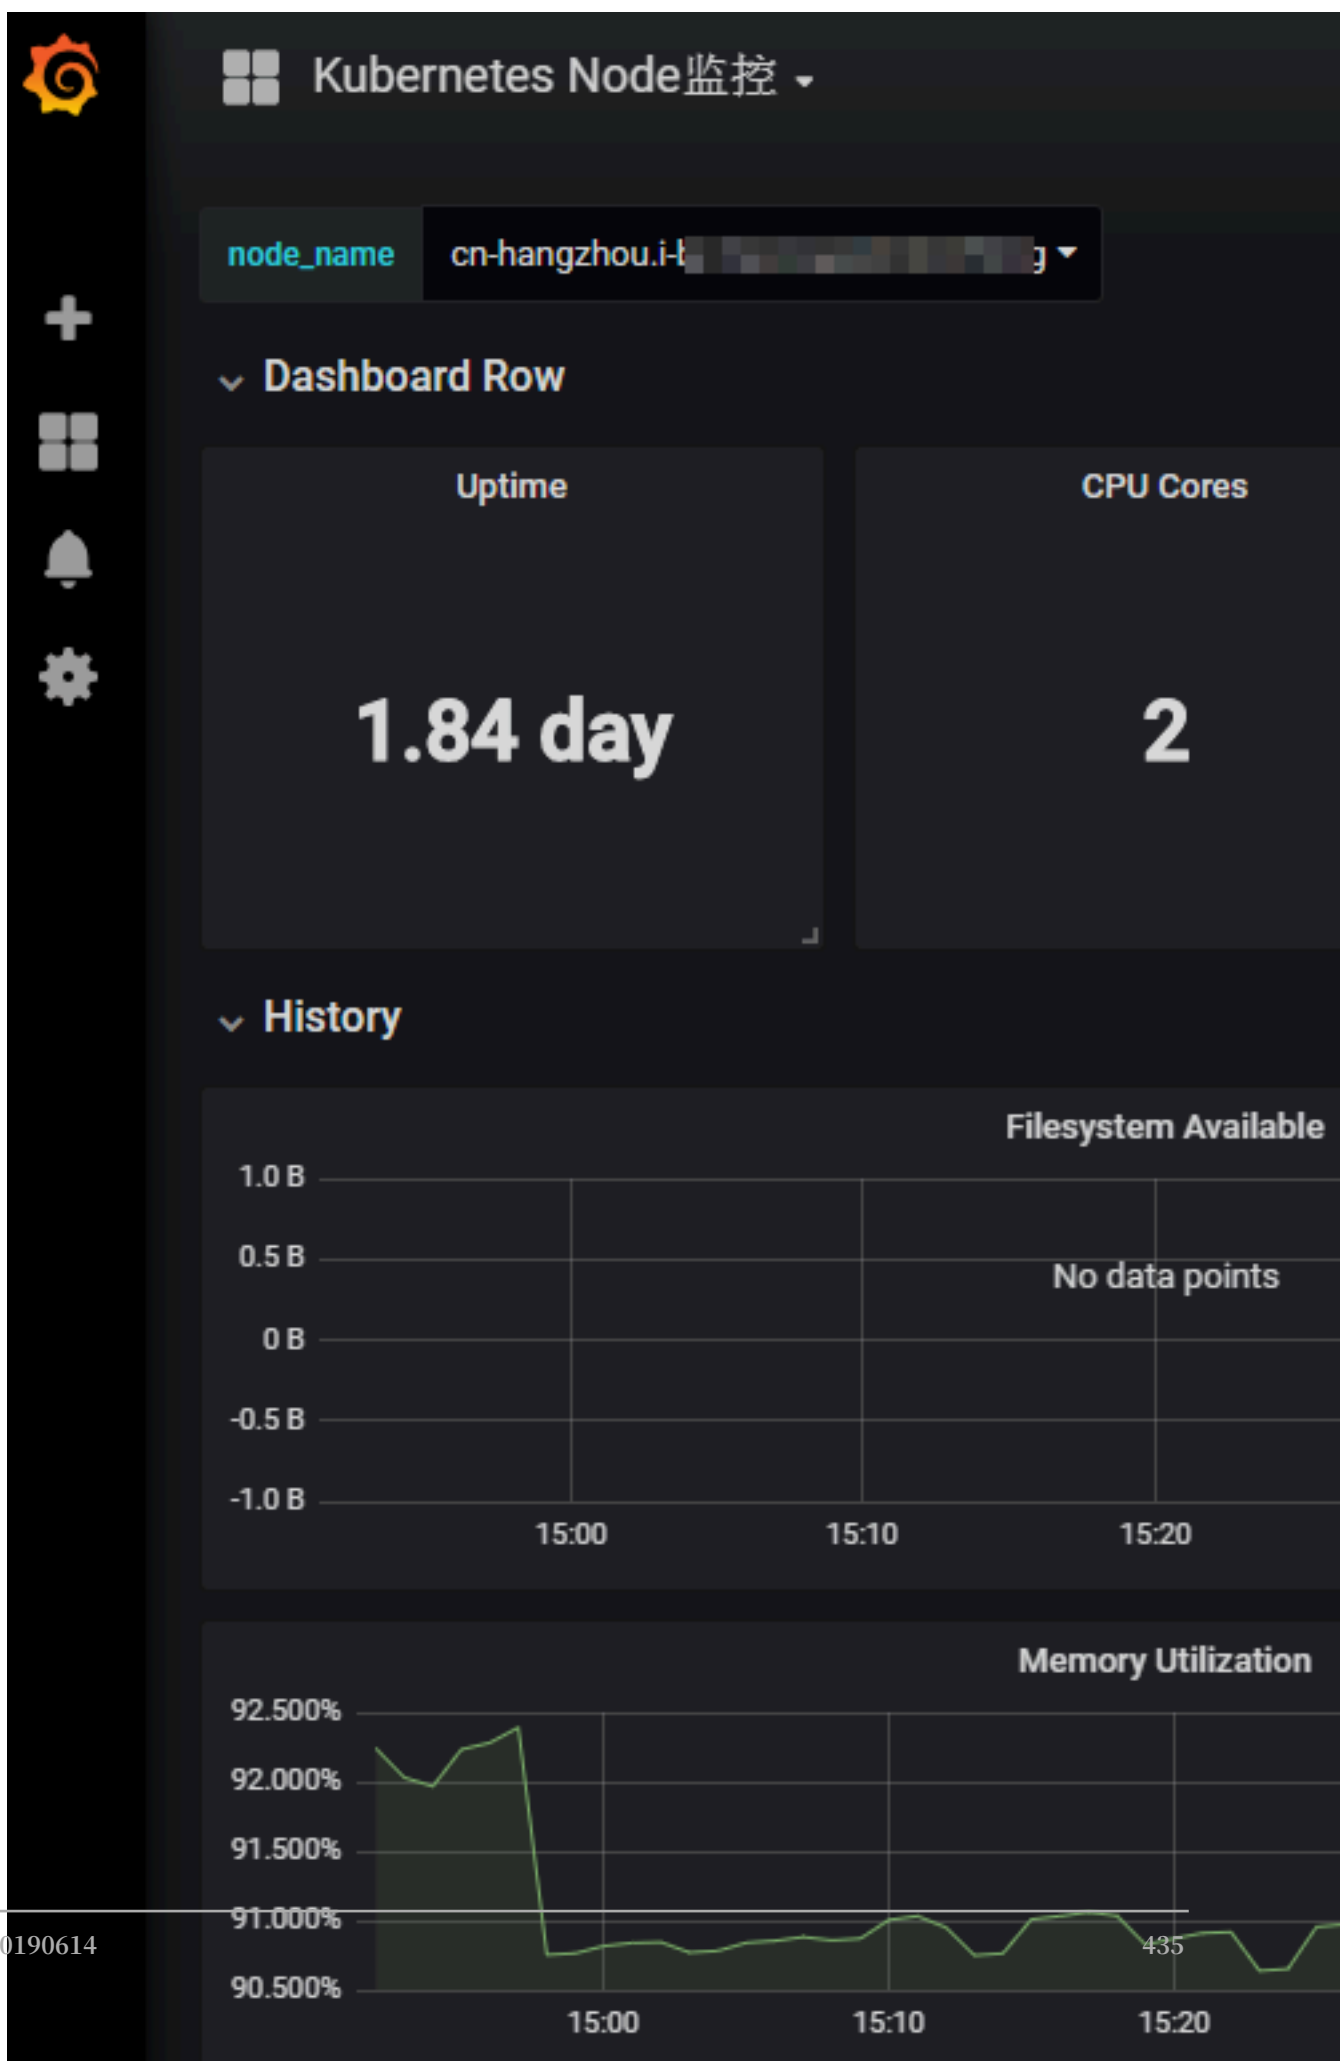
By default, the username and password of Grafana are both admin. We recommend that you change the password after the logon.

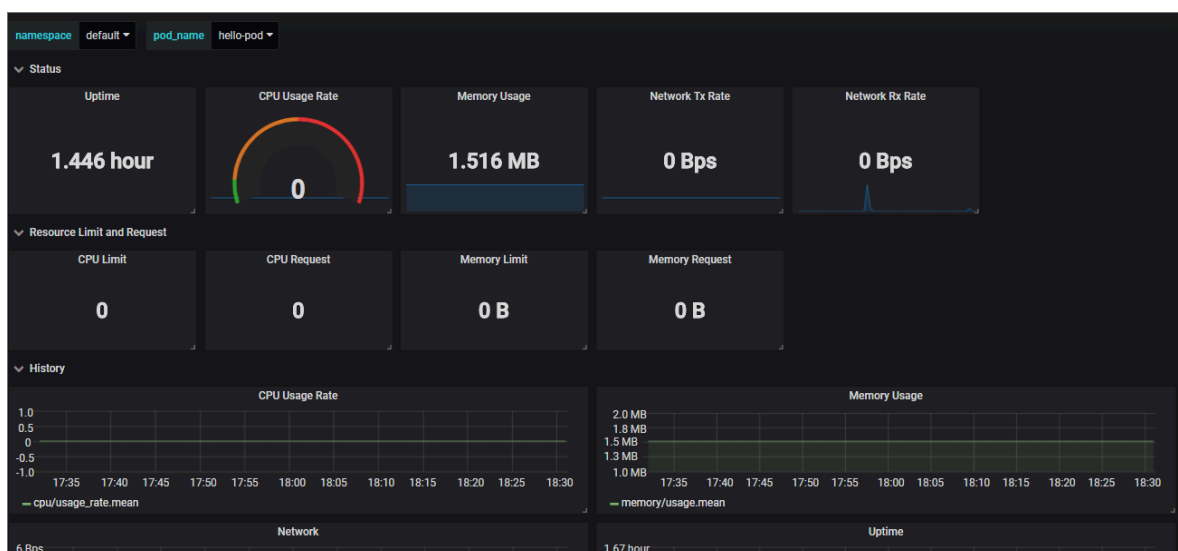


9. Select the built-in monitoring templates to view the monitoring dashboards of the pod and node.

In this example, the Grafana has two built-in templates, one for displaying physical resources at the node level, and one for displaying resources related to the pod.

Developers can also perform more complex presentations by adding custom dashboards or configure resource alarms based on Grafana.





1.15.5 Use an HPA auto scaling container

Alibaba Cloud Container Service supports the rapid creation of HPA-enabled applications on the console interface to achieve auto scaling of container resources. You can also configure it by defining the yaml configuration of Horizontal Pod Autoscaling (HPA).

Prerequisites

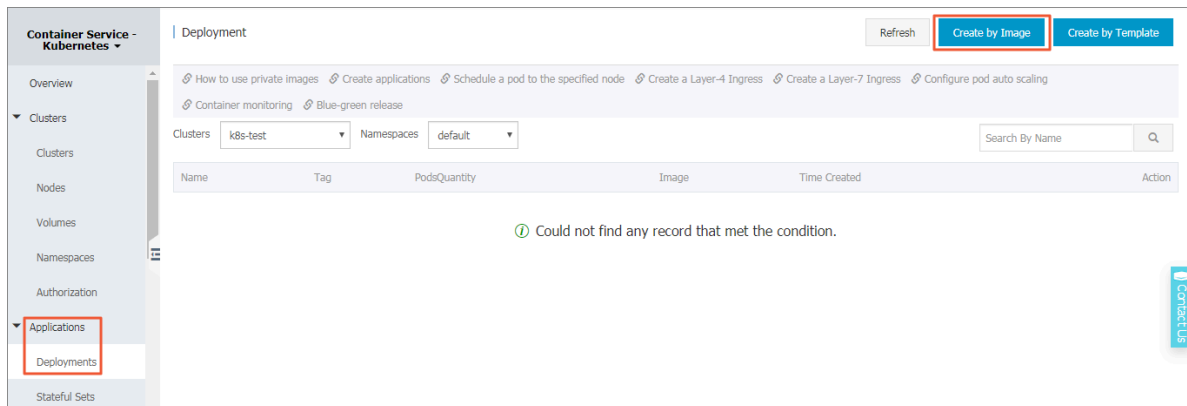
- You have created a Kubernetes cluster. For more information, see [Create a Kubernetes cluster](#).
- You have successfully connected to the master node of the Kubernetes cluster.

Method 1 Create an HPA application in the Container Service console

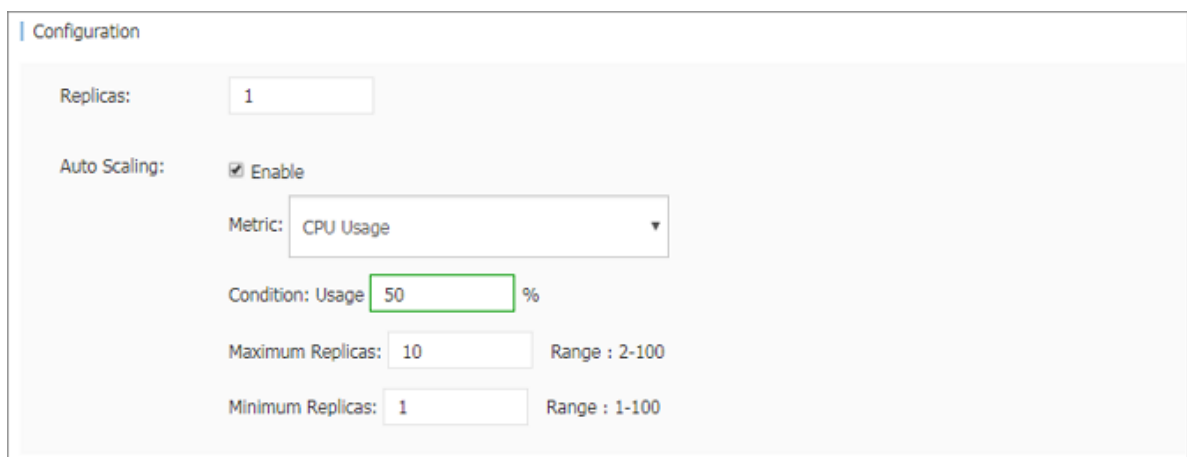
In Alibaba Cloud Container Service, HPA has been integrated. You can easily create it through the Container Service console.

1. Log on to the [Container Service console](#).

2. In the left-side navigation pane under Container Service-Kubernetes, choose **Applications > Deployments**. Then, click **Create by Image** in the upper-right corner.



3. Enter the application name, select the cluster and namespace, and click **Next**.
4. Configure the application settings. Set the number of replicas, select the **Enable** box for **Automatic Scaling**, and configure the settings for scaling.
 - **Metric:** CPU and memory. Configure a resource type as needed.
 - **Condition:** The percentage value of resource usage. The container begins to expand when the resource usage exceeds this value.
 - **Maximum Replicas:** The maximum number of replicas that the deployment can expand to.
 - **Minimum Replicas:** The minimum number of replicas that the deployment can contract to.



5. Configure the container. Select an image and configure the required resources. Click **Next**.



Note:

You must configure the required resources for the deployment. Otherwise, container auto scaling cannot be achieved.

container0

Image Name: Image Version: ☐ Always pull image

[Select image](#) [Select image version](#)

Resource Limit: CPU Memory

Resource Request: CPU Memory

Init Container ☐

6. In the Access Control page, do not configure any settings in this example. Click **Create directly**.

Now a deployment that supports HPA has been created. You can view the auto scaling group information in the details of your deployment.

Overview

Name:	nginx-deployment
Namespace:	default
Time Created:	2018-08-23 10:07:23
Label:	app:nginx
annotation:	deployment.kubernetes.io/revision:1
Selector:	app:nginx
Strategy:	RollingUpdate
Status:	Updated:1 , Unavailable:1 , Replica:1

Trigger 1. You can only have one of each trigger type. [Create Trigger](#)

No trigger is available at the moment. Click "Create Trigger" in the upper-right corner.

CPU

0 10:09:00

● CPU usage(Cores)

Memory

0 10:08:00 10:09:00

● Memory usage(Gi)

Horizontal Pod Autoscaler

Name	Target Utilization	Minimum Replicas	Maximum Replicas	Created At	Action
nginx	cpu:70%	1	10	08/23/2018,10:07:23	Edit Delete

7. In the actual environment, the application scales according to the CPU load. You can also verify auto scaling in the test environment. By performing a CPU pressure

test on the pod, you can find that the pod can complete the horizontal expansion in half a minute.

Horizontal Pod Autoscaler		
Name	Status	Image
k8s-hpa-deployment-f87696b8b-jpr15	Running	nginx:latest

Method 2 Use kubectl commands to configure container auto scaling

You can also manually create an HPA by using an orchestration template and bind it to the deployment object to be scaled. Use the `kubectl` command to complete the container auto scaling configuration.

The following is an example of an Nginx application. Execute the `kubectl`

`create -f xxx . yml` command to create an orchestration template for the deployment as follows:

```
apiVersion : apps / v1beta2 # for versions before 1 . 8 . 0
use apps / v1beta1
kind : Deployment
metadata :
  name : nginx
  labels :
    app : nginx
spec :
  replicas : 2
  selector :
    matchLabels :
      app : nginx
  template :
    metadata :
      labels :
        app : nginx
    spec :
      containers :
        - name : nginx
          image : nginx : 1 . 7 . 9 # replace it with your
exactly < image_name : tags >
          ports :
            - containerPort : 80
          resources :
            requests :
              # This parameter
must be configured . Otherwise , the HPA cannot operate .
              cpu : 500m
```

Create an HPA. Configure an object to which the current HPA is bound by using `scaleTargetRef`. In this example, the object is the deployment named `nginx`.

```
apiVersion : autoscaling / v2beta1
kind : HorizontalPodAutoscaler
metadata :
  name : nginx - hpa
  namespace : default
spec :
```

```

scaleTargetRef :                                ## Bind the HPA
to a deployment named nginx
  apiVersion : apps / v1beta2
  kind : Deployment
  name : nginx
  minReplicas : 1
  maxReplicas : 10
  metrics :
  - type : Resource
    resource :
      name : cpu
      targetAverageUtilization : 50

```

**Note:**

The HPA needs to configure the request resource for the pod. The HPA does not operate without the request resource.

Warnings similar to the following are displayed when you execute `kubectl`

```
describe hpa [ name ]:
```

```

Warning   FailedGetResourceMetric   2m (x6 over 4m )
horizontal - pod - autoscaler missing request for cpu on
container nginx in pod default / nginx - deployment - basic
- 75675f5897 - mqzs7

Warning   FailedComputeMetricsReplicas 2m (x6 over
4m ) horizontal - pod - autoscaler failed to get cpu
utilization : missing request for cpu on container
nginx in pod default / nginx - deployment - basic - 75675f5

```

After creating the HPA, execute the `kubectl describe hpa [name]` command again. You can see the following message, which indicates that the HPA is running normally.

```

Normal Successful Rescale 39s horizontal - pod - autoscaler
New size : 1 ; reason : All metrics below target

```

When the usage of Nginx pod exceeds 50% set in this example, the container expands horizontally. When the usage of Nginx pod drops below 50%, the container contracts.

1.15.6 Monitor a Kubernetes cluster and send alarm notifications by using DingTalk

After you deploy a robot in a DingTalk group, the cluster sends a notification of an exception event to the DingTalk group through the robot, implementing real-time monitoring and alarming for cluster exception events.

Context

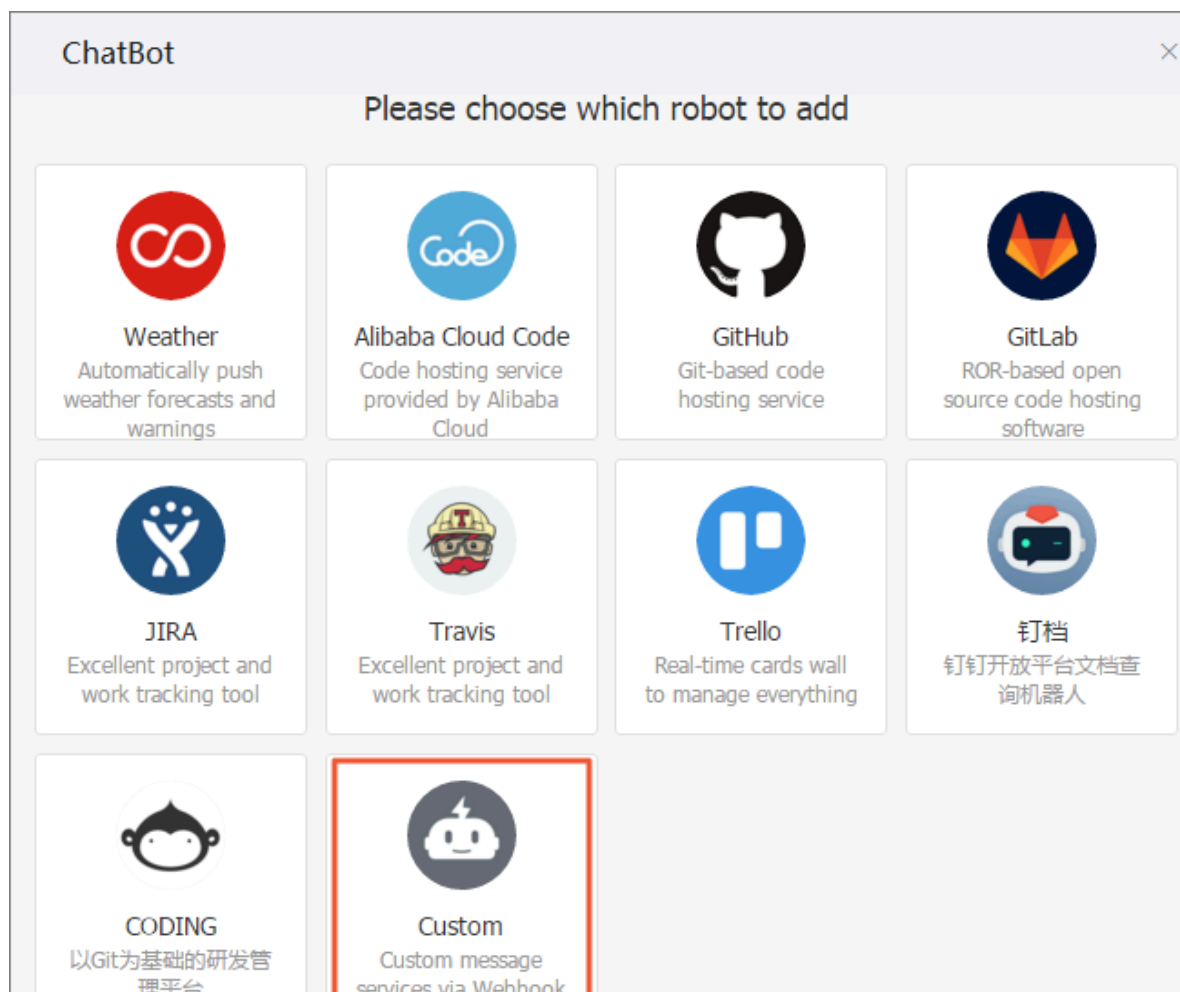
- You have created a DingTalk group .

- You have created a Kubernetes cluster. For more information, see [Create a Kubernetes cluster](#).

Procedure


1. Click the  icon in the upper-right corner of the DingTalk group.

2. Click ChatBot. On the ChatBot page, select a robot. Select a Custom robot.




3. On the Robot details page, click Add.


Robot details


Custom

Brief Info: Use the Chatbot API to push any service message that you need to DingTalk

Message preview:


VIP monitoring alarm ChatBot
The failure rate of sending message is higher than 5%, module 202, network type 4G.
@Yinan Urgent task.



Plan reminder ChatBot
[P3][Online][Early plan]
- The number of the mobile home tab shows downgrade
- Operator: Xumo

Cancel

Add


4. Configure the following parameters for a robot and then click Finished:

Configuration	Description
Edit profile picture	(Optional) Set a profile picture for the robot.
ChatBot Name	The robot name.
Add to Group	The DingTalk group to which the robot is added to.
Enable the outgoing function	<p>(Optional) By perform the @robot operation, you can send messages to a specified external service as well as return response results of the external service to the group.</p> <div> Note: We recommend that you do not enable this function.</div>

Configuration	Description
POST address	<p>The HTTP service address that receives messages.</p> <p> Note: You can configure this parameter after you enable the outgoing function.</p>
Token	<p>The key used to verify that a request is from DingTalk.</p> <p> Note: You can configure this parameter after you enable the outgoing function.</p>

5. Click Copy to copy the webhook address.

Add Robot



1. Add robot✓

2. Set up webhook, click setting instruction and check how to make robot effective

webhook :

Copy

Finished

Setting ins...



Note:

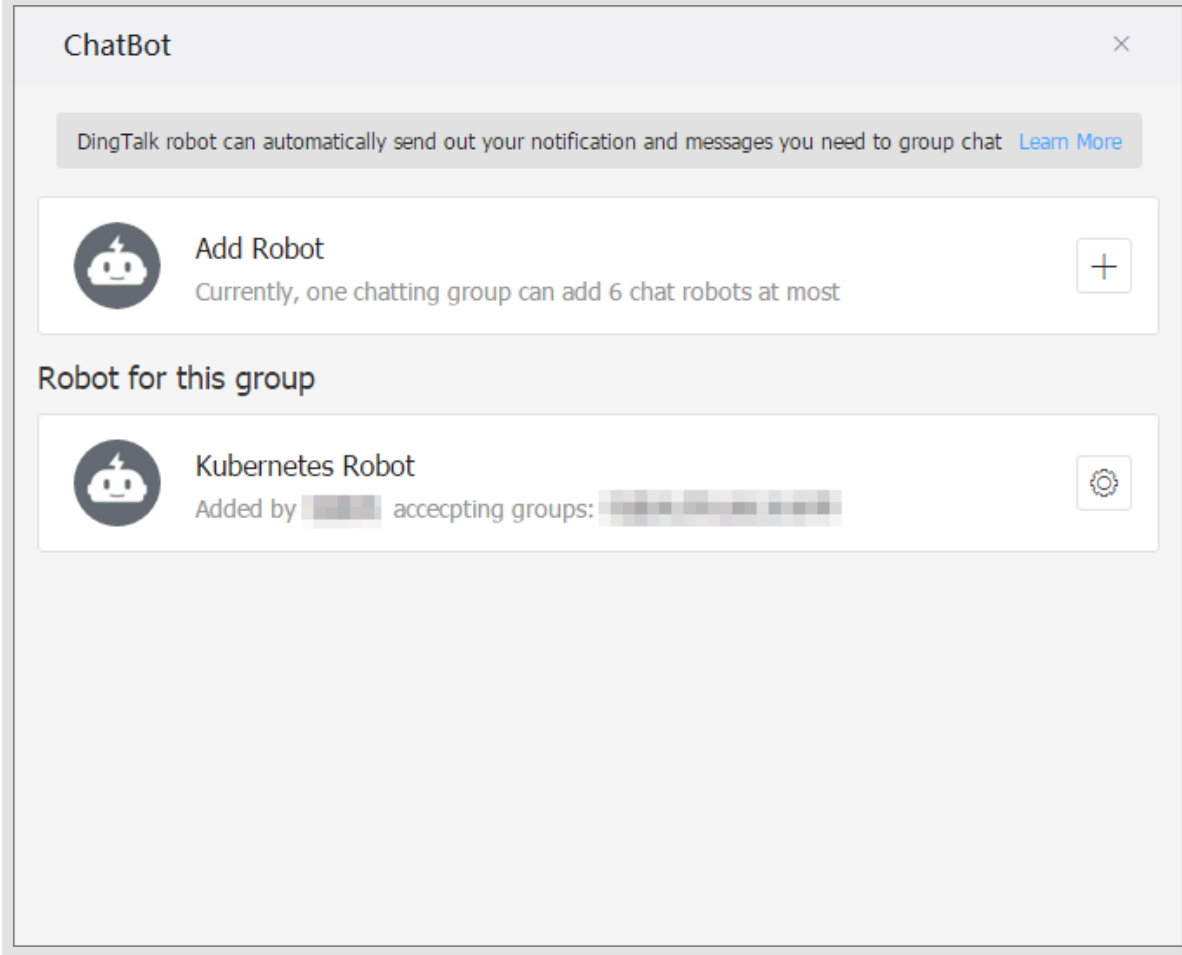
On the ChatBot page, click the



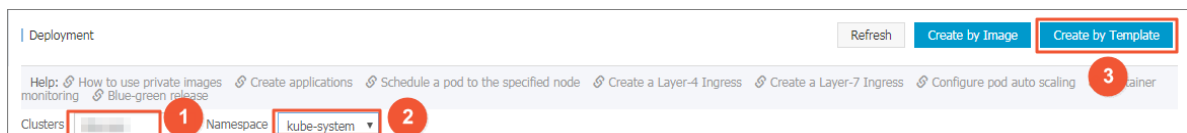
icon at the right of a robot and then you

can perform following operations:

- Modify the profile picture and name of the robot.
- Open or Close notifications.
- Reset the webhook address.
- Remove the robot.



6. Log on to the [Container Service console](#).
7. Under the Kubernetes menu, click Application > Deployment in the left-side navigation pane.
8. Select a cluster, select the kube-system namespace, and click Create by Template in the upper-right corner.



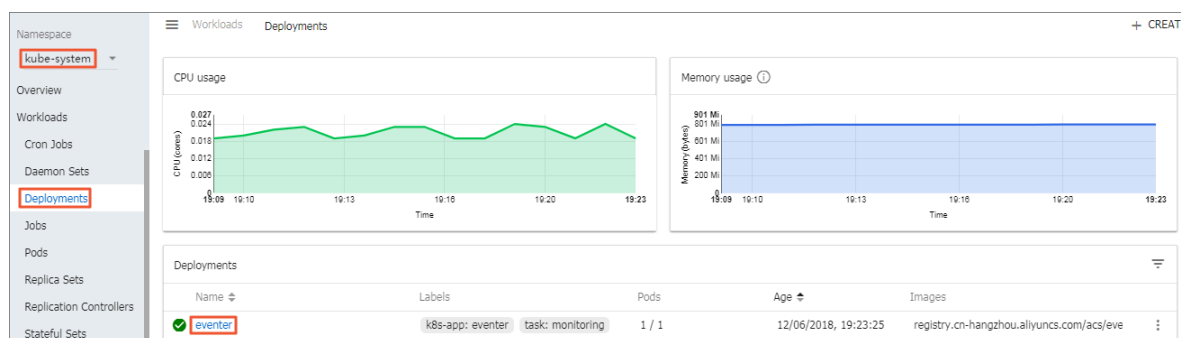
9. Configure a template based on the following parameters, and then click Deploy.

Configuration	Description
Clusters	Select a cluster.

Configuration	Description
Namespace	Select a namespace to which resource object belongs. The default namespace is default. Select kube-system.
Sample template	Alibaba Cloud Container Service provides Kubernetes YAML sample templates of many resource types for you to deploy resource objects quickly. You can write your own template based on the format requirements of Kubernetes YAML orchestration to describe the resource type you want to define. Select Custom.

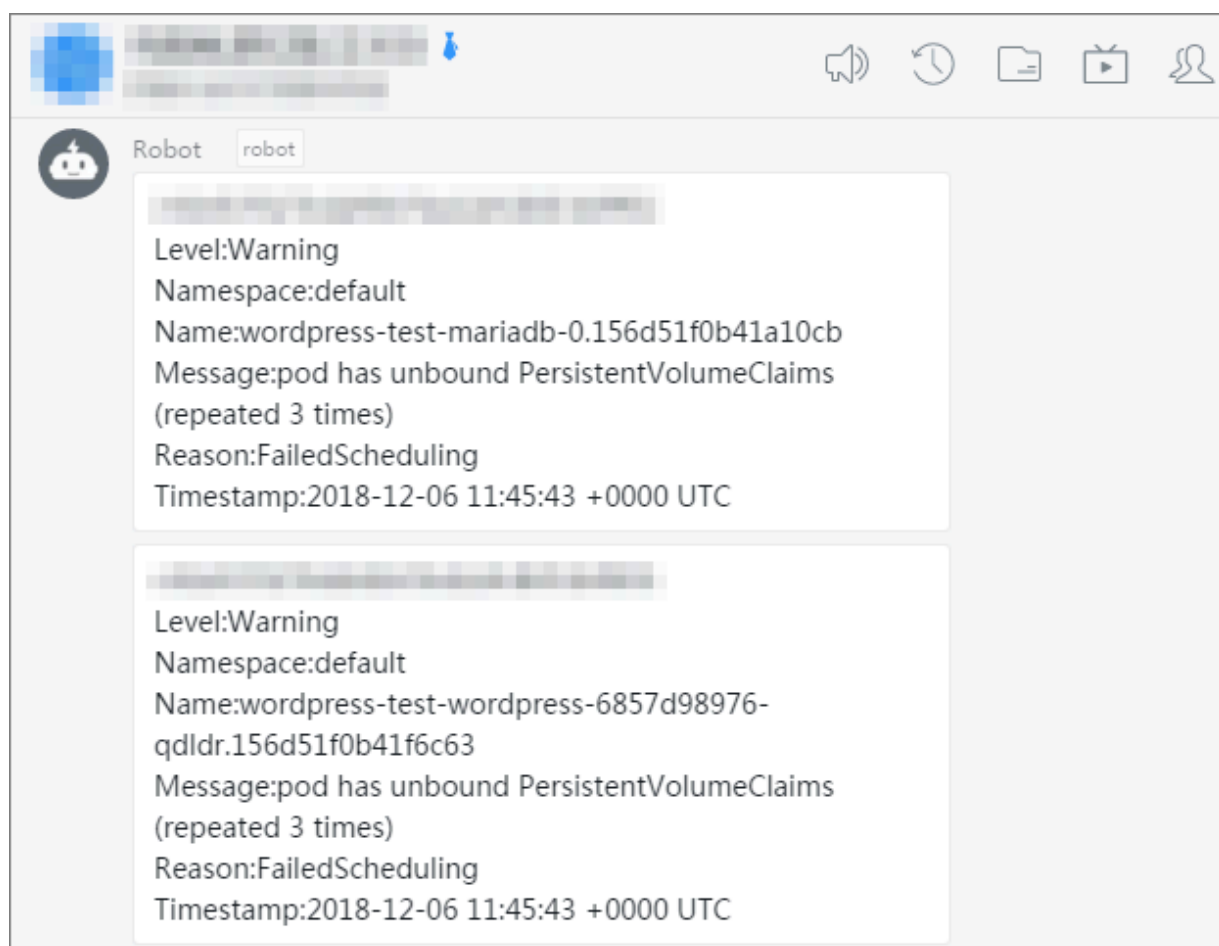
Configuration	Description
Template	<p>Enter the following custom content:</p> <pre> apiVersion : extensions / v1beta1 kind : Deployment metadata : name : eventer namespace : kube - system spec : replicas : 1 template : metadata : labels : task : monitoring k8s - app : eventer annotations : scheduler.alpha.kubernetes.io / critical - pod : '' spec : serviceAccount : admin containers : - name : eventer image : registry.cn-hangzhou.aliyuncs.com / acs / eventer : v1.6.0 imagePullPolicy : IfNotPresent command : - / eventer - -- source = kubernetes : https : // kubernetes . default - -- sink = dingtalk : [your_webhook_url]& label = [your_cluster_id]& level = [Optional parameters are : Normal and Warning . The default is Warning .] # The level field can be set to Normal or Warning . The default is Warning . When the level field is set to Normal , alarm notifications of the Normal and Warning levels can be received in the DingTalk group . When you do not set the level field or set the level field to Warning , only alarm notifications of the Warning level can be received in the DingTalk group . </pre>

On the Cluster List page, click Dashboard at the right of the cluster. On the Dashboard, select kube-system from the drop-down list of Namespace, and click Deployments in the left-side navigation pane. The deployed eventer is displayed.



Result

The eventer takes effect 30 seconds after you complete the deployment. When an event exceeds the threshold level, you receive the following alarm notifications in the DingTalk group.



1.16 Security management

1.16.1 Security

Authorization

Kubernetes clusters support authorizing RAM users to perform operations on clusters.

For more information, see [Use the Container Service console as a RAM user](#).

Full-link TLS certificates

The following communication links in Container Service Kubernetes clusters are verified by TLS certificates to prevent the communication from being eavesdropped or tampered:

- `kubelet` on worker nodes actively communicates with `apiserver` on master nodes
- `apiserver` on master nodes actively communicates with `kubelet` on worker nodes

During initialization, the master node uses SSH tunnels to connect to the SSH service of other nodes (port 22) for initialization.

Native secret & RBAC support

Kubernetes secrets are used to store sensitive information such as passwords, OAuth tokens, and SSH keys. Using plain text to write sensitive information to a pod YAML file or a Docker image may leak the information, while using secrets avoids such security risks effectively.

For more information, see [Secret](#).

Role-Based Access Control (RBAC) uses the Kubernetes built-in API group to drive authorization and authentication, which allows you to use APIs to manage pods that correspond to different roles, and the access permissions of roles.

For more information, see [Using RBAC authorization](#).

Network policy

In a Kubernetes cluster, pods on different nodes can communicate with each other by default. In some scenarios, to reduce risks, the network intercommunication among different business services is not allowed and you must introduce the network

policy. In Kubernetes clusters, you can use the Canal network driver to implement the support for network policy.

Image security scan

Kubernetes clusters can use Container Registry to manage images, which allows you to perform image security scan.

Image security scan identifies the security risks in images quickly and reduces the possibility of applications running on your Kubernetes cluster being attacked.

For more information, see [Image security scan](#).

Security group and Internet access

By default, each newly created Kubernetes cluster is assigned a new security group with the minimal security risk. This security group only allows ICMP for the Internet inbound.

By default, you cannot use Internet SSH to access your clusters. To use Internet SSH to connect to the cluster nodes, see [Access Kubernetes clusters by using SSH](#).

The cluster nodes access the Internet by using the NAT Gateway, which further reduces the security risks.

1.16.2 Kube-apiserver audit logs

In a Kubernetes cluster, apiserver audit logs are important for cluster Operation & Maintenance (O&M) because they record daily operations of different users. This topic describes how to configure the apiserver audit logs of an Alibaba Cloud Kubernetes cluster, and how to collect and analyze audit logs through Log Service, and how to customize audit log alarm rules.

Configurations of apiserver audit logs

The apiserver audit function is enabled by default when you create a Kubernetes cluster. Relevant parameters and description are as follows:



Note:

Log on to the Master node, and the directory of the apiserver configuration files is /
`etc / kubernetes / manifests / kube - apiserver . yaml .`

Configuration	Description
<code>audit-log-maxbackup</code>	The maximum fragment of audit logs stores 10 log files.
<code>audit-log-maxsize</code>	The maximum size of a single audit log is 100 MB.
<code>audit-log-path</code>	The audit log output path is <code>/var/log/kubernetes/kubernetes.audit</code> .
<code>audit-log-maxage</code>	The longest storage period of audit logs is seven days.
<code>audit-policy-file</code>	Configuration policy file of audit logs. The directory is <code>/etc/kubernetes/audit-policy.yml</code> .

Log on to the Master node machine. The directory of the audit log configuration policy file is `/etc/kubernetes/audit-policy.yml`. The content of the file is as follows:

```
apiVersion: audit.k8s.io/v1beta1 # This is required.
kind: Policy
# We recommend that you do not generate audit events
# for all requests in RequestReceived stage.
omitStages:
- "RequestReceived"
rules:
# The following requests are manually identified as
# high - volume and low - risk .
# Therefore , we recommend that you drop them .
- level: None
  users: ["system:kube-proxy"]
  verbs: ["watch"]
  resources:
    - group: "" # core
      resources: ["endpoints", "services"]
- level: None
  users: ["system:unsecured"]
  namespaces: ["kube-system"]
  verbs: ["get"]
  resources:
    - group: "" # core
      resources: ["configmaps"]
- level: None
  users: ["kubelet"] # legacy kubelet identity
  verbs: ["get"]
  resources:
    - group: "" # core
      resources: ["nodes"]
- level: None
  userGroups: ["system:nodes"]
  verbs: ["get"]
  resources:
```

```

- group : "" # core
  resources : [" nodes "]
- level : None
  users :
    - system : kube - controller - manager
    - system : kube - scheduler
    - system : serviceaccount : kube - system : endpoint -
controller
  verbs : [" get ", " update "]
  namespaces : [" kube - system "]
  resources :
    - group : "" # core
      resources : [" endpoints "]
- level : None
  users : [" system : apiserver "]
  verbs : [" get "]
  resources :
    - group : "" # core
      resources : [" namespaces "]
# We recommend that you do not log these read -
only URLs .
- level : None
  nonResourceURLs :
    - / healthz *
    - / version
    - / swagger *
# We recommend that you do not log events requests
.
- level : None
  resources :
    - group : "" # core
      resources : [" events "]
# Secrets , ConfigMaps , and TokenReviews can contain
sensitive and binary data .
# Therefore , they are logged only at the Metadata
level .
- level : Metadata
  resources :
    - group : "" # core
      resources : [" secrets ", " configmaps "]
    - group : authentication.k8s.io
      resources : [" tokenreviews "]
# Get responses can be large ; skip them .
- level : Request
  verbs : [" get ", " list ", " watch "]
  resources :
    - group : "" # core
    - group : " admissionregistration.k8s.io "
    - group : " apps "
    - group : " authentication.k8s.io "
    - group : " authorization.k8s.io "
    - group : " autoscaling "
    - group : " batch "
    - group : " certificates.k8s.io "
    - group : " extensions "
    - group : " networking.k8s.io "
    - group : " policy "
    - group : " rbac.authorization.k8s.io "
    - group : " settings.k8s.io "
    - group : " storage.k8s.io "
# Default level for known APIs .
- level : Response
  resources :
    - group : "" # core

```

```

- group : " admissionr egistratio n . k8s . io "
- group : " apps "
- group : " authentica tion . k8s . io "
- group : " authorizat ion . k8s . io "
- group : " autoscalin g "
- group : " batch "
- group : " certificat es . k8s . io "
- group : " extensions "
- group : " networking . k8s . io "
- group : " policy "
- group : " rbac . authorizat ion . k8s . io "
- group : " settings . k8s . io "
- group : " storage . k8s . io "
# Default level for all other requests .
- level : Metadata

```



Note:

- Logs are not recorded immediately after requests are received. Log recording starts only after the response body header is sent.
- The following requests or operations are not audited: redundant kube-proxy watch requests, GET requests from kubelet and system:nodes for nodes, operations performed on endpoints by kube components in the kube-system, and GET requests from the apiserver for namespaces.
- Read-only urls such as `/ healthz *`, `/ version *`, and `/ swagger *` are not audited.
- Logs of interfaces of secrets, configmaps, and tokenreviews are set to the metadata level because they might contain sensitive information or binary files . For logs of this level, only the user, timestamp, request resources, and request actions of the request event are audited. The request body and the response body are not audited.
- For sensitive interfaces such as authentication, rbac, certificates, autoscaling, and storage, the corresponding request bodies and response bodies are audited according to the read and write requests.

View audit log reports

A Kubernetes cluster that runs on Alibaba Cloud Container Service has three audit log reports that provide the following information:

- Operations performed by all users and system components on the cluster
- The source IP address of each operation, the area to which a source IP addresses belongs, and the source IP address distribution
- Detailed operation charts of all resources

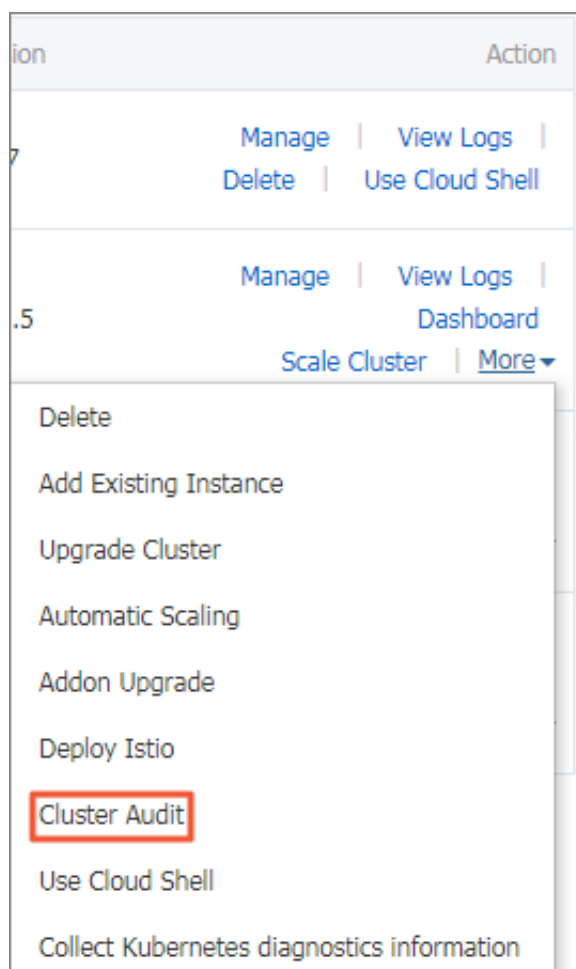
- Operation charts of each sub-account
- Charts of important operations such as logging on to a container, accessing a secret, and removing resources

**Note:**

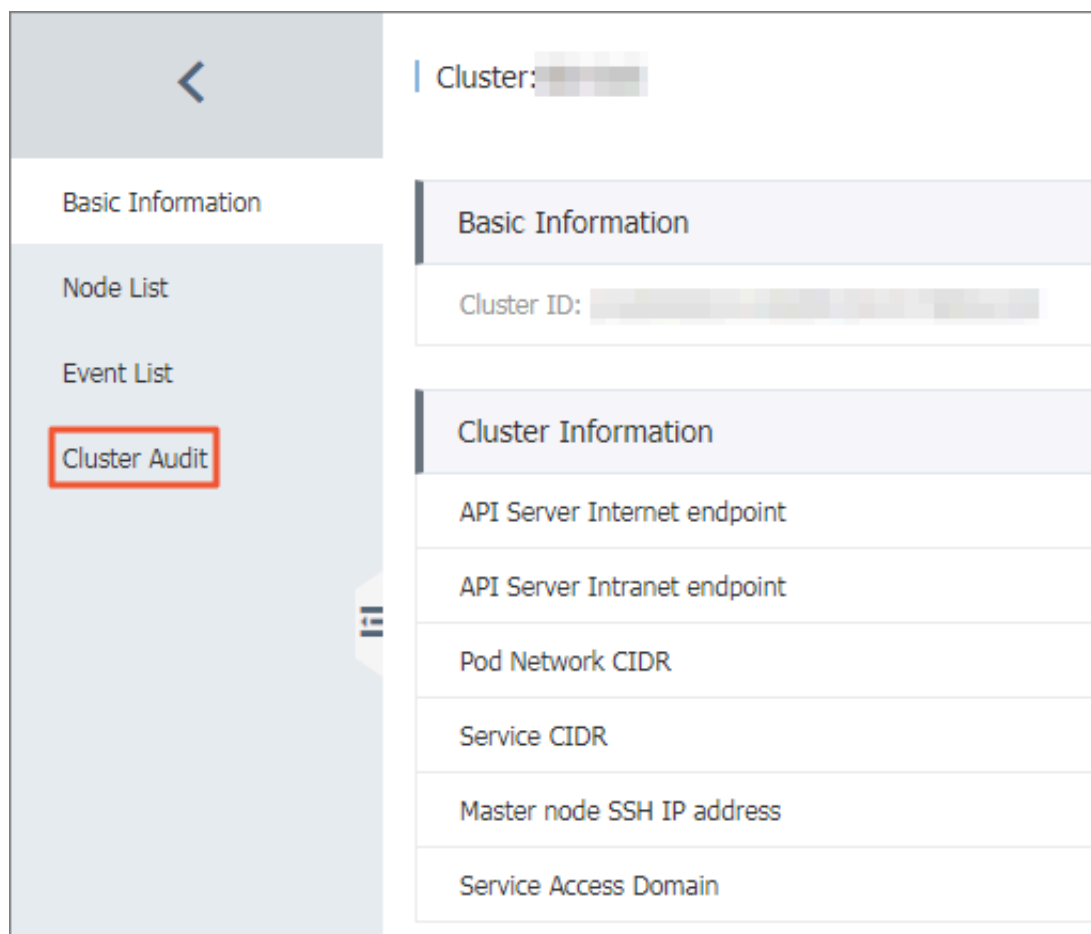
- For Kubernetes clusters created after January 13, 2019, if you activate Log Service for the clusters, the system automatically enables audit log report functions. If audit log report functions are disabled for a Kubernetes cluster, see [Manually enable audit log report functions](#).
- We recommend that you do not modify audit log reports. If you want to customize audit log reports, you can create new reports in the [Log Service console](#).

You can access audit log reports by using either of the following two methods:

- Log on to the [Container Service console](#). In the action column of the target cluster, choose **More > Cluster Audit**.



- Log on to the [Container Service console](#). Click the target cluster name, and then click Cluster Audit in the left-side navigation pane.

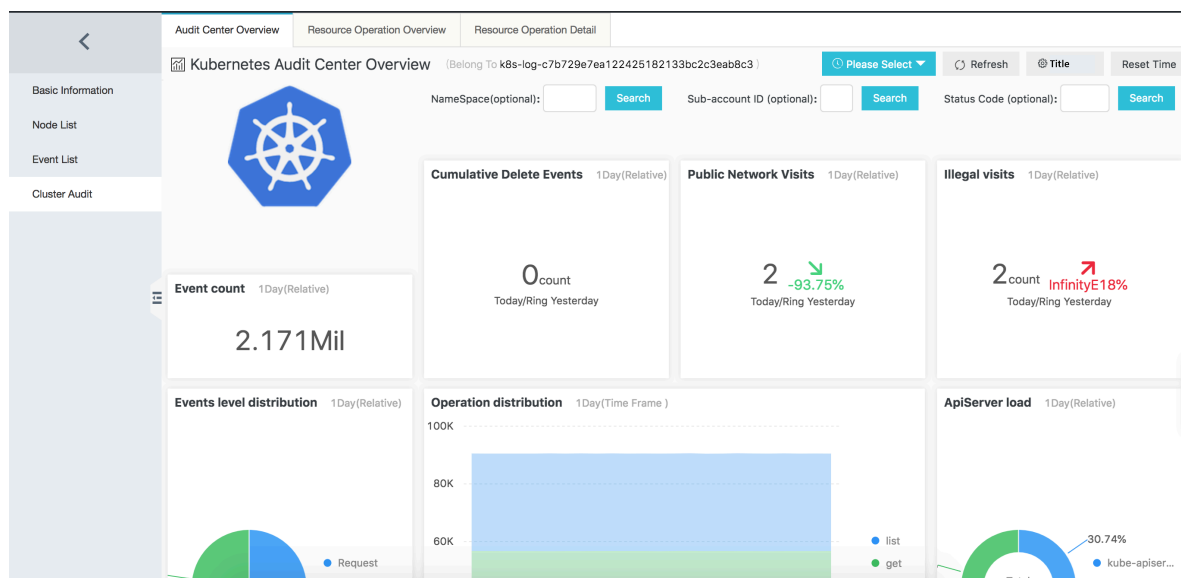


Audit log report overview

The following three apiserver audit log reports are available: Audit Center Overview, Resource Operation Overview, and Resource Operation Detail.

- Audit Center Overview

This report displays an overview of the Kubernetes cluster events and the detailed information about important events, such as public network visits, command execution, resource removal, and secret visits.

**Note:**

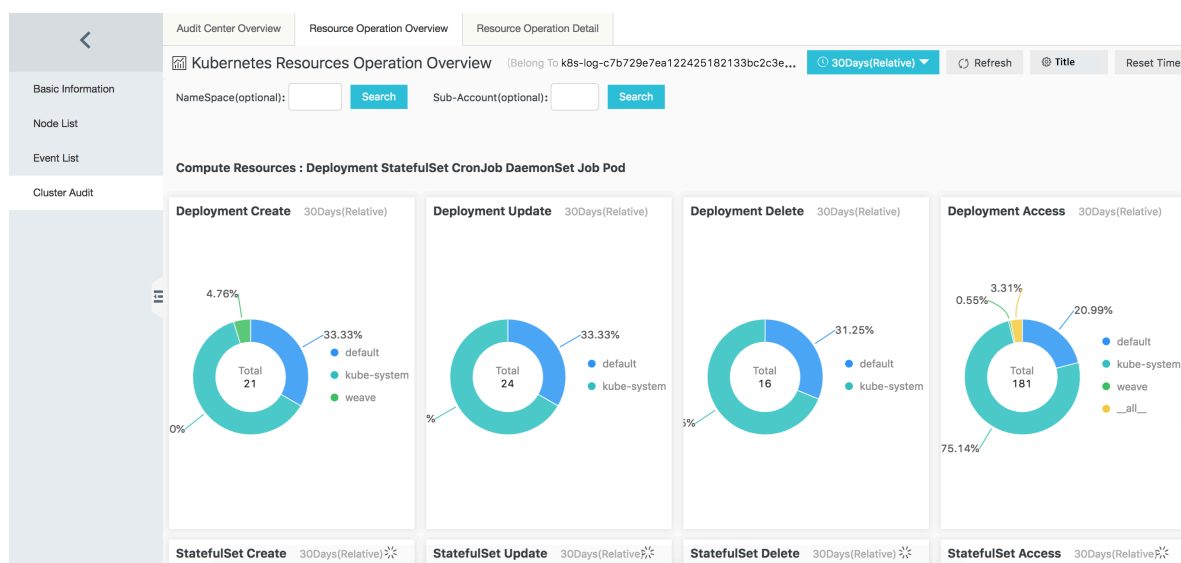
By default, this report displays statistics for one week. You can customize a statistics time range. In addition, you can filter events by specifying one or multiple factors, such as a namespace, a sub-account ID, and a status code.

• Resource Operation Overview

This report displays the operation statistics information about computing resources, network resources, and storage resources of a Kubernetes cluster.

Operations include creation, update, removal, and access.

- Computing resources include deployment, StatefulSet, CronJob, DaemonSet, Job, and pod.
- Network resources include service and Ingress.
- Storage resources include ConfigMap, secret, and Persistent Volume Claim.



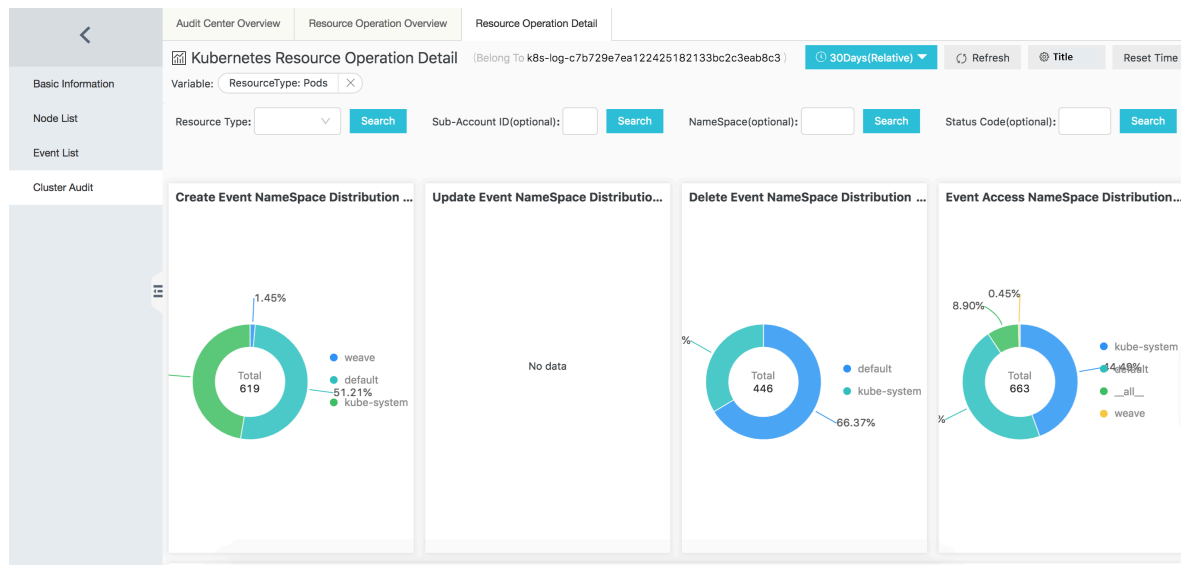
Note:

- By default, this report displays statistics for one week. You can customize a statistics time range. In addition, you can filter events by specifying one or both of the following factors: a namespace or a sub-account ID.
- If you want to view the detailed operation events of a resource, we recommend that you use Resource Operation Detail.

• Resource Operation Detail

This report displays detailed operation information of a Kubernetes cluster resource. You must select or enter a resource type to view detailed operation

information in time. This report displays the total number of operation events, namespace distribution, success rate, timing trend, and specific operation charts.

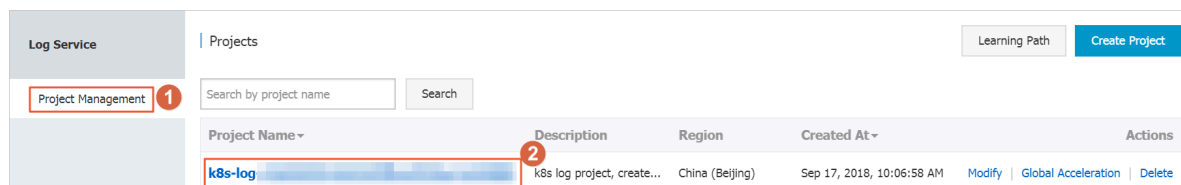


Note:

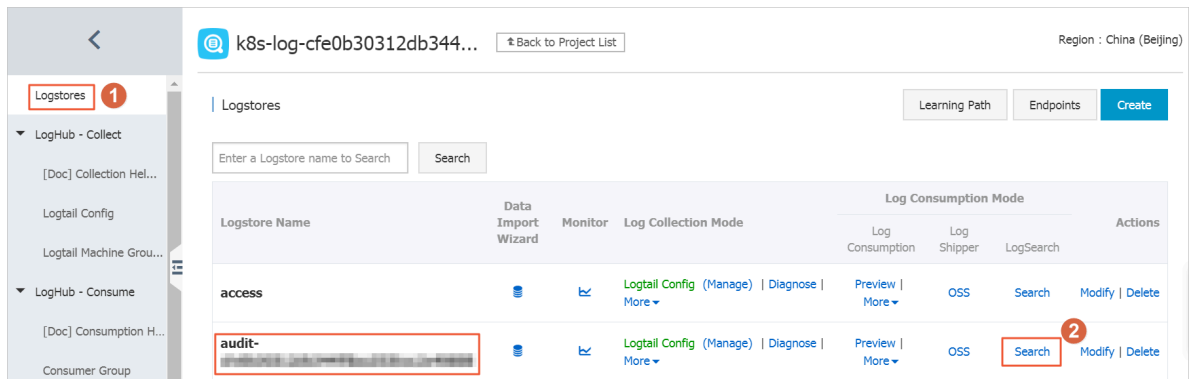
- If you want to view a CRD resource registered in Kubernetes or any other resources not listed in the report, you can enter the plural form of the target resource. For example, to view a CRD resource named AliyunLogConfig, you can enter AliyunLogConfigs.
- By default, this report displays statistics for one week. You can customize a statistics time range. In addition, you can filter events by specifying one or multiple factors, such as a namespace, a sub-account ID, and a status code.

logs, you can log on to Log Service to view detailed log records.

1. Log on to the [Log Service console](#).
2. In the left-side navigation pane, click Project Management, select the Project configured when you create the cluster, and then click the Project name.



3. On the Logstores page, find the Logstore named `audit-${clusterid}` and click Search at the right side of the Logstore. The audit logs of the cluster are stored in this Logstore.



Note:

- When you create a Kubernetes cluster, your specified log Project automatically creates a Logstore named `audit-${clusterid}`.
- The audit log Logstore index is set by default. We recommend that you do not modify the index. Otherwise, the audit log reports become invalid.

To search for an audit log, you can use one of the following methods:

- To query a sub-account operation record, enter the sub-account ID and then click Search & Analysis.
- To query operations on a resource, enter the resource name and click Search & Analysis.
- To filter out operations performed by system components, enter `NOT user .`
`username : node NOT user . username : serviceaccount NOT`
`user . username : apiserver NOT user . username : kube -`
`scheduler NOT user . username : kube - controller - manager` , and then click Search & Analysis.

For more information, see [Log Service search and analysis methods](#).

Set resource alarms

You can use the alarm function of Log Service to set resource alarms. Alarm notifications can be sent through a DingTalk group robot, a customized Webhook, and the Message Center.

**Note:**

Audit log reports provide multiple query statements. On the Logstores page, click Dashboard in the left-side navigation pane and then click a dashboard (namely, an audit log report) to display all charts. Click the menu in the upper-right corner of a chart, and then click View Details.

Example 1: Set an alarm notification for running a command on a container

To prevent Kubernetes cluster users from logging on to any container to run a command, you must set an alarm notification for running a command on any container. Furthermore, the alarm notification must include detailed information such as the container to which the user logged on, commands, user name, the event ID, the operation time, and the user IP address.

- The query statement is as follows:

```
verb : create and objectRef . subresource : exec and
stage : ResponseStarted | SELECT auditID as "event ID",
date_format ( from_unixtime ( __time__ ), '%Y-%m-%d %T' ) as "operation time",
regexp_extract ("requestURI", '([^\?]*)/exec\?.*', 1) as "resource",
regexp_extract ("requestURI", '\?(.*)', 1) as "command", "responseStatus"
as "status code",
CASE
WHEN "user . username" != 'kubernetes-admin' then "user . username"
WHEN "user . username" = 'kubernetes-admin' and
regexp_like ("annotation.s.authorization.k8s.io / reason", 'RoleBinding') then
regexp_extract ("annotation.s.authorization.k8s.io / reason", 'to User
"(\w+)"', 1)
ELSE 'kubernetes-admin' END
as "operation account",
CASE WHEN json_array_length ( sourceIPs ) = 1 then
json_format ( json_array_get ( sourceIPs , 0 )) ELSE
sourceIPs END
as "source IP address" limit 100
```

- The condition expression is `operation event =~ ".*"`.

Example 2: Set an alarm notification for failed Internet access to apiserver

To prevent malicious attacks on a Kubernetes cluster for which Internet access is enabled, you need to monitor the number of Internet access times and the failed access rate. Specifically, an alarm notification must be sent, when the number of Internet access times reaches a specified threshold and the failed access rate exceeds a specified threshold. Furthermore, the alarm notification must include detailed information such as to which the user IP address belongs, the user IP address, and the

high risk IP address. For example, to receive an alarm notification when the number of Internet access times reaches 10 and the failed access rate exceeds 50%, configure the following settings:

- Query statement.

```
* | select ip as "source IP address", total as "
number of access times", round ( rate * 100 , 2 ) as
"failed access rate %", failCount as "number of
illegal access times", CASE when security_c heck_ip
( ip ) = 1 then 'yes' else 'no' end as "high
risk IP address", ip_to_coun try ( ip ) as "country ",
ip_to_prov ince ( ip ) as "province ", ip_to_city ( ip ) as
"city ", ip_to_prov ider ( ip ) as "network operator "
from ( select CASE WHEN json_array _length ( sourceIPs )
= 1 then json_forma t ( json_array _get ( sourceIPs , 0 ))
ELSE sourceIPs END
as ip , count ( 1 ) as total ,
sum ( CASE WHEN " responseSt atus . code " < 400 then 0
ELSE 1 END ) * 1 . 0 / count ( 1 ) as rate ,
count_if ( " responseSt atus . code " = 403 ) as failCount
from log group by ip limit 10000 ) where
ip_to_doma in ( ip ) != 'intranet ' having " number of
access times " > 10 and " failed access rate %" > 50
ORDER by " number of access times " desc limit 100
```

- Condition expression is `source IP address =~ ".*"`.

Manually enable audit log report functions

You can manually enable audit log report functions.

1. Enable API server audit log.

View the API server pod settings of the three Master nodes. That is, check whether audit log settings are configured for the startup parameters, the policy file, the environment variable, and the mounting directory.

- Startup parameters

```
containers :
- command :
- kube - apiserver
- -- audit - log - maxbackup = 10
- -- audit - log - maxsize = 100
- -- audit - log - path = / var / log / kubernetes / kubernetes
. audit
- -- audit - log - maxage = 7
```

```
- -- audit - policy - file = / etc / kubernetes / audit -
policy . yml
```

- Policy file (stored in the `/ etc / kubernetes / audit - policy . yml` directory)

For more information, see [Configure a policy file](#).



Note:

If the `/ etc / kubernetes /` directory does not have any policy file, you need to run the `vi audit - policy . yml` command to create a file, and then copy the content of the policy file and paste the content to the created file.

- Environment variable

```
env :
- name : aliyun_log s_audit - c12ba20 ***** 9f0b
  value : / var / log / kubernetes / kubernetes . audit
- name : aliyun_log s_audit - c12ba20 *****
9f0b_tags
  value : audit = apiserver
- name : aliyun_log s_audit - c12ba20 *****
9f0b_produ ct
  value : k8s - audit
- name : aliyun_log s_audit - c12ba20 *****
9f0b_jsonf ile
  value : " true "
```

- Mounting directory

```
volumeMoun ts :
- mountPath : / var / log / kubernetes
  name : k8s - audit
- mountPath : / etc / kubernetes / audit - policy . yml
  name : audit - policy
  readOnly : true
volumes :
- hostPath :
  path : / var / log / kubernetes
  type : Directory0 rCreate
  name : k8s - audit
- hostPath :
  path : / etc / kubernetes / audit - policy . yml
  type : FileOrCrea te
  name : audit - policy
```

Backup the original YAML file and then restart the API server by using a new `kube - apiserver . yml` YAML file. This action will overwrite the original YAML file

stored in the `/etc/kubernetes/manifests/kube-apiserver.yaml` directory.

If the API server pod settings does not contain the preceding settings, you must upgrade the Kubernetes cluster to the latest version. For more information, see [Upgrade a Kubernetes cluster](#).

2. Use the latest version of the Log Service component.

- For how to install the Log Service component, see [Manually install the Log Service component](#).
- If you have installed the Log Service component, but the audit log function is disabled, you must upgrade the component to the latest version and you must ensure that your Logtail version is not earlier than v0.16.16 and can run on Master nodes. For more information, see [Upgrade the Log Service component](#).

3. Update audit log parsing methods.

- Log on to the [Log Service console](#).
- In the left-side navigation pane, click Project Management, and then click the name of the Project specified when creating your Kubernetes cluster.
- The Logstores page is displayed by default. Click Manage on the right of the Logstore named `audit-${clustered}`, and then click the configuration name. On the Specify Collection Mode tab page, select the JSON Mode.

Environment Variable Whitelist:

EnvKey +	EnvValue	-
aliyun_logs_audit-c3730910447504a6692ca1	/var/log/kubernetes/kubernetes.audit	✕

Collects log entries that contain the environment variables in the whitelist. If the whitelist is empty, all log entries will be collected.

Environment Variable Blacklist:

EnvKey +	EnvValue	-
----------	----------	---

Collects log entries that do not contain environment variables in the blacklist. If the blacklist is empty, all log entries will be collected.

Mode: JSON Mode ▼

[How to set JSON configuration](#)

Use System Time: ☒

Use a thirty-party log solution

Log on to the Master node of the cluster, and you can find the source file of the audit logs in the path of `/var/log/kubernetes/kubernetes.audit`. The source

file is in standard JSON format. When deploying a cluster, you can use other log solutions to collect and search audit logs, instead of using Alibaba Cloud Log Service.

1.16.3 Implement secure access through HTTPS in Kubernetes

A Container Service Kubernetes cluster supports multiple application access methods. The most common methods include `SLB : Port` access, `NodeIP : NodePort` access, and domain name access. By default, a Kubernetes cluster does not support HTTPS access. To access applications through HTTPS, you can use the secure HTTPS access method provided by Container Service and Alibaba Cloud Server Load Balancer (SLB) service. This document explains how to configure a certificate in Container Service Kubernetes by using HTTPS access configuration as an example.

Depending on different access methods, your certificate can be configured with the following two methods:

- Configure the certificate on the frontend SLB.
- Configure the certificate on Ingress.

Prerequisites

- You have created a Kubernetes cluster. For more information, see [Create a Kubernetes cluster](#).
- You have connected to the Master node through SSH. For more information, see [Access Kubernetes clusters by using SSH](#).
- After connecting to the Master node, you have created the server certificates for the cluster, including the public key certificate and the private key certificate by running the following commands :

```
$ openssl genrsa - out tls . key 2048
Generating RSA private key , 2048 bit long modulus
..... +++
+++
e is 65537 ( 0x10001 )

$ openssl req - sha256 - new - x509 - days 365 - key
tls . key - out tls . crt

You are about to be asked to enter information
that will be incorporated
...
-----
Country Name ( 2 letter code ) [ XX ]: CN
State or Province Name ( full name ) []: zhejiang
```

```

Locality Name ( eg , city ) [ Default City ]: hangzhou
Organization Name ( eg , company ) [ Default Company Ltd ]: alibaba
Organizational Unit Name ( eg , section ) []: test
Common Name ( eg , your name or your server 's hostname ) []: foo . bar . com # you must configure the domain name correctly
Email Address []: a @ alibaba . com

```

Method 1: Configure the HTTPS certificate on SLB

This method has the following advantages and disadvantages:

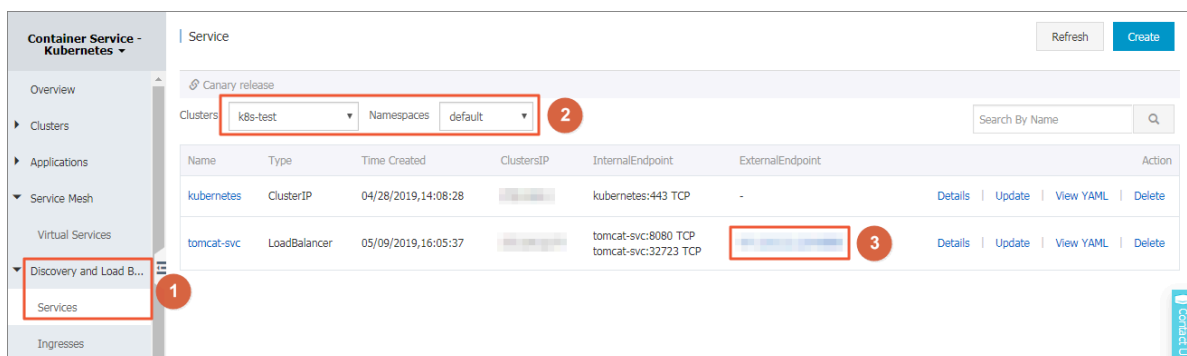
- **Advantages:** The certificate is configured on SLB and it is the external access portal of applications. The access to applications in the cluster still uses the HTTP access method.
- **Disadvantages:** You need to maintain many associations between domain names and their corresponding IP addresses.
- **Scenarios:** This method is applicable to applications that use LoadBalancer service rather than Ingress to expose access methods.

Preparations

You have created a Tomcat application in the Kubernetes cluster. The application provides external access by using the LoadBalancer service. For more information, see [Create a service](#).

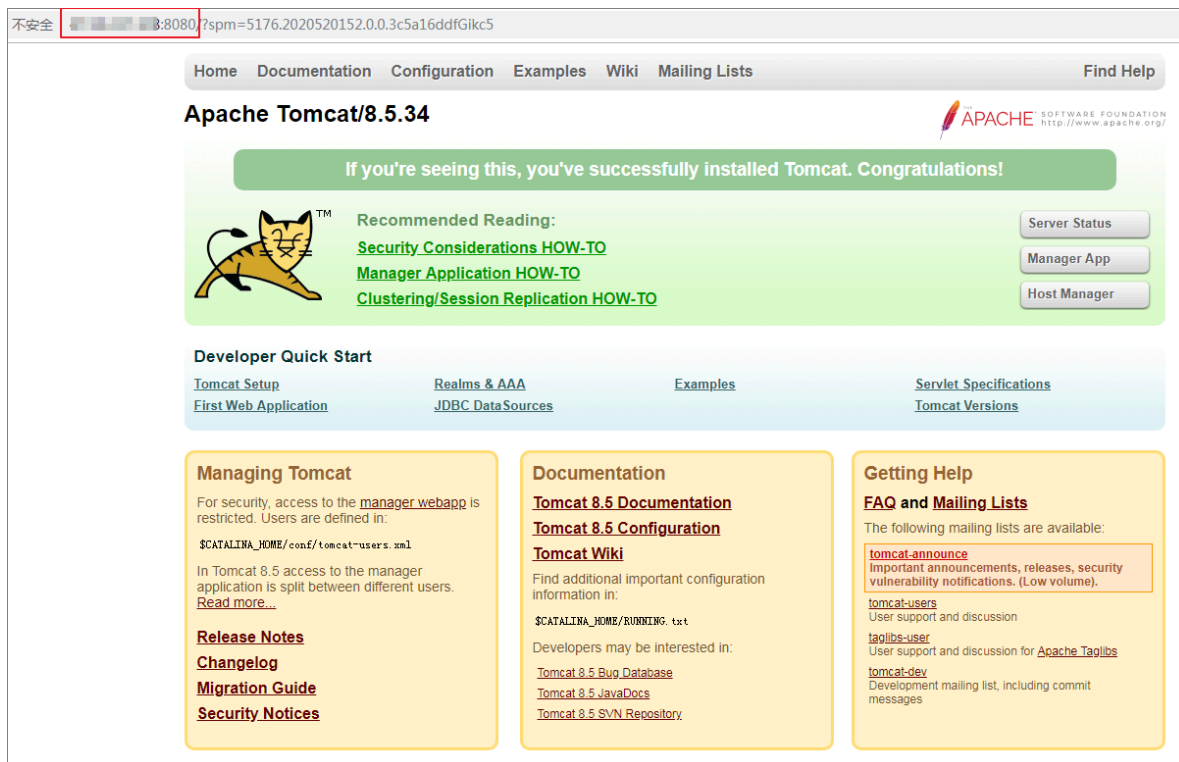
Example

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Container Service-Kubernetes, choose **Discovery and Load Balancing > Services**. Then, select the cluster and the namespace to view the pre-created Tomcat application. As shown in the following figure, the created Tomcat application is named tomcat and the service name is tomcat-svc. The service type of the application is LoadBalancer, and the service port exposed by the application is 8080.



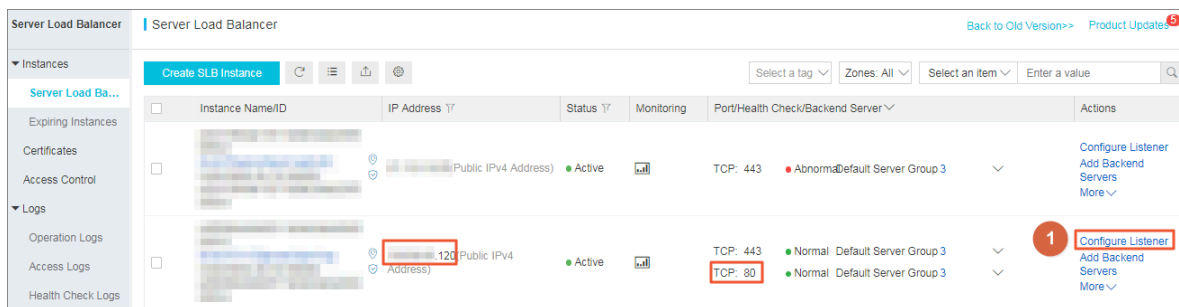
3. By clicking the external endpoint, you can access the Tomcat application through

IP : Port .



4. Log on to the [SLB console](#).

5. By default, the Server Load Balancer page is displayed. In the IP address column, find the server load balancer that corresponds to the external endpoint of the tomcat-svc service, and click Configure Listener in the actions column.



6. Configure the server load balancer. Select a listener protocol first. Select HTTPS, set the listening port to 443, and then click Next.

7. Configure the SSL certificate.

a. Click Create Server Certificate.

Configure Server Load Balancer [Back](#)

Protocol and Listener **SSL Certificates** Backend Servers Health Check Submit

Configure SSL Certificates

① Configure SSL certificates to ensure that your business is protected by encryptions and authenticated by a trusted certificate authority.

Select Server Certificate

Select [Create Server Certificate](#) [Buy Certificate](#)

Advanced [Modify](#)

Enable Mutual Authentication Disabled CA Certificate None Selected

[Previous](#) [Next](#) [Cancel](#)

- b. On the displayed page, select a certificate source. In this example, select Upload Third-Party Certificate, and then click Next.
- c. On the uploading third-party certificate page, set the certificate name and select the region in which the certificate is deployed. In the Certificate Content and

the Private Key columns, enter the server public key certificate and private key created in [Prerequisites](#), and then click OK.

Upload Third-Party Certificate

Certificate Name ?

cert-tomcat

Regions

China East 1 (Hangzhou) ×

Certificate Content ?

17 [REDACTED] zb3

18 [REDACTED] rF

19 [REDACTED] 31J

20 [REDACTED] ned

21 [REDACTED] 0vz

22 [REDACTED]

23 [REDACTED]

(NGINX-compatible) Upload [View Sample Certificate](#)

Private Key: ?

22 [REDACTED] J1

23 [REDACTED] 6K

24 [REDACTED] 0y

25 [REDACTED] Hs

26 [REDACTED]

27 -----END RSA PRIVATE KEY-----

28 [REDACTED]

(NGINX-compatible) Upload [View Sample Certificate](#)

Previous

OK

Cancel

468

Issue: 20190614

- d. From the Select Server Certificate drop-down list, select the created server certificate.
 - e. Click Next.
8. Configure Backend Servers. By default, servers are added. You need to configure a port for each backend server to listen to the tomcat-svc service, and then click Next.



Note:

You need to find the NodePort number of this service in the Container Service Web interface, and configure the number as the port number of each backend server.

Configure Server Load Balancer Back

Protocol and Listener SSL Certificates **Backend Servers** Health Check Submit

Add Backend Servers

① Add backend servers to handle the access requests received by the SLB instance.

Forward Requests To

Default Server Group VServer Group Active/Standby Server Group

Servers Added

ECS Instance ID/Name	Public/Internal IP Address	Port	Weight	Actions
node-0003-k8s-for-cs-cad7e15c0784848a5be02443e9186ccb7-i-bp1a4u8pao171d36zfg1	192.168.0.78(Private) vpc-bp1k1kyevdjjerqs0u4vb vsw-bp1qb1yn2nbzm4kck66xx	32529	100	Delete
node-0001-k8s-for-cs-cad7e15c0784848a5be02443e9186ccb7-i-bp1hk1m08e5rxkgrae8a	192.168.0.37(Private) vpc-bp1k1kyevdjjerqs0u4vb vsw-bp1qb1yn2nbzm4kck66xx	32529	100	Delete
node-0002-k8s-for-cs-cad7e15c0784848a5be02443e9186ccb7-i-bp1hk1m08e5rxkgrae8b	192.168.0.38(Private) vpc-bp1k1kyevdjjerqs0u4vb vsw-bp1qb1yn2nbzm4kck66xx	32529	100	Delete

4 servers have been added. 0 servers are to be added, and 1 servers are to be deleted. Add More

Previous **Next** Cancel

9. Configure Health Check, and then click Next. In this example, use the default settings.
10. Confirm the Submit tab. When you make sure that all configurations are correct, click Submit.

11. After completing the configuration, click OK.

The screenshot shows the 'Configure Server Load Balancer' dialog box. It has a progress bar at the top with four steps: 'Protocol and Listener', 'Backend Servers', 'Health Check', and 'Submit'. The 'Submit' step is currently active. Below the progress bar, there are two status indicators: 'Layer-7 listener' and 'Start Listener', both showing 'Success' in green. At the bottom, there are two buttons: 'OK' (highlighted with a red rectangle) and 'Cancel'.

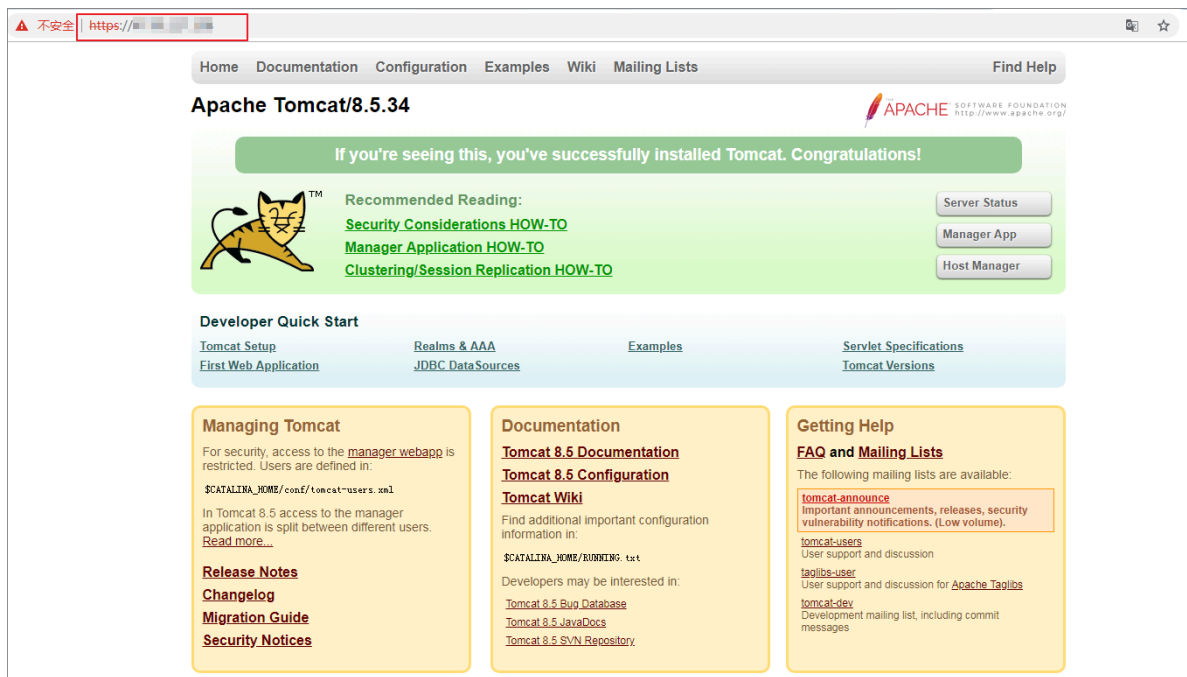
12. Return to the Server Load Balancer page to view the instance. The listening rule of `HTTPS : 443` is generated.

13. Access the Tomcat application through HTTPS. In the address bar of the browser, enter `https :// slb_ip` to access the application.



Note:

If the domain name authentication is included in the certificate, you can access the application by using the domain name. You can also access the application through `slb_ip : 8080` because `tcp : 8080` is not deleted.



Method 2: Configure the certificate on Ingress

This method has the following advantages and disadvantages:

- **Advantages:** You do not need to modify the SLB configuration. All applications can manage their own certificates through Ingress without interfering with each other.
- **Disadvantages:** Each application can be accessed by using a separate certificate or the cluster has applications that can be accessed by only using a certificate.

Preparations

You have created a Tomcat application in the Kubernetes cluster. The service of the application provides access through ClusterIP. In this example, use Ingress to provide the HTTPS access service.

Example

1. Log on to the Master node of the Kubernetes cluster and create a secret according to the prepared certificate.



Note:

You must set the domain name properly. Otherwise, you will encounter exceptions when accessing the application through HTTPS.

```
kubectl create secret tls secret - https -- key tls .  
key -- cert tls . crt
```

2. Log on to the [Container Service console](#).
3. In the left-side navigation pane under Container Service-Kubernetes, choose **Discovery and Load Balancing > Ingresses**, select a cluster and namespace, and click **Create** in the upper-right corner.

4. In the displayed dialog box, configure the Ingress to make it accessible through HTTPS, and then click OK.

For more information about Ingress configuration, see [Create an Ingress in the Container Service console](#). The configuration in this example is as follows:

- **Name:** Enter an Ingress name.
- **Domain:** Enter the domain name set in the preceding steps. It must be the same as that configured in the SSL certificate.
- **Service:** Select the service corresponding to the tomcat application. The service port is 8080.
- **Enable TLS:** After enabling TLS, select the existing secret.

Create

Name:

tomcat-https

Rule:

+

 Add

Domain

foo.bar.com

Select *.cd5f29d03dcd544d3943a4c2cb45bb4ec.cn-hangzhou.alicontainer.com or Custom

path

e.g./

Service

+

 Add

Name	Port	Weight	Percent of Weight
tomcat-svc	8080	100	100.0%

☒ Enable TLS
 ☒ Exist secret
 ☐ Create secret

secret-https

Service weight:

☒ Enable

Grayscale release:

+

 Add

After the gray rule is set, the request meeting the rule will be routed to the new service. If you set a weight other than 100, the request to satisfy the gamma rule will continue to be routed to the new and old version services according to the weights.

annotation:

+

 Add

rewrite annotation

Tag:

+

 Add

Create

Cancel

You can also use a YAML file to create an Ingress. In this example, the YAML sample file is as follows:

```
apiVersion : extensions / v1beta1
kind : Ingress
metadata :
  name : tomcat - https
spec :
  tls :
  - hosts :
    - foo . bar . com
    secretName : secret - https
  rules :
  - host : foo . bar . com
    http :
      paths :
      - path : /
        backend :
          serviceName : tomcat - svc
          servicePort : 8080
```


- Return to the Ingress list to view the created Ingress, the endpoint, and the domain name. In this example, the domain name is `foo . bar . com` . You can also enter the Ingress detail page to view the Ingress.



Note:

In this example, `foo . bar . com` is used as a testing domain name, and you need to create a record in the hosts file.

```
47 . 110 . 119 . 203   foo . bar . com           # where
, the IP address is the Ingress endpoint .
```

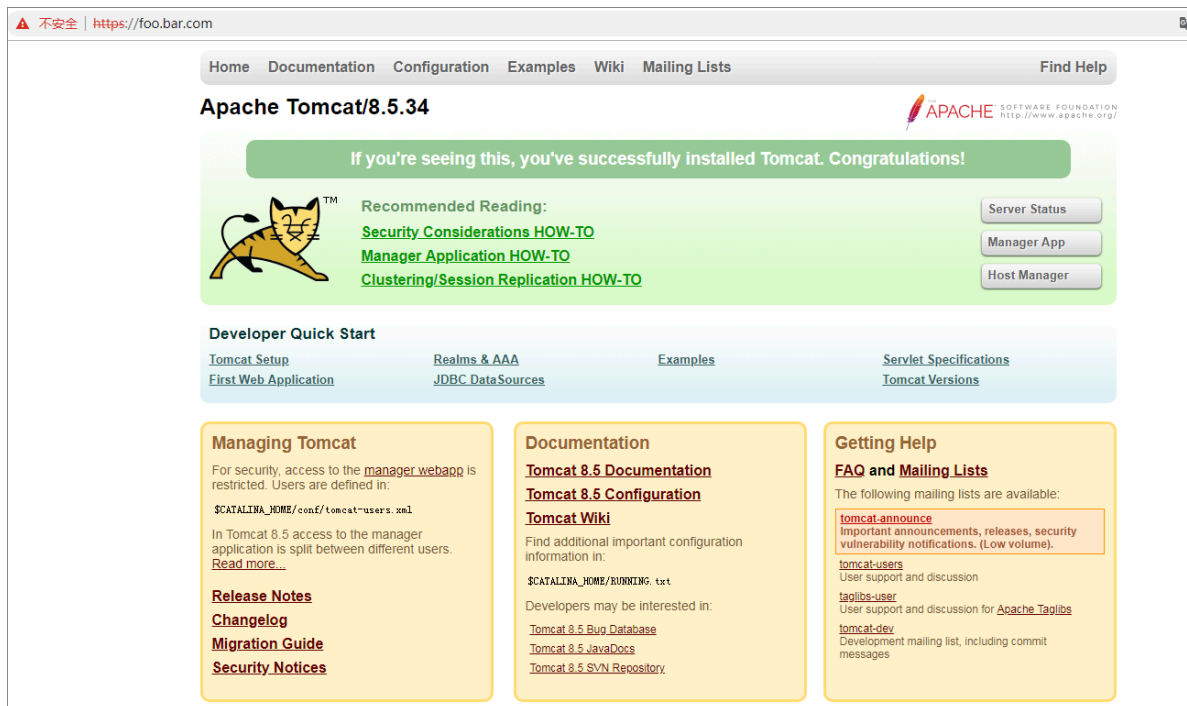
Ingress						Refresh	Create
Help: Blue-green release							
Clusters	k8s-test111	Namespace	default				
Name	Endpoint	Rule	Time Created	Action			
tomcat-https		foo.bar.com/ -> tomcat-svc	11/07/2018,15:37:00	Details	Update	View YAML	Delete

6. In the browser, access `https://foo.bar.com`.



Note:

You need to access the domain name by using HTTPS because you have created a TLS access certificate. This example uses `foo.bar.com` as a sample domain name to be parsed locally. In your specific configuration scenarios, you need to use the registered domain names.



1.17 Release management

1.17.1 Manage a Helm-based release

Alibaba Cloud Container Service for Kubernetes is integrated with the package management tool Helm to help you quickly deploy applications on the cloud. However, Helm charts can be released multiple times and the release version must be managed. Container Service for Kubernetes provides a release function, which allows you to manage the applications released by using Helm in the Container Service console.

Prerequisites

- You have created a Kubernetes cluster. For more information, see [Create a Kubernetes cluster](#).

- You have installed a Helm application by using the App Catalog function or Service Catalog function. For more information, see [Simplify Kubernetes application deployment by using Helm](#). In this topic, the wordpress-default application is used as an example.

View release details

- Log on to the [Container Service console](#).
- In the left-side navigation pane, select Container Service - Kubernetes. Then, select Application > Release and click the Helm tab. Select the target cluster from the Clusters drop-down list.

In the displayed release list, you can view the applications and services released through Helm in the selected cluster.

Release Name	Status	Namespace	Chart Name	Chart Version	App Version	Update Time	Action
wordpress-default	Deployed	default	wordpress	0.6.13	4.8.2	11/23/2018,11:43:51	Details Update Delete
wordpress-test	Deployed	default	wordpress	4.0.0	4.9.8	11/23/2018,10:49:40	Details Update Delete
istio	Deployed	istio-system	ack-istio	1.0.2	1.0.2	10/26/2018,18:18:28	Details Update Delete
alibaba-log-controller	Deployed	default	alibaba-cloud-log	0.1.1	1.0	10/22/2018,15:20:34	Details Update Delete

- Find your target release (wordpress-default in this example) and click Details to view the release details.

You can view such release details as the current version and history version. In this example, the current version is 1 and no history version exists. On the Resource tab page, you can view the resource information of wordpress-default, such as the resource name and the resource type, and view the YAML information.



Note:

You can view the running status of the resource in details by clicking the resource name and going to the Kubernetes dashboard page.

Container Service | Release List - wordpress-default | Refresh

Kubernetes | Swarm

Clusters

Nodes

Storage

Application

Deployment

Service

Release

Config Maps

Store

App Catalog

Service Catalog

Current Version

Release Name : wordpress-default | Namespace : default | Deployed at : 04/20/2018,15:27:55

Current Version : 1 | Time Updated : 04/20/2018,15:27:55

Resource	Kind	Values
wordpress-default-mariadb	Secret	View YAML
wordpress-default-wordpress	Secret	View YAML
wordpress-default-mariadb	ConfigMap	View YAML
wordpress-default-mariadb	PersistentVolumeClaim	View YAML
wordpress-default-wordpress	PersistentVolumeClaim	View YAML
wordpress-default-mariadb	Service	View YAML
wordpress-default-wordpress	Service	View YAML
wordpress-default-mariadb	Deployment	View YAML
wordpress-default-wordpress	Deployment	View YAML

History Version

4. Click the Values tab to view the release parameters.

Container Service | Release List - wordpress-default | Refresh

Kubernetes | Swarm

Clusters

Nodes

Storage

Application

Deployment

Service

Release

Config Maps

Store

App Catalog

Service Catalog

Current Version

Release Name : wordpress-default | Namespace : default | Deployed at : 04/20/2018,15:27:55

Current Version : 1 | Time Updated : 04/20/2018,15:27:55

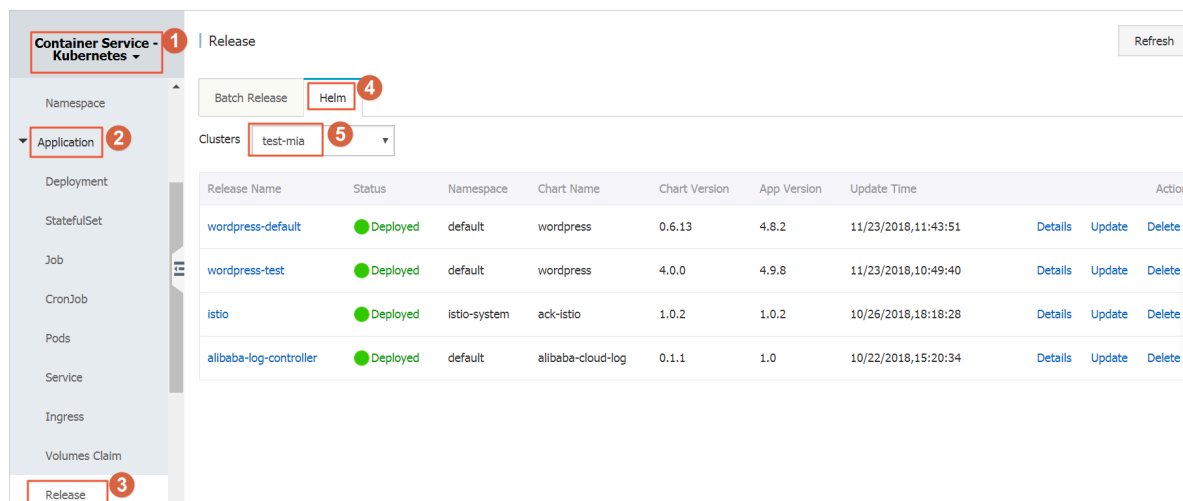
Resource	Values
1	## Bitnami WordPress image version
2	## ref: https://hub.docker.com/r/bitnami/wordpress/tags/
3	##
4	image: bitnami/wordpress:4.8.2-r0
5	##
6	## Specify a imagePullPolicy
7	## ref: http://kubernetes.io/docs/user-guide/images/#pre-pulling-images
8	##
9	imagePullPolicy: IfNotPresent
10	##
11	## User of the application
12	## ref: https://github.com/bitnami/bitnami-docker-wordpress#environment-variables
13	##
14	wordpressUsername: user
15	##
16	## Application password
17	## Defaults to a random 10-character alphanumeric string if not set
18	## ref: https://github.com/bitnami/bitnami-docker-wordpress#environment-variables
19	##
20	wordpressPassword:
21	##
22	## Admin email
23	## ref: https://github.com/bitnami/bitnami-docker-wordpress#environment-variables
24	##
25	wordpressEmail: user@example.com

Update a release version

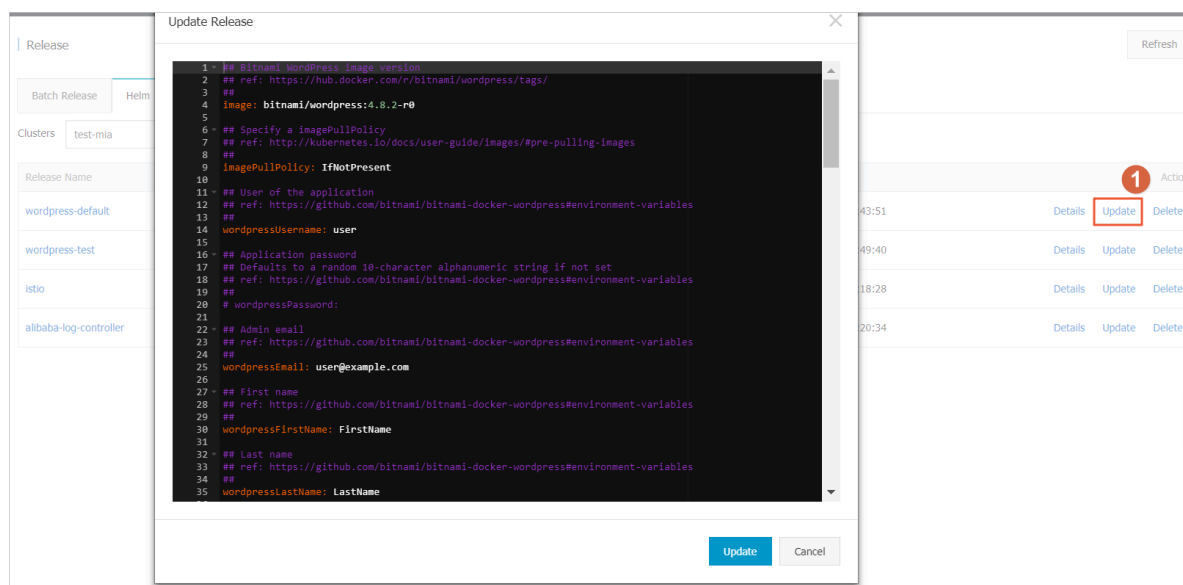
1. Log on to the [Container Service console](#).

2. In the left-side navigation pane, select Container Service - Kubernetes. Then, select Application > Release and click the Helm tab. Select the target cluster from the Clusters drop-down list.

In the displayed release list, you can view the applications and services released through Helm in the selected cluster.



3. Find your target release (wordpress-default in this example). Click Update and the Update Release dialog box appears.



4. Modify the parameters and then click Update.

Update Release

```

124     ##
125     # tls:
126     #   - secretName: wordpress.local-tls
127     #   hosts:
128     #     - wordpress.local
129
130   ## Enable persistence using Persistent Volume Claims
131   ## ref: http://kubernetes.io/docs/user-guide/persistent-volumes/
132   ##
133   persistence:
134     enabled: true
135     ## wordpress data Persistent Volume Storage Class
136     ## If defined, storageClassName: <storageClass>
137     ## If set to "-", storageClassName: "", which disables dynamic provisioning
138     ## If undefined (the default) or set to null, no storageClassName spec is
139     ## set, choosing the default provisioner. (gp2 on AWS, standard on
140     ## GKE, AWS & OpenStack)
141     ##
142     storageClass: "alicloud-disk-efficiency"
143     accessMode: ReadWriteOnce
144     size: 20Gi
145
146   ## Configure resource requests and limits
147   ## ref: http://kubernetes.io/docs/user-guide/compute-resources/
148   ##
149   resources:
150     requests:
151       memory: 512Mi
152       cpu: 300m
153
154   ## Node labels for pod assignment
155   ## Ref: https://kubernetes.io/docs/user-guide/node-selection/
156   ##
157   nodeSelector: {}
158

```

Update

Cancel

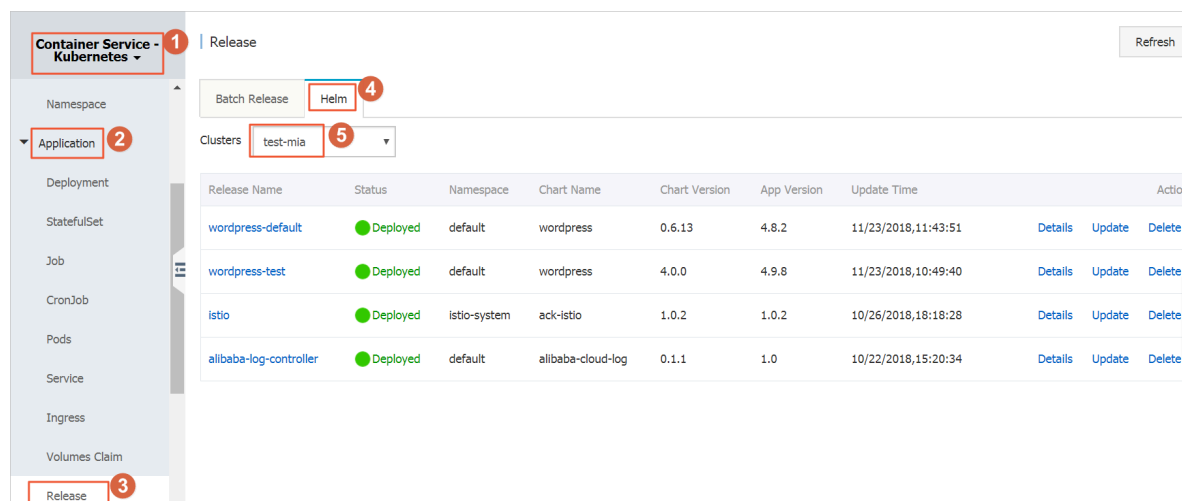
On the release list page, you can see that the current version changes to 2. To roll back to version 1, click Details and in the History Version area, click Rollback.

Current Version		
Release Name : wordpress-default	Namespace : default	Deployed at : 04/20/2018,17:45:35
Current Version : 2		Time Updated : 04/20/2018,17:45:46
Resource	Kind	Values
wordpress-default-mariadb	Secret	View YAML
wordpress-default-wordpress	Secret	View YAML
wordpress-default-mariadb	ConfigMap	View YAML
wordpress-default-mariadb	PersistentVolumeClaim	View YAML
wordpress-default-wordpress	PersistentVolumeClaim	View YAML
wordpress-default-mariadb	Service	View YAML
wordpress-default-wordpress	Service	View YAML
wordpress-default-mariadb	Deployment	View YAML
wordpress-default-wordpress	Deployment	View YAML
History Version		
Version : 1 Rollback		Time Updated : 04/20/2018,17:45:35

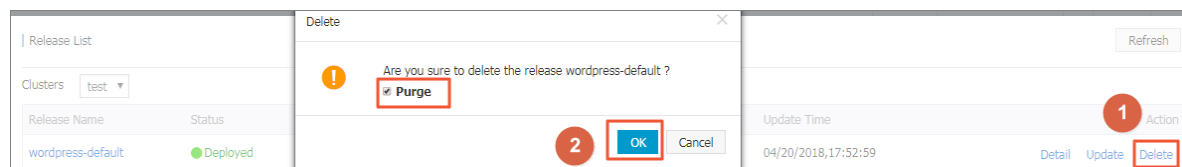
Delete a release

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane, select Container Service - Kubernetes. Then, select Application > Release and click the Helm tab. Select the target cluster from the Clusters drop-down list.

In the displayed release list, you can view the applications and services released through Helm in the selected cluster.



3. Find your target release (wordpress-default in this example). Click Delete and the Delete dialog box appears.



4. Select the Purge check box if you want to clear the release records, and then click OK. After you delete a release, the related resources such as the services and deployments are deleted too.

1.17.2 Use batch release on Alibaba Cloud Container Service for Kubernetes

You can use Alibaba Cloud Container Service for Kubernetes to release application versions in batches, achieving fast version verification and rapid iteration of applications.

Context



Note:

The latest Kubernetes cluster has installed alicloud-application-controller by default. For older versions of clusters, only versions of 1.9.3 and later are currently supported, and you can upgrade old versions of clusters through the prompt link on the console.

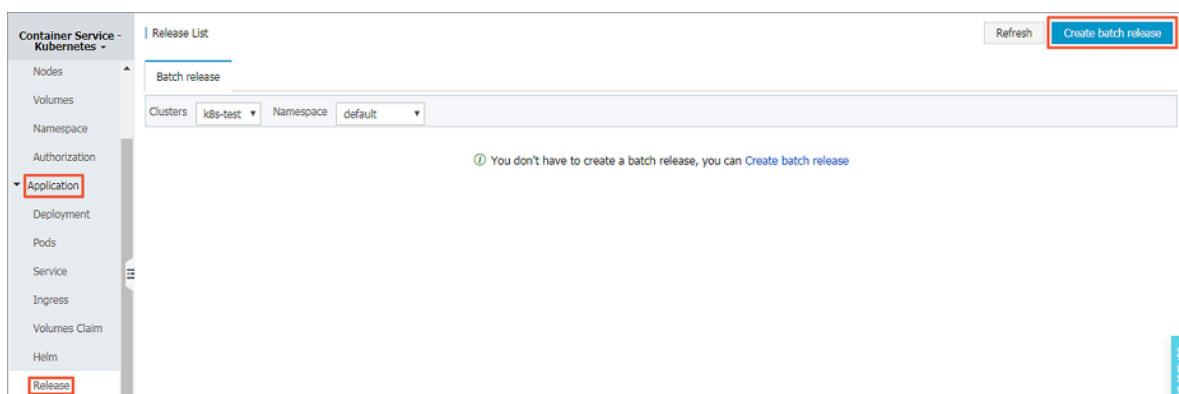
Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Application > Release in the left-side navigation pane. Click Create batch release in the upper-right corner.

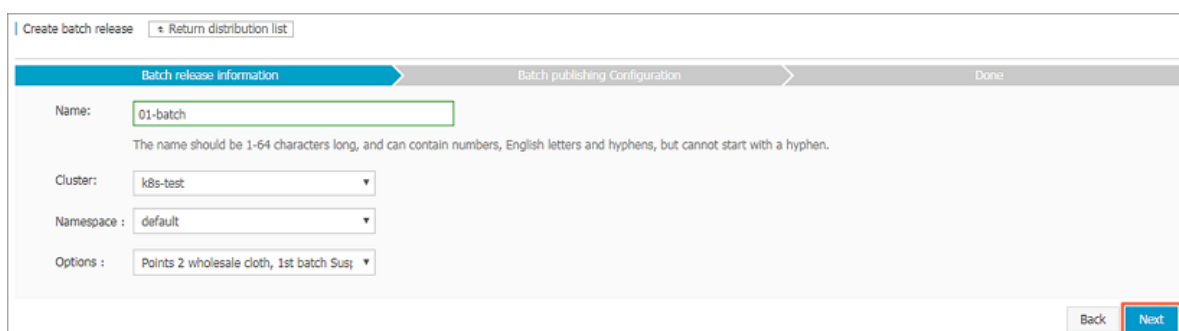


Note:

If the button is gray, you can upgrade the cluster by following the upgrade link.



3. Configure batch release information, including the application name, cluster, namespace, and options. Click Next.



4. On the batch publishing configuration page, configure the backend pod and service, and then click Update to create an application.

The screenshot shows the 'Batch publishing Configuration' page. The 'General' section includes fields for 'Image Name' (nginx), 'Image Version' (latest), 'Scale' (4), 'Resource Limit' (CPU: 500m, Memory: 128Mi), and 'Resource Request' (CPU: 500m, Memory: 128Mi). The 'Access Control' section includes 'Service' (Server Load Balancer) and 'Port Mapping' (Port: 80, target port: 80, Protocol: TCP). The 'Update' button is highlighted with a red circle and the number 3. Other red circles with numbers 1 and 2 highlight the 'Image Name' field and the 'Port Mapping' table respectively.

5. Return to the release list, an application is displayed in the Not started status. Click Detail on the right.

The screenshot shows the 'Release List' page. It displays a table with the following data:

Release Name	Namespace	Update Time	Status	Action
01-batch	default	2018-09-03 16:00:56	Not started	Detail Update Delete

6. On the application detail page, you can view more information. Click Change Configuration in the upper-right corner of the page to make a batch release change.

The screenshot shows the 'Release List - nginx' page. It displays the 'Details' tab. The 'Overview' section shows the following information:

- Release Name: nginx
- Release Type: Batch Release
- Created At: 2018-09-28 16:24:04
- RelatedService: [batchrelease-nginx-svc](#)
- Status: Not Started

Below the overview is a table showing the status of individual pods:

Name	App Version	Status	Pod IP	Time Created	Action
batchrelease-nginx-0	v1	Running	172.16.1.159	09/28/2018,16:24:04	Terminal Logs
batchrelease-nginx-1	v1	Running	172.16.1.160	09/28/2018,16:24:07	Terminal Logs
batchrelease-nginx-2	v1	Running	172.16.1.161	09/28/2018,16:24:10	Terminal Logs
batchrelease-nginx-3	v1	Running	172.16.1.162	09/28/2018,16:24:12	Terminal Logs

7. Configure changes for the new version of the application, and then click Update.

8. By default, you return to the release list page, where you can view the batch release status of the application. After completing the first deployment, click Detail.

Release List

Refresh

Create Batch Release

Batch Release

Clusters

k8s-test

Namespace

default

Release Name

Namespace

Update Time

Status

1

Action

nginx

default

2018-09-28 16:24:04

Wait for Deployment to Complete (Total Batches: 2. Currently, batch 0 is in process and the batch status is Deploying)

Details

Update

Delete

9. You can see that the Not Started list is has two pods and the Completed list has two pods, which indicates that the first batch has been completed in batch release. Click Continue, you can release the second batch of pods. Click Roll Back to roll back to the previous version.

Details

History

Overview

Release Name:

nginx

Release Type:

Batch Release

Created At:

2018-09-28 16:24:04

RelatedService:

[batchrelease-nginx-svc](#)

Status:

To be confirmed (Total Batches: 2. Currently, batch 1 is in process and the batch status is Completed)

Not Started

In Progress

Completed

Refresh

Continue

Roll Back

Complete

Name	App Version	Status	Pod IP	Time Created	Action
batchrelease-nginx-2	v2	<div><div></div>Running</div>	172.16.1.164	09/28/2018,16:26:57	Terminal Logs
batchrelease-nginx-3	v2	<div><div></div>Running</div>	172.16.1.163	09/28/2018,16:26:53	Terminal Logs

10. When completing the release, click **History** to roll back to history versions.



What's next

You can use batch release to quickly verify your application version without traffic consumption. Batch release is more resource-saving than blue-green release. Currently, batch release can be performed on only web pages. The yaml file editing is to be opened later to support more complex operations.

1.18 Istio management

1.18.1 Overview

Istio is an open platform that provides connection, protection, control and monitors microservices.

Microservices are currently being valued by more and more IT enterprises.

Microservices are multiple services divided from a complicated application. Each service can be developed, deployed, and scaled. Combining the microservices and container technology simplifies the delivery of microservices and improves the liability and scalability of applications.

As microservices are extensively used, the distributed application architecture composed of microservices becomes more complicated in dimensions of operation and maintenance, debugging, and security management. Developers have to deal with greater challenges, such as service discovery, load balancing, failure recovery, metric collection and monitoring, A/B testing, gray release, blue-green release, traffic limiting, access control, and end-to-end authentication.

Istio emerged. Istio is an open platform for connecting, protecting, controlling, and monitoring microservices. It provides a simple way to create microservices networks and provides capabilities such as load balancing, inter-service authentication, and monitoring. Besides, Istio can provide the preceding functions without modifying services.

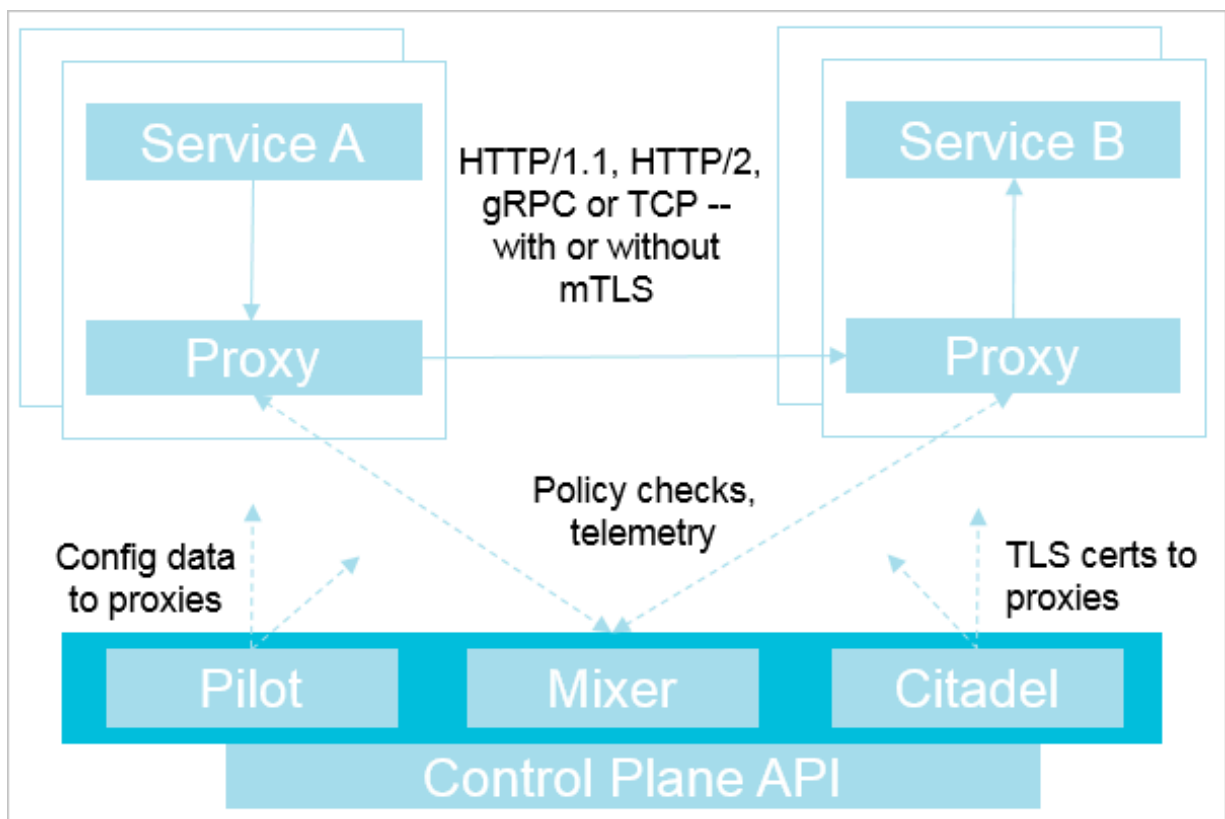
Istio provides the following functions:

- **Traffic management:** Controls traffic and API calls between services to enhance the system reliability.
- **Authentication and security protection:** Provides authentication for services in meshes, and protects the traffic of services to enhance the system security.
- **Policy execution:** Controls access policies between services without requiring changes to the services.
- **Observability:** Obtains traffic distribution and call relationships between services to quickly locate problems.

Istio architecture

Istio is logically divided into a control plane and a data plane:

- **Control plane:** Administration proxy (the default is Envoy) for managing traffic routing, runtime policy execution, and more
- **Data plane:** Consists of a series of proxys (the default is Envoy) for managing and controlling network communication between services.



Istio is composed of the following components:

- **Istio Pilot:** Collects and validates configurations, and propagates them to various Istio components. It extracts environment-specific implementation details from the policy execution module (Mixer) and the intelligent proxy (Envoy), providing them with an abstract representation of user services, independent of the underlying platform. In addition, traffic management rules (that is, generic Layer-4 rules and Layer-7 HTTP/gRPC routing rules) can be programmed through Pilot at runtime.
- **Policy execution module (Mixer):** Executes access control and usage policies across the service mesh, and collects telemetry data from the intelligent proxy (Envoy) and other services. Mixer executes policies based on the Attributes provided by the intelligent proxy (Envoy).
- **Istio security module:** Provides inter-service and inter-user authentication to guarantee enhanced security between services without modifying service codes. Includes three components:
 - **Identification:** When Istio runs on Kubernetes, it identifies the principal that runs the service according to the service account provided by container Kubernetes.
 - **Key management:** Provides CA automated generation, and manages keys and certificates.
 - **Communication security:** Provides a tunnel between the client and the server through the intelligent proxy (Envoy) to secure services.
- **Intelligent proxy (Envoy):** Deployed as an independent component in the same Kubernetes pod along with relevant microservice, and provides a series of attributes to the policy execution module (Mixer). The policy execution module (Mixer) uses these attributes as the basis to execute policies, and sends them to monitoring systems.

1.18.2 Deploy Istio

The distributed application architecture composed of microservices has disadvantages in aspects such as operation and maintenance (O&M), debugging, and security management. To eliminate the disadvantages, you can deploy Istio to create microservice network and to provide load balancing, service-to-service authentication, monitoring, and other functions. Istio provides the functions without requiring any changes to services.

Prerequisites

- A Kubernetes cluster is created. For more information, see [Create a Kubernetes cluster](#).
- You have logged on to the Container Service console by using the primary account or by using a sub-account that has been granted sufficient permissions. For example, if the `cluster - admin` permission is granted to a sub-account then Istio can be deployed. Other combinations of permissions are also sufficient. For more information, see [Grant Kubernetes permissions to a RAM user](#).

Background information

- Alibaba Cloud Container Service for Kubernetes in versions of 1.10.4 and later support Istio deployment. If your Container Service for Kubernetes is in any version prior to 1.10.4, update the version to 1.10.4 or later.
- To guarantee sufficient resources, the number of Worker nodes in a cluster must be greater than or equal to 3.



Procedure

You can deploy Istio through the Clusters page or through the App Catalog page.

- Deploy Istio through the Clusters page

1. Deploy Istio.



- Log on to the [Container Service console](#).
- In the left-side navigation pane, choose Clusters > Clusters.
- On the right of the target cluster, choose More > Deploy Istio.


Cluster Name/ID	Cluster Type	Region (All)	Network Type	Cluster Status	Number of Nodes	Time Created	Kubernetes Version	Action
	Kubernetes	China North 2 (Beijing)	VPC vpc-2ze6clfm15d...	Running	6	10/22/2018,19:41:50	1.11.2	Manage View Logs Dashboard Scale Cluster More
	Kubernetes	China North 2 (Beijing)	VPC vpc-2ze8sk2raka...	Failed to delete	4	09/18/2018,14:00:45	1.11.2	Delete Add Existing Instance Upgrade Cluster Automatic Scaling Addon Upgrade Deploy Istio

- Set the following Istio parameters.

Configuration	Description
Clusters	Target cluster in which Istio is deployed.
Namespace	Namespace in which Istio is deployed.

Configuration	Description
Release Name	Name of Istio to be released.
Enable Prometheus for metrics/ logs collection	Whether to enable Prometheus for metrics/logs collection. Enabled by default.
Enable Grafana for metrics display	Whether to enable Grafana for metrics display. Enabled by default.
Enable automatic Istio Sidecar injection	Whether to enable automatic Istio Sidecar injection. Enabled by default.
Enable the Kiali Visualization Service Mesh	Whether to enable the Kiali Visualization Service Mesh. Disabled by default. <ul style="list-style-type: none">- Username: Set a user name. The default is admin.- Password: Set a password. The default is admin.

Configuration	Description
Tracing Analysis Settings	<ul style="list-style-type: none"> - Enable Distributed Tracing with Jaeger : indicates whether to enable Jaeger (the distributed tracing system). To use Jaeger, select this radio button, and activate Alibaba Cloud Log Service. <div>  <p>Note: If you select this radio button, Log Service will automatically create a project named <code>istio - tracing - { ClusterID }</code> that is used to store the tracking data.</p> </div> <ul style="list-style-type: none"> - Activate Tracing Analysis: indicates whether to activate the Tracing Analysis service. To activate this service, select this radio button, and then click Activate now. Additionally, you need to enter an endpoint address in the format of <code>http :// tracing - analysis - dc - hz . aliyuncs . com /.../ api / v1 / spans .</code> <div>  <p>Note:</p> <ul style="list-style-type: none"> ■ An address of this format indicates an Internet or intranet endpoint used by a Zipkin client to transmit collected data in a region to the Tracing Analysis service that uses the API from v1 release. ■ If you use an intranet endpoint, you must ensure that your Kubernetes cluster and the Tracing Analysis instance are in the same region to maintain stable network performance. </div>
Pilot Settings	Set the trace sampling percentage in the range of 0 to 100. The default value is 1.

Configuration	Description
Control Egress Traffic	<ul style="list-style-type: none"> - Permitted Addresses for External Access: range of IP addresses that can be used to directly access services in the Istio service mesh. By default, this field is left blank. Use commas (,) to separate multiple IP address ranges. - Blocked Addresses for External Access: range of IP addresses that are blocked against external accesses. By default, this IP address range contains the cluster pod CIDR block and service CIDR block. Use commas (,) to separate multiple IP address ranges. - ALL: Select this check box to block all the IP addresses used to access the Internet. <div>  Note: <p>If the settings of these two parameters conflict with each other, the Permitted Addresses for External Access prevails.</p> <p>For example, if an IP address is listed in both IP address ranges that you set for these two parameters, the IP address can be still accessed. That is, the setting of Permitted Addresses for External Access prevails.</p> </div>

e. Click Deploy Istio.

At the bottom of the deployment page, you can view the deployment progress and status in real time.

Step	Status
Create Istio Resource Definition	● Succeeded 53 / 53
Deploy Istio	⚙️ Running

Deploy Istio

Verify the result

You can view your deployment results in the following ways:

- At the bottom of the Deploy Istio page, Deploy Istio is changed to Deployed.

Step	Status
Create Istio Resource Definition	pending
Deploy Istio	pending

Deployed

- ■ In the left-side navigation pane, choose **Application > Pods**.
- Select the cluster and namespace in which Istio is deployed, and you can see the relevant pods in which Istio is deployed.

Container Service - Kubernetes

Overview

Clusters

Nodes

Volumes

Namespace

Authorization

Application

Deployment

StatefulSet

Pods

Service

Pods

Clusters

k8s

Namespace

istio-system

Name	Status	Pod IP	Node	Time Created	CPU	Memory	
grafana-6d786489b9-87288	Running			10/22/2018,20:14:53	<div></div> 0.006	<div></div> 24.57 Mi	<div>Details</div> <div>More</div>
istio-citadel-765964b585-nc4lg	Running			10/22/2018,20:14:53	<div></div> 0	<div></div> 10.598 Mi	<div>Details</div> <div>More</div>
istio-egressgateway-5ddcf6d6f4-6p5kw	Running			10/22/2018,20:14:53	<div></div> 0.002	<div></div> 27.918 Mi	<div>Details</div> <div>More</div>
istio-egressgateway-5ddcf6d6f4-wtwt4t	Running			10/22/2018,20:18:38	<div></div> 0.002	<div></div> 26.758 Mi	<div>Details</div> <div>More</div>
istio-galley-85666b6578-qlnfq	Running			10/22/2018,20:14:53	<div></div> 0.017	<div></div> 11.895 Mi	<div>Details</div> <div>More</div>
istio-ingressgateway-85669f5c5-g48q2	Running			10/22/2018,20:14:53	<div></div> 0.002	<div></div> 25.086 Mi	<div>Details</div> <div>More</div>

- ■ In the left-side navigation pane, choose **Application > Service**.
- Select the cluster and namespace in which Istio is deployed, and you can see the access addresses provided by the relevant services in which Istio is deployed.

Container Service - Kubernetes

Overview

Clusters

Nodes

Volumes

Namespace

Authorization

Application

Deployment

StatefulSet

Pods

Service

Ingress

Volumes Claim

Helm

Release

Service List


Help: [Canary release](#)

Clusters

k8s

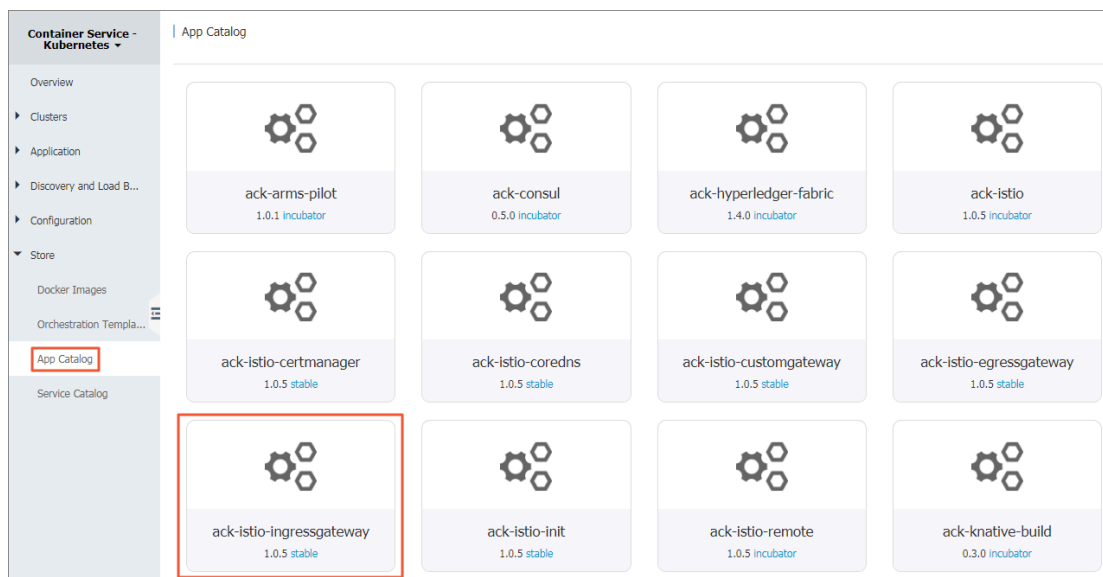
 Namespace

istio-system

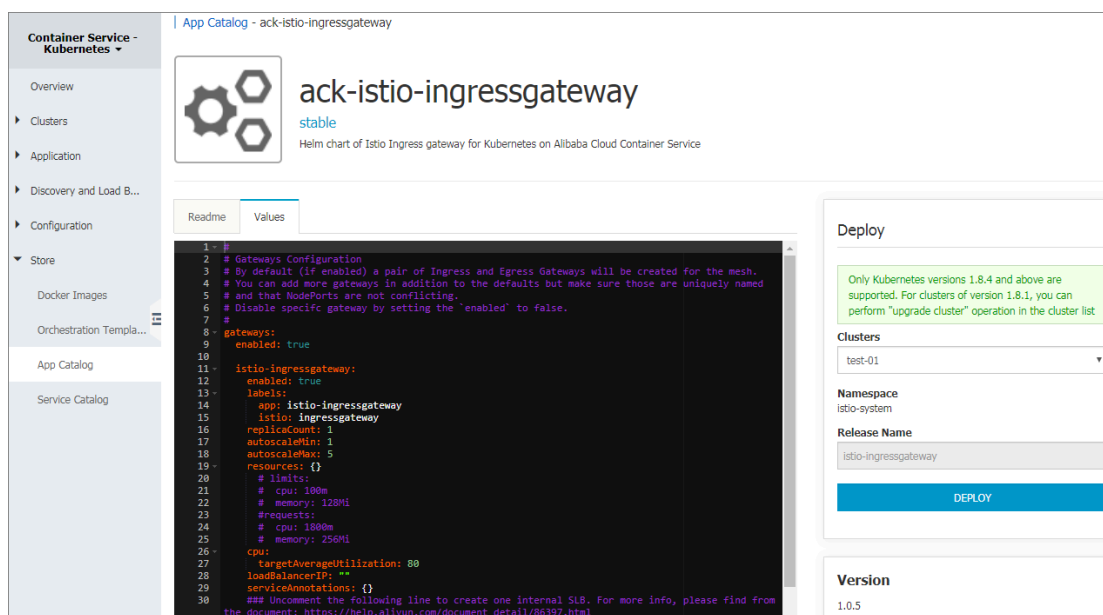
Name	Type	Time Created	ClustersIP	InternalEndpoint	ExternalEndpoint	Action
grafana	ClusterIP	10/22/2018,20:14:52		grafana:3000 TCP	-	Details Update View YAML Delete
istio-citadel	ClusterIP	10/22/2018,20:14:53		istio-citadel:8060 TCP istio-citadel:9093 TCP	-	Details Update View YAML Delete
istio-egressgateway	ClusterIP	10/22/2018,20:14:52		istio-egressgateway:80 TCP istio-egressgateway:443 TCP	-	Details Update View YAML Delete
istio-galley	ClusterIP	10/22/2018,20:14:52		istio-galley:443 TCP istio-galley:9093 TCP	-	Details Update View YAML Delete
istio-ingressgateway	LoadBalancer	10/22/2018,20:14:52		istio-ingressgateway:80 TCP istio-ingressgateway:31380 TCP istio-ingressgateway:443 TCP istio-ingressgateway:31390 TCP istio-ingressgateway:31400 TCP istio-ingressgateway:31400 TCP istio-ingressgateway:15011 TCP istio-ingressgateway:32009 TCP istio-ingressgateway:8060 TCP istio-ingressgateway:32522 TCP istio-ingressgateway:853 TCP istio-ingressgateway:32512 TCP istio-ingressgateway:15030 TCP istio-ingressgateway:30282 TCP istio-ingressgateway:15031 TCP istio-ingressgateway:31006 TCP		Details Update View YAML Delete

By default, the system does not create the Istio Ingress gateway after you have deployed Istio.

- In the left-side navigation pane, choose Store > App Catalog.
- Click ack-istio-ingressgateway.



- Click the Values tab, and then set the parameters.



Note:

- For more information about the description, the values, and the default parameters, see the Configuration section on the Readme tab page.

- You can set customized parameters, including indicating whether to enable a specific port, or whether to use the intranet SLB or the Internet SLB by setting the `serviceAnnotations` parameter.

d. In the Deploy area on the right, select the target Cluster from the drop-down list, and then click DEPLOY.



Note:

Namespace is fixed as `istio-system`, and Release Name is fixed as `istio-ingressgateway`.

Verify the result.

- In the left-side navigation pane, choose **Application > Pod**.
- Select the target cluster and the `istio-system` namespace to view the pod to which the Istio Ingress gateway has been deployed.

Container Service - Kubernetes

Overview

Clusters

Nodes

Volumes

Namespace

Authorization

Application

Deployment

StatefulSet

DaeomonSet

Job

CronJob

Pods

Volumes Claim

Pod

Clusters

test-01

Namespace

istio-system

Search By Name

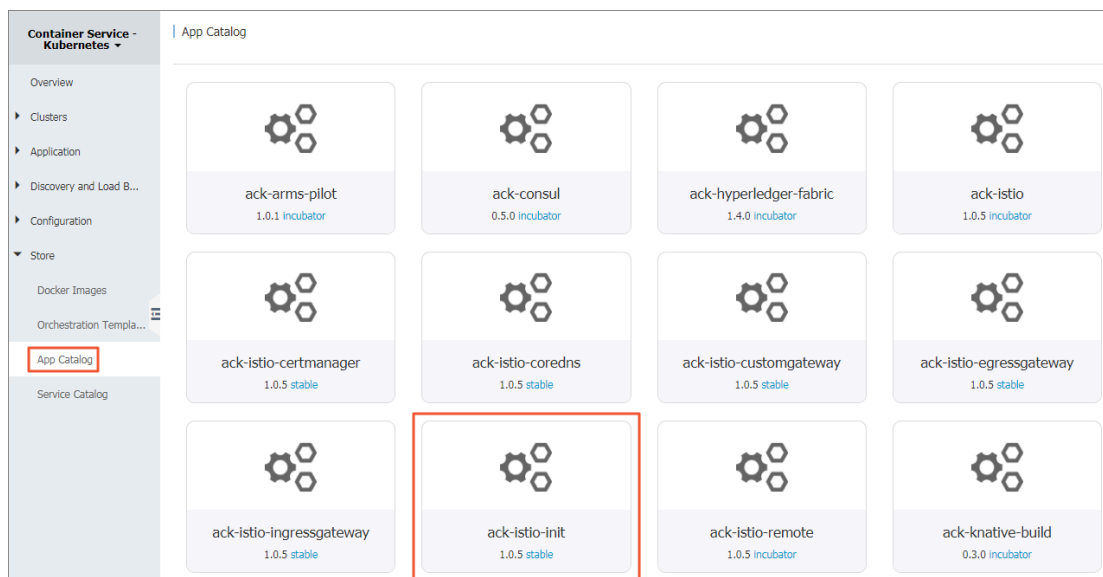
Name	Status	max attempts	Pod IP	Node	Time Created	CPU	Memory	
<div>grafana-6f5c4c7c6d-8h84s</div> <div>grafana:5.2.3</div>	Running	0			04/01/2019,16:24:12	0.003	<div>22.297 Mi</div>	<div>Details</div> <div>More</div>
<div>istio-citadel-6d658f8f88-r1pmt</div> <div>citadel:1.0.5</div>	Running	0			04/01/2019,16:24:12	0.36	<div>9.441 Mi</div>	<div>Details</div> <div>More</div>
<div>istio-galley-c87c85d69-87snf</div> <div>galley:1.0.5</div>	Running	0			04/01/2019,16:24:12	0	<div>7.664 Mi</div>	<div>Details</div> <div>More</div>
<div>istio-grafana-post-install-kvqhf</div> <div>hyperkubev1.7.6_coreos.0</div>	Succeeded	0			04/01/2019,16:24:13			<div>Details</div> <div>More</div>
<div>istio-ingressgateway-84b4868f6b-2tz55</div> <div>proxyv2:1.0.5</div>	Running	0			04/01/2019,16:21:52	0.002	<div>29.48 Mi</div>	<div>Details</div> <div>More</div>
<div>istio-init-crd-10-csq4</div> <div>kubecti:1.0.5</div>	Succeeded	0			04/01/2019,16:23:44			<div>Details</div> <div>More</div>

- Deploy Istio through the App Catalog page

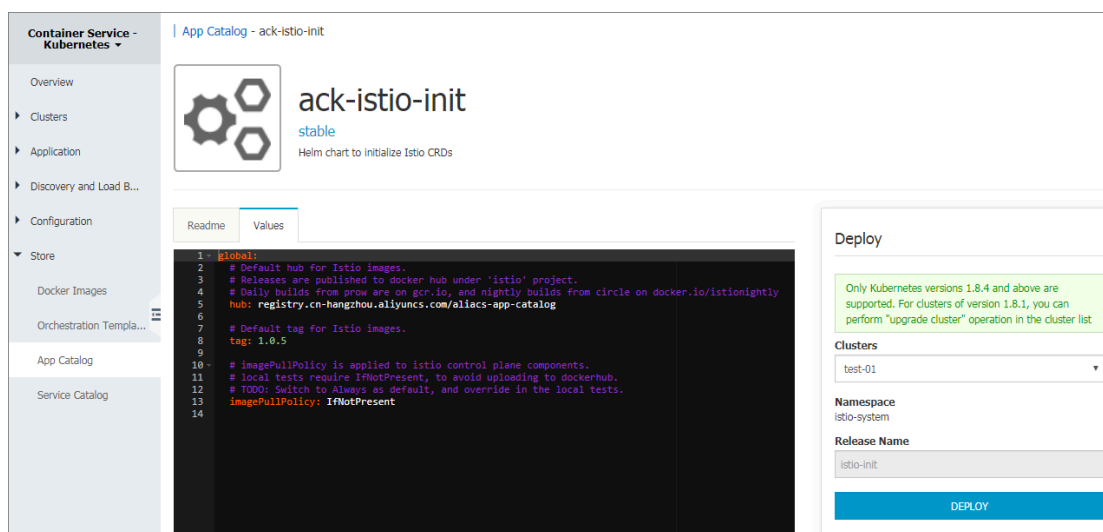
1. Deploy the CRD of Istio.

- a. Log on to the [Container Service console](#).

- b. In the left-side navigation pane, choose Store > App Catalog.



- c. Click ack-istio-init.



- d. In the Deploy area on the right, select the target Cluster from the drop-down list, and then click DEPLOY.

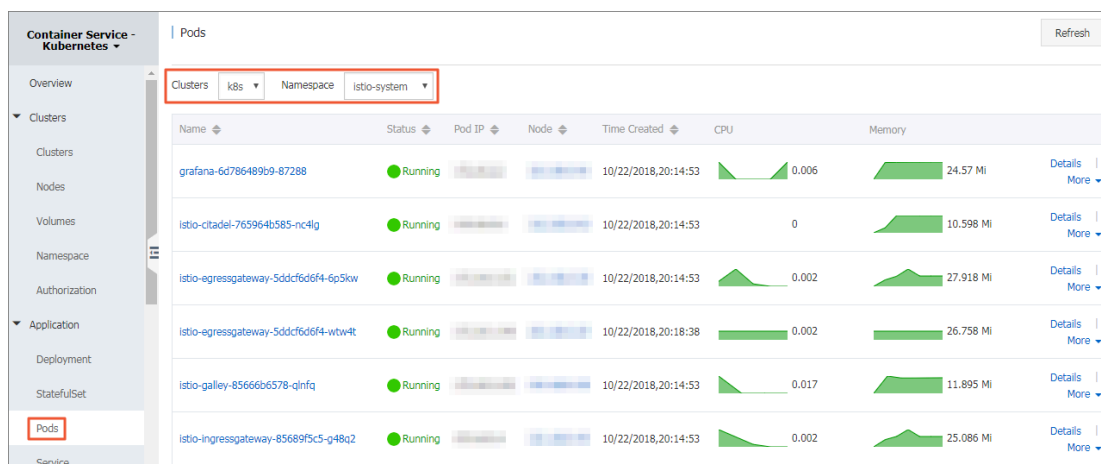


Note:

Namespace is fixed as istio-system, and Release Name is fixed as istio-init.

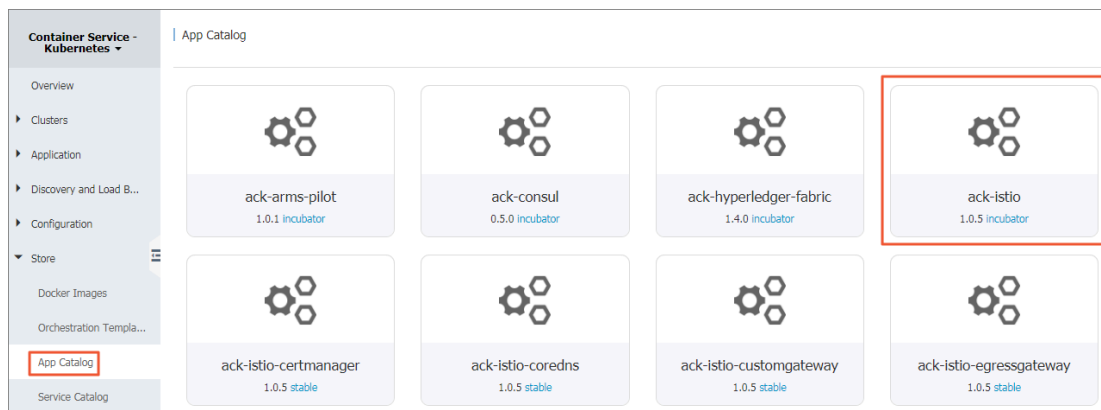
Verify the result.

- In the left-side navigation pane, choose Application > Pods.
- Select the target cluster and the istio-system namespace to view the pod to which the Istio CRD has been deployed.



2. Deploy Istio.

- In the left-side navigation pane, choose Store > App Catalog.
- Click ack-istio.



- Click the Values tab, and then set the parameters.

Container Service - Kubernetes

Overview

Clusters

Application

Discovery and Load B...

Configuration

Store


Docker Images

Orchestration Templa...

App Catalog

Service Catalog

App Catalog - ack-istio



ack-istio

incubator

Helm chart of all istio components for Kubernetes on Alibaba Cloud Container Service

Readme

Values

```

1 # Common settings.
2 - global:
3   # Default hub for Istio images.
4   hub: registry.cn-hangzhou.aliyuncs.com/aliacs-app-catalog
5   # Default tag for Istio images.
6   tag: 1.0.5
7
8   # Gateway used for legacy k8s Ingress resources. By default it is
9   # using 'istio:ingress', to match 0.8 config. It requires that
10  # ingress.enabled is set to true. You can also set it
11  # to ingressgateway, or any other gateway you define in the 'gateway'
12  # section.
13  k8sIngressSelector: ingress
14
15  # k8sIngressHttps will add port 443 on the Ingress and ingressgateway.
16  # It REQUIRES that the certificates are installed in the
17  # expected secrets - enabling this option without certificates
18  # will result in LDS rejection and the ingress will not work.
19  k8sIngressHttps: false
20
21  proxy:
22    image: proxyv2
23
24

```

Deploy

Only Kubernetes versions 1.8.4 and above are supported. For clusters of version 1.8.1, you can perform "upgrade cluster" operation in the cluster list

Clusters

test-01

Namespace

istio-system

Release Name

istio

DEPLOY



Note:

- For more information about the description, the values, and the default parameters, see the Configuration section on the Readme tab page.
- You can set customized values for parameters, including the grafana , prometheus , tracing , kiali , and other parameters, to better meet your requirements.

d. In the Deploy area on the right, select the target Cluster from the drop-down list, and then click DEPLOY.



Note:

Namespace is fixed as istio-system, and Release Name is fixed as istio.

e. Click DEPLOY.

Verify the result.

- Verify that Istio has been deployed to a pod.
 - a. In the left-side navigation pane, choose Application > Pods.
 - b. Select the target cluster and namespace to view the pod to which Istio has been deployed.

Name	Status	Pod IP	Node	Time Created	CPU	Memory
grafana-6d796489b9-87288	Running			10/22/2018,20:14:53	0.006	24.57 Mi
istio-citadel-765964b585-nc4fg	Running			10/22/2018,20:14:53	0	10.598 Mi
istio-egressgateway-5ddcf6d6f4-5p5kw	Running			10/22/2018,20:14:53	0.002	27.918 Mi
istio-egressgateway-5ddcf6d6f4-wtw4t	Running			10/22/2018,20:18:38	0.002	26.758 Mi
istio-galley-85666b6578-q1nfq	Running			10/22/2018,20:14:53	0.017	11.895 Mi
istio-ingressgateway-85689f5c5-g48q2	Running			10/22/2018,20:14:53	0.002	25.086 Mi

- Verify that Istio has been deployed to a service.
 - a. In the left-side navigation pane, choose Application > Service.
 - b. Select the cluster and namespace in which Istio is deployed to view the IP addresses provided by the services to which Istio has been deployed.

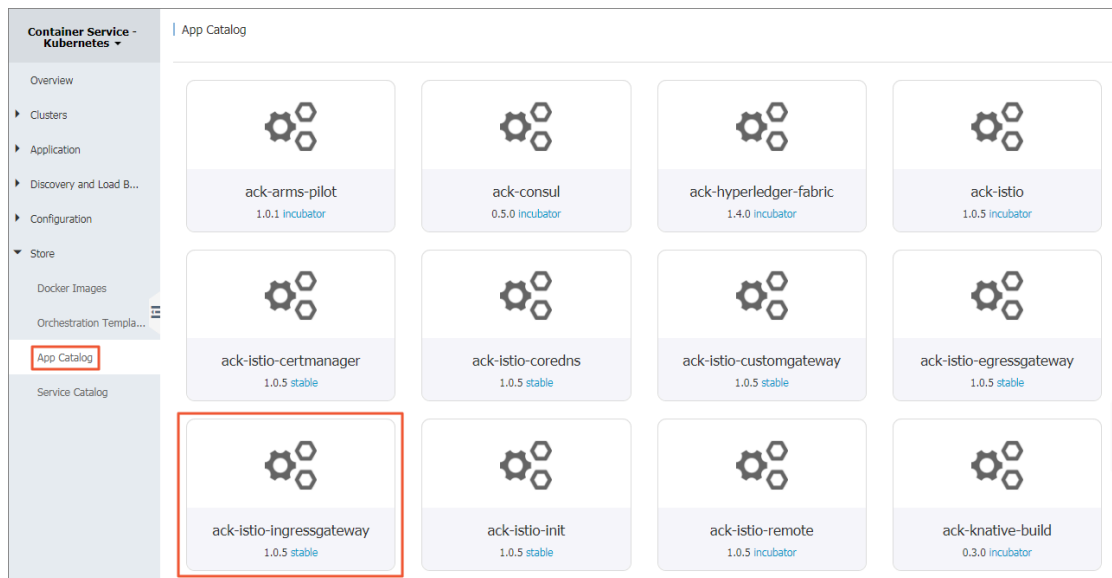
Name	Type	Time Created	ClustersIP	InternalEndpoint	ExternalEndpoint
grafana	ClusterIP	10/22/2018,20:14:52		grafana:3000 TCP	-
istio-citadel	ClusterIP	10/22/2018,20:14:53		istio-citadel:8060 TCP istio-citadel:9093 TCP	-
istio-egressgateway	ClusterIP	10/22/2018,20:14:52		istio-egressgateway:80 TCP istio-egressgateway:443 TCP	-
istio-galley	ClusterIP	10/22/2018,20:14:52		istio-galley:443 TCP istio-galley:9093 TCP	-
istio-ingressgateway	LoadBalancer	10/22/2018,20:14:52		istio-ingressgateway:80 TCP istio-ingressgateway:31380 TCP istio-ingressgateway:443 TCP istio-ingressgateway:31390 TCP istio-ingressgateway:31400 TCP istio-ingressgateway:31400 TCP istio-ingressgateway:15011 TCP istio-ingressgateway:32009 TCP istio-ingressgateway:8060 TCP istio-ingressgateway:32522 TCP istio-ingressgateway:853 TCP istio-ingressgateway:32512 TCP istio-ingressgateway:15030 TCP istio-ingressgateway:30282 TCP istio-ingressgateway:15031 TCP istio-ingressgateway:31006 TCP	-

3. Create the Istio Ingress gateway.

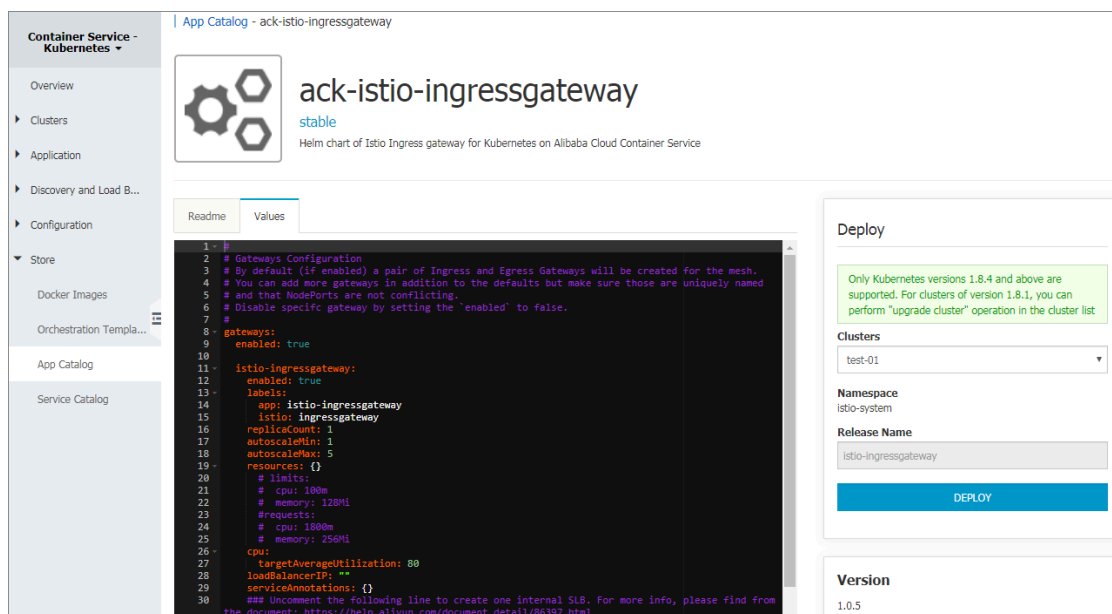
**Note:**

By default, the system does not create the Istio Ingress gateway after you have deployed Istio.

- In the left-side navigation pane, choose Store > App Catalog.
- Click ack-istio-ingressgateway.



- Click the Values tab, and then set the parameters.

**Note:**

- For more information about the description, the values, and the default of the parameters, see the Configuration section on the Readme tab page.

- You can set customized parameters. For example, you can enable a specific port, or use the intranet SLB or the Internet SLB by set the serviceAnnotations parameter.

d. In the Deploy area on the right, select the target Cluster from the drop-down list, and then click DEPLOY.



Note:

Namespace is fixed as istio-system, and Release Name is fixed as istio-ingressgateway.

Verify the result.

- In the left-side navigation pane, choose Application > Pod.
- Select the target cluster and the istio-system namespace to view the pod to which the Istio Ingress gateway has been deployed.

Container Service - Kubernetes

Overview

Clusters

Nodes

Volumes

Namespace

Authorization

Application

Deployment

StatefulSet

DaemonSet

Job

CronJob

Pods

Volumes Claim

Pod

Clusters

test-01

Namespace

istio-system

Search By Name

Name	Status	max attempts	Pod IP	Node	Time Created	CPU	Memory	
grafana-6f5c47c6d-8h84s grafana:5.2.3	Running	0			04/01/2019,16:24:12	0.003	22.297 Mi	<div>Details</div> <div>More</div>
istio-citadel-6d658f88-ripmt citadel:1.0.5	Running	0			04/01/2019,16:24:12	0.36	9.441 Mi	<div>Details</div> <div>More</div>
istio-galley-c87c85d69-87snf galley:1.0.5	Running	0			04/01/2019,16:24:12	0	7.664 Mi	<div>Details</div> <div>More</div>
istio-grafana-post-install-kvghf hyperkube:v1.7.6_coreos.0	Succeeded	0			04/01/2019,16:24:13			<div>Details</div> <div>More</div>
istio-ingressgateway-84b4868f6b-2tz55 proxyv2:1.0.5	Running	0			04/01/2019,16:21:52	0.002	29.48 Mi	<div>Details</div> <div>More</div>
istio-init-crd-10-csq4 kubecti:1.0.5	Succeeded	0			04/01/2019,16:23:44			<div>Details</div> <div>More</div>

1.18.3 Update Istio

You can modify the deployed Istio through updates.

Prerequisites

- You have created an Kubernetes cluster. For more information, see [Create a Kubernetes cluster](#).
- You have created an Istio. For more information, see [Deploy Istio](#).

Procedure

1. Log on to the [Container Service console](#).

2. Under the Kubernetes menu, click Application > Helm in the left-side navigation pane.
3. Select a cluster, select the Istio to be updated, and click Update in the action column.

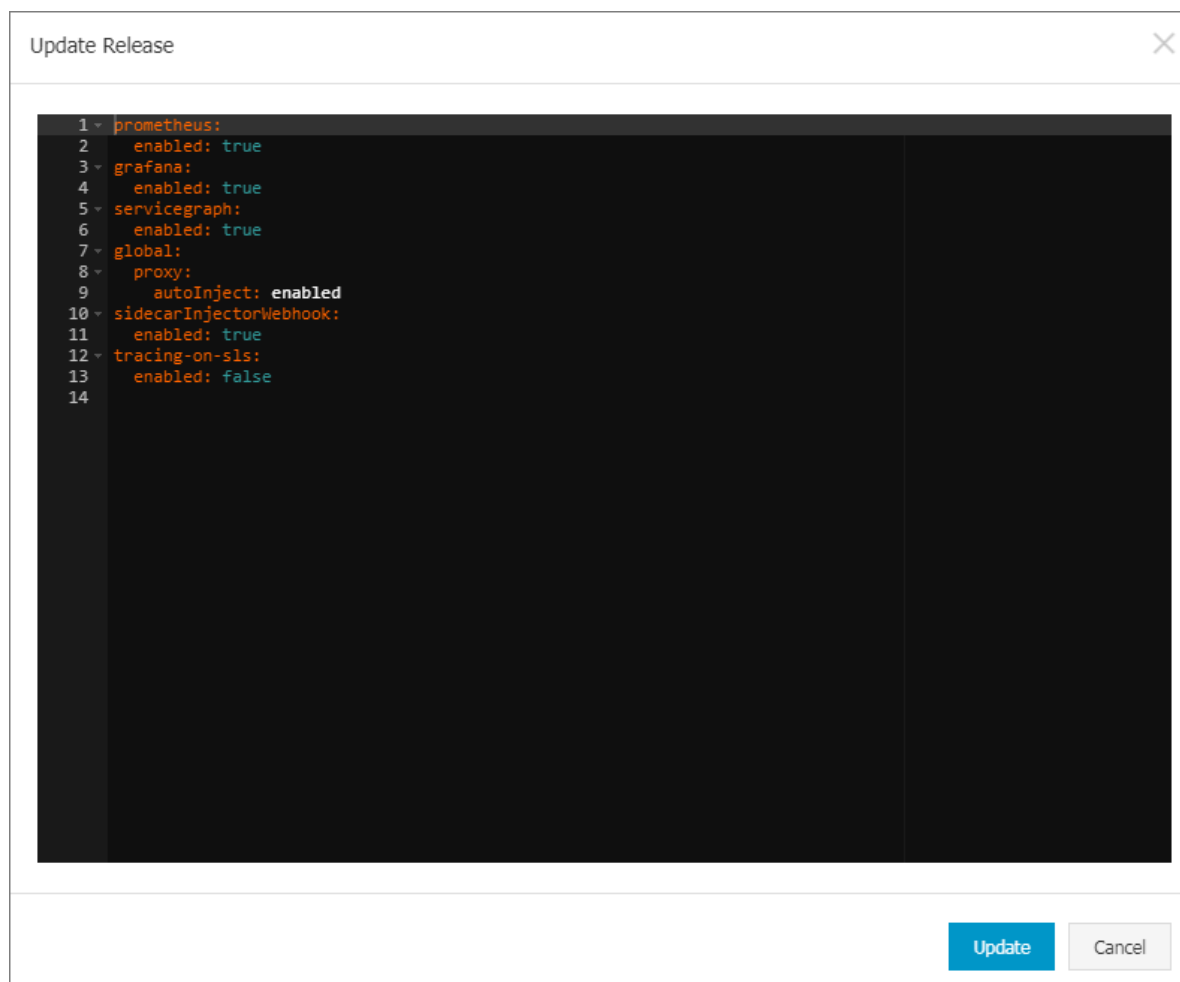
**Note:**

- The release name of the Istio that is deployed through the cluster interface is istio. Configurations to be updated are the same as the options configured in deployment.
- The release name of the Istio that is deployed through the application catalog is the name specified when you create the Istio. Configurations to be updated are the same as the options configured in deployment.

Release List							Refresh
Clusters k8s							
Release Name	Status	Namespace	Chart Name	Chart Version	App Version	Update Time	Action
ack-istio-default	Deployed	istio-system	ack-istio	1.0.2	1.0.2	10/22/2018,20:50:57	Details Update Delete
istio	Deployed	istio-system	ack-istio	1.0.2	1.0.2	10/22/2018,20:14:52	Details Update Delete

4. In the displayed dialog box, modify parameters of the Istio, and then click Update.

In this example, update the Istio that is deployed through the cluster interface:



Result

You can view updated content in two ways:

- After you complete the update, the page automatically jumps to the Release List page. On the Resource tab, you can view updated content.
- Under the Kubernetes menu, click Application > Pods, and select the target cluster and namespace to view updating results.

1.18.4 Delete Istio

You can delete a deployed Istio through the deleting operation.

Prerequisites

- You have created a Kubernetes cluster. For more information, see [Create a Kubernetes cluster](#).
- You have created an Istio. For more information, see [Deploy Istio](#).

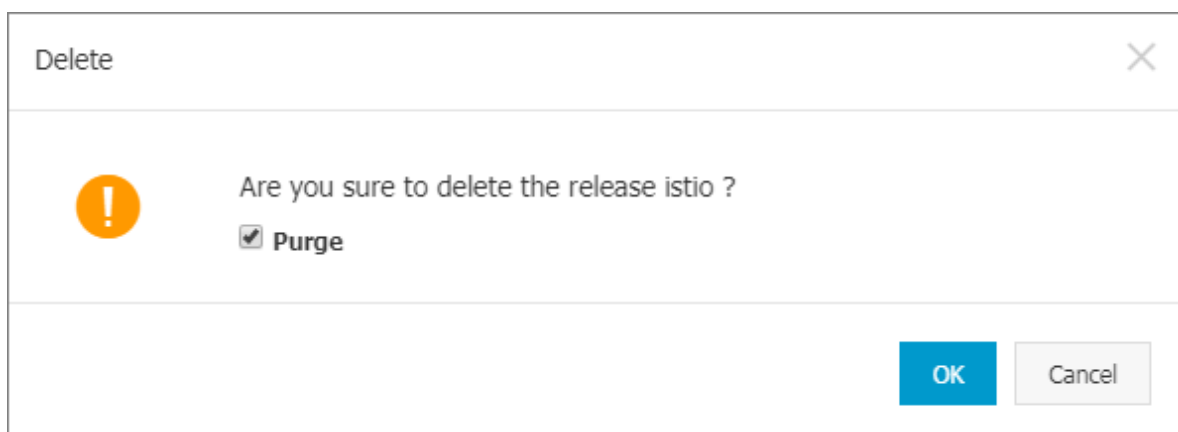
Procedure

1. Log on to the [Container Service console](#).
2. Under the Kubernetes menu, click Application > Helm in the left-side navigation pane.
3. Select a cluster, select the Istio to be deleted, and click Delete in the action column.




Release List							Refresh
Clusters: k8s							
Release Name	Status	Namespace	Chart Name	Chart Version	App Version	Update Time	Action
ack-istio-default	Deployed	istio-system	ack-istio	1.0.2	1.0.2	10/22/2018,20:50:57	Details Update Delete
istio	Deployed	istio-system	ack-istio	1.0.2	1.0.2	10/22/2018,20:14:52	Details Update Delete

4. In the displayed dialog box, click OK.



Delete



Are you sure to delete the release istio ?
☒ Purge

OK

Cancel



Note:

- Do not select the Purge box:
- Releasing records are not deleted:

Release List							Refresh
Clusters k8s							
Release Name	Status	Namespace	Chart Name	Chart Version	App Version	Update Time	Action
ack-istio-default	Deployed	istio-system	ack-istio	1.0.2	1.0.2	10/22/2018,20:50:57	Details Update Delete
istio	Deleted	istio-system	ack-istio	1.0.2	1.0.2	10/22/2018,20:14:52	Details Update Delete


- The name of this Istio cannot be used again.

When you redeploy the Istio through the cluster interface, the deployment status is deployed.

Step	Status
Create Istio Resource Definition	pending
Deploy Istio	pending
Deployed	

When you redeploy the Istio through the application catalog, the system prompts you that deployment or resource with the same name already exists and please modify the Istio name.

Note



A release with the same name already exists, please edit the name

Can't install release with errors: rpc error: code = Unknown desc = a release named ack-istio-default already exists. Run: helm ls --all ack-istio-default; to check the status of the release Or run: helm del --purge ack-istio-default; to delete it

OK

- Selecting the Purge box deletes all releasing records and the Istio name can be reused.

We recommend that you keep the Purge box selected.

Result

Back to the Release List page, you can see that the Istio is removed.

1.18.5 Upgrade Istio components

This topic describes how to upgrade Istio components.

Background information

- The Istio upgrade may install new binaries, and change configurations and API schemas.
- The upgrade process may cause service downtime.
- To minimize downtime, use multiple replicas to ensure that your Istio control plane components and your applications remain highly available.



Note:

In the following example, assume that the Istio components are installed and upgraded in the istio-system namespace.

Procedure

To complete the upgrade process, you need to upgrade CRD files, the control plane, and the data plane sidecar.

Upgrade CRD files

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Container Service-Kubernetes, choose Applications > Releases. Then, click the Helm tab, find the istio-init release, and delete it.



Note:

If no release named istio-init is displayed on the Helm tab page, you can directly perform the next step.

3. In the left-side navigation pane under Container Service-Kubernetes, choose Store > App Catalog. Then, click ack-istio-init.
4. In the Deploy area, select the target cluster. Then, click DEPLOY.



Note:

By default, the istio-system namespace is selected, and the release name is set to istio-init.

Upgrade the control plane

The Istio control plane components include the Citadel, Pilot, Policy, Telemetry, and Sidecar injector.

1. In the left-side navigation pane under Container Service-Kubernetes, choose Applications > Releases.
2. Select the target cluster, select the target Release Name, and click Upgrade in the Action column.
3. In the configurations of the deployed Istio, specify the version number of the Istio to be deployed.



Note:

On this page, you can also modify other parameters.

```
global :
    tag : < enter the version number >
```

4. Click Update.

Upgrade the data plane sidecar

Note that after you upgrade the control plane, the applications that have already run Istio will still use the sidecar of an earlier version. To upgrade the sidecar, you need to re-inject it.

Automatic sidecar injection

If you use automatic sidecar injection, you can upgrade the sidecar by performing a rolling update for all pods. Then, the sidecar of the new version will be automatically re-injected.

You can use the following script to trigger the rolling update by patching the termination grace period.

```
NAMESPACE =$ 1
DEPLOYMENT _LIST=$( kubectl -n $ NAMESPACE get
deployment -o jsonpath='{. items [*]. metadata . name }')
echo " Refreshing pods in all Deployment s : $
DEPLOYMENT _LIST "
for deployment _name in $ DEPLOYMENT _LIST ; do
# echo " get TERMINATIO N_GRACE_PE RIOD_SECON DS from
deployment : $ deployment _name "
TERMINATIO N_GRACE_PE RIOD_SECON DS=$( kubectl -n $
NAMESPACE get deployment "$ deployment _name " -o jsonpath
='{. spec . template . spec . terminatio nGracePeri odSeconds }')
if [ "$ TERMINATIO N_GRACE_PE RIOD_SECON DS " -eq 30 ];
then
TERMINATIO N_GRACE_PE RIOD_SECON DS =' 31 '
else
TERMINATIO N_GRACE_PE RIOD_SECON DS =' 30 '
```

```

fi
patch_strategy="{\"spec\":{\"template\":{\"spec\":{\"terminationGracePeriodSeconds\":$TERMINATION_GRACE_PERIOD_SECONDS}}}}\"
# echo $patch_strategy
kubectl -n $NAMESPACE patch deployment $deployment_name -p $patch_strategy
done
echo " done ."

```

Manual sidecar injection

Run the following command to manually upgrade the sidecar:

```

kubectl apply -f <(istioctl kube-inject -f $ORIGINAL_DEPLOYMENT_YAML)

```

If the sidecar was previously injected with some customized injection configuration files, run the following command to manually upgrade the sidecar:

```

kubectl apply -f <(istioctl kube-inject --injectConfigFile inject-config.yaml --filename $ORIGINAL_DEPLOYMENT_YAML)

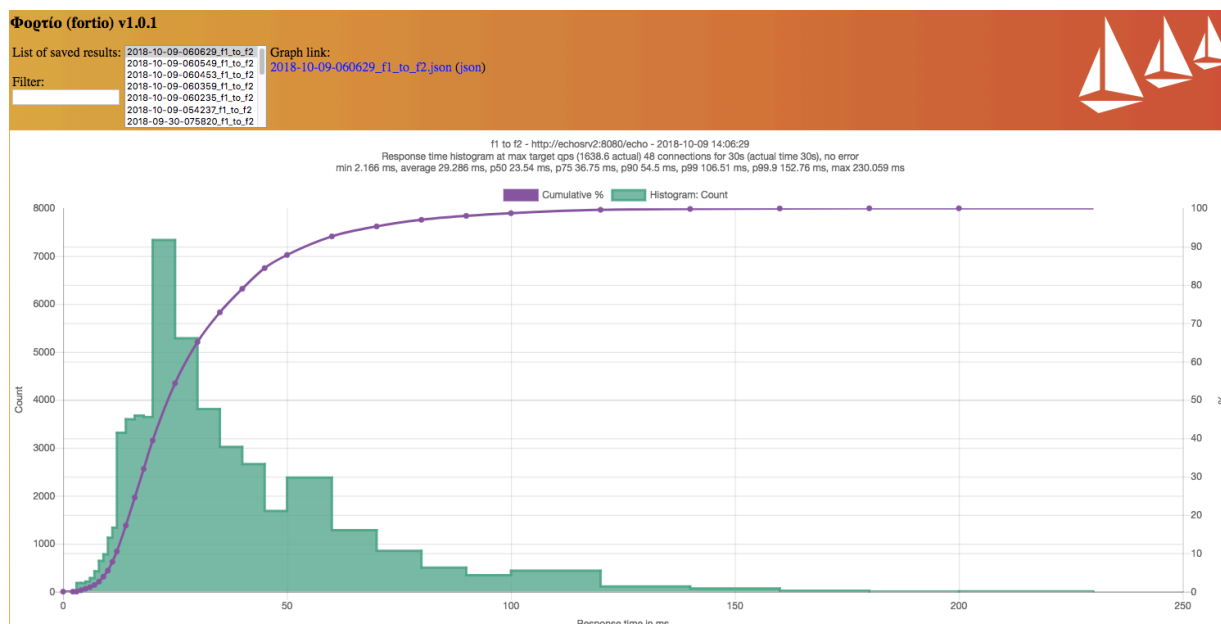
```

Impacts caused by the Istio upgrade

Impacts caused by the CRD file upgrade

The upgrade process does not impact the calls between services within the cluster or the calls from the gateway to services.

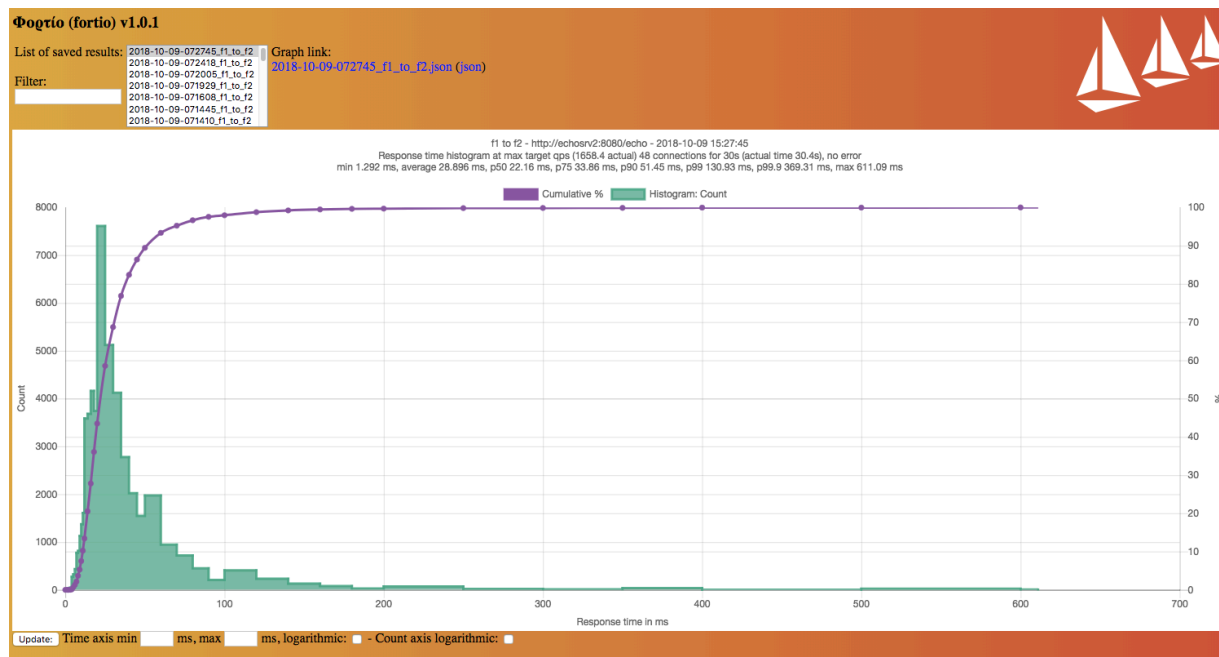
Calls between services within the cluster.



Impacts caused by the control plane upgrade

If HA is enabled, that is, the replicas of Pilot is 2, the HPA setting of `istio - pilot` / `istio - policy` / `istio - telemetry` is `minReplicas: 2`.

If you have changed the Istio version multiple times by upgrading or rolling back the component version, testing results will indicate that the QPS of calls between services remains unchanged and the calls proceed normally.



Impacts caused by the control plane sidecar upgrade

No obvious change occurs to both the QPS of the calls between services within the cluster and the QPS of the calls from the gateway to services. But these calls will terminate temporarily. We recommend that you use multiple replicas to upgrade the sidecar to reduce the impacts.

1.19 Template management

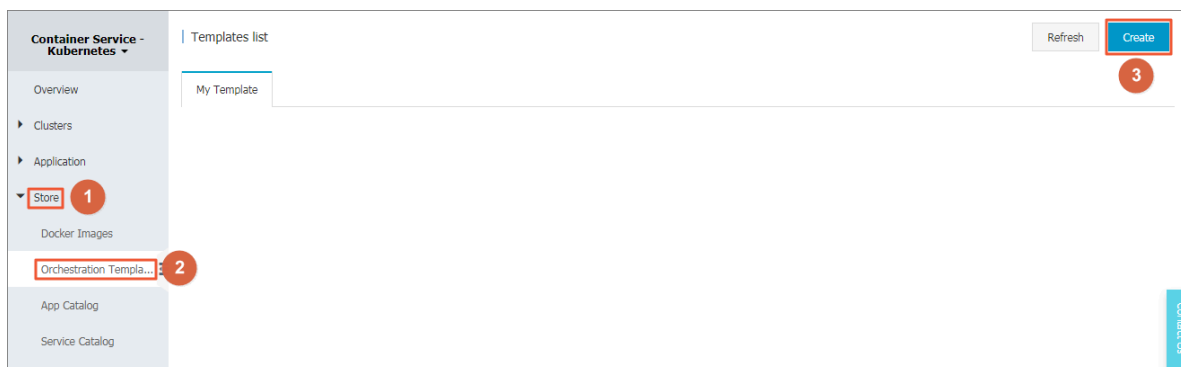
1.19.1 Create an orchestration template

You can use multiple methods to create orchestration templates through the Container Service console.

Procedure

1. Log on to the [Container Service console](#).

2. In the left-side navigation pane under Container Service-Kubernetes, choose Store > Orchestration Templates. Then, in the upper-right corner, click Create.



3. In the displayed dialog box, configure the orchestration template, and then click Save. In this example, build a tomcat application template that contains a deployment and a service.

- **Name:** Set the template name.
- **Description:** Enter the description for the template. This parameter is optional.
- **Template:** Configure the template that conforms to Kubernetes yaml syntax rules. The template can contain multiple resource objects that are separated by `---`.

Create

Name:

tomcat

The name should be 1-64 characters long, and can contain numbers, English letters, Chinese characters and hyphens.

Description:

tomcat application

Template:

```

1  apiVersion: apps/v1beta2 # for versions before 1.8.0 use
   apps/v1beta1
2  kind: Deployment
3  metadata:
4    name: tomcat-deployment
5    labels:
6      app: tomcat
7  spec:
8    replicas: 1
9    selector:
10     matchLabels:
11       app: tomcat
12   template:
13     metadata:
14       labels:
15         app: tomcat
16     spec:
17       containers:
18         - name: tomcat
19           image: tomcat # replace it with your exactly
20           <image_name:tags>
21           ports:
22             - containerPort: 8080

```

Save

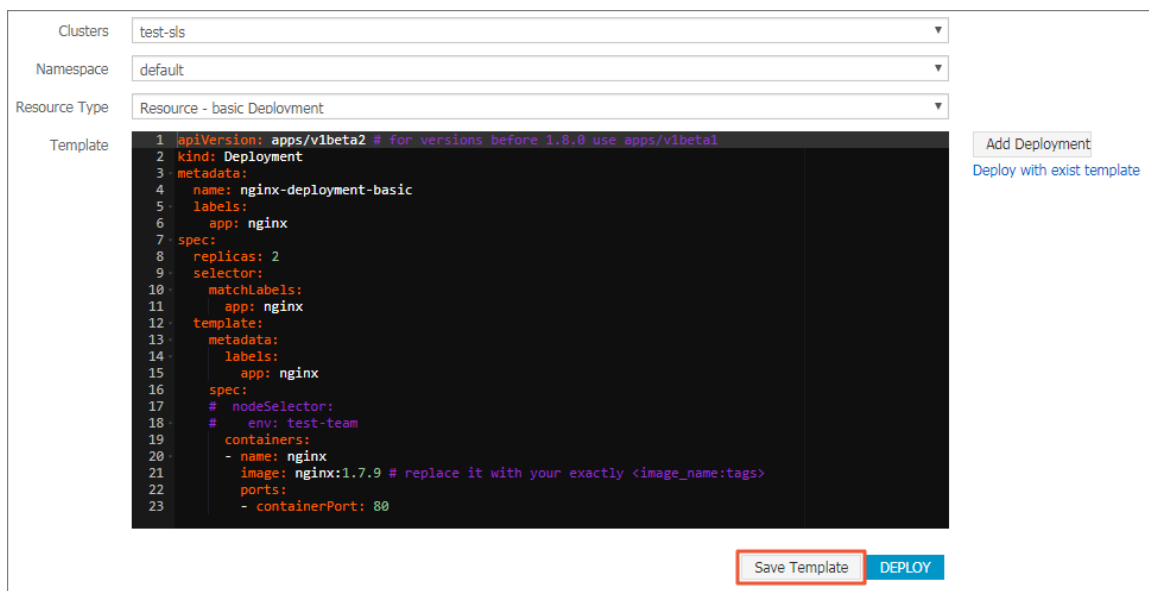
Cancel

4. After the template is created, the Template List page is displayed. You can see the template under My Template.



5. Optional: In the left-side navigation pane under Container Service-Kubernetes, choose Applications > Deployments. Then, in the upper-right corner, click Create by Template. Save one of orchestration templates built-in Container Service as your custom template.

- a) Select a built-in template and click Save Template.



- b) In the displayed dialog box, configure the name, description, and template. After completing the configurations, click Save.



Note:

You can modify the built-in template.

- c) Choose Store > Orchestration Templates, the created template is displayed under My Template.



What's next

You can quickly create an application by using the orchestration template under My Template.

1.19.2 Edit an orchestration template

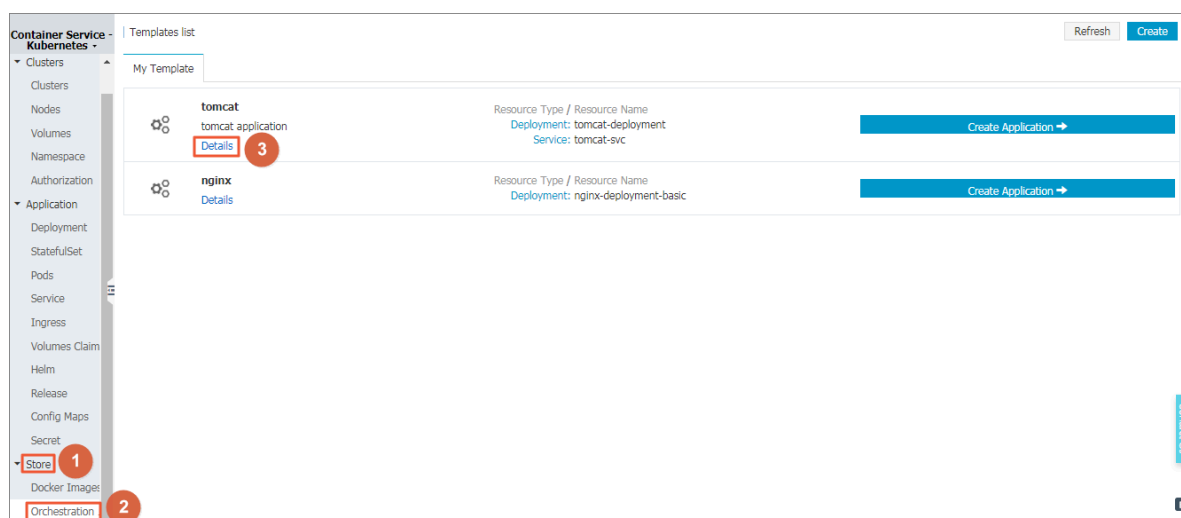
You can edit an orchestration arrangement template.

Prerequisites

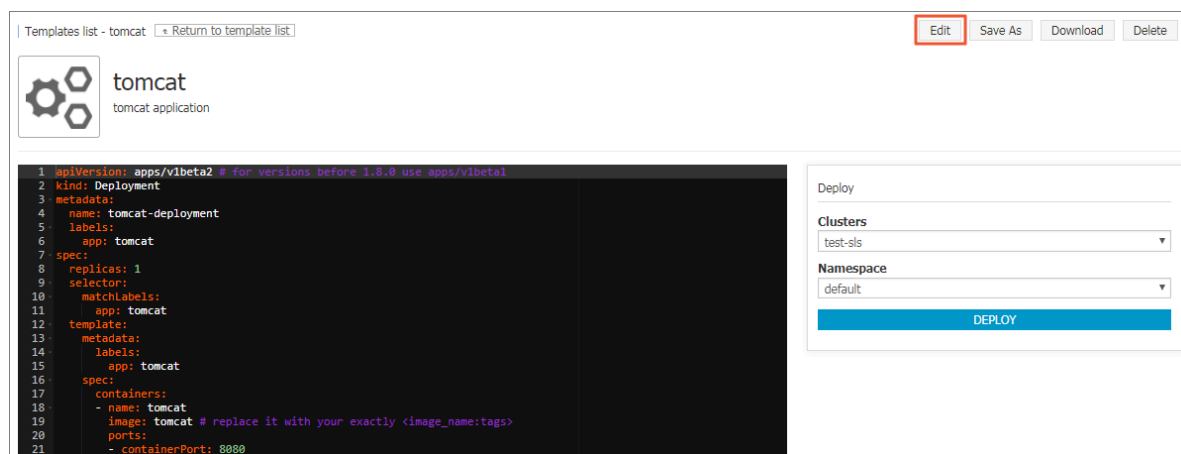
You have created an orchestration template, see [Create an orchestration template](#).

Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Store > Orchestration Templates. Existing orchestration templates are displayed under My Template.
3. Select a template and click Details.



4. Click Edit in the upper-right corner.



5. In the displayed dialog box, edit the name, description, and template, and click Save.

Modify template

Name:

tomcat-V2

The name should be 1-64 characters long, and can contain numbers, English letters, Chinese characters and hyphens.

Description:

tomcat application

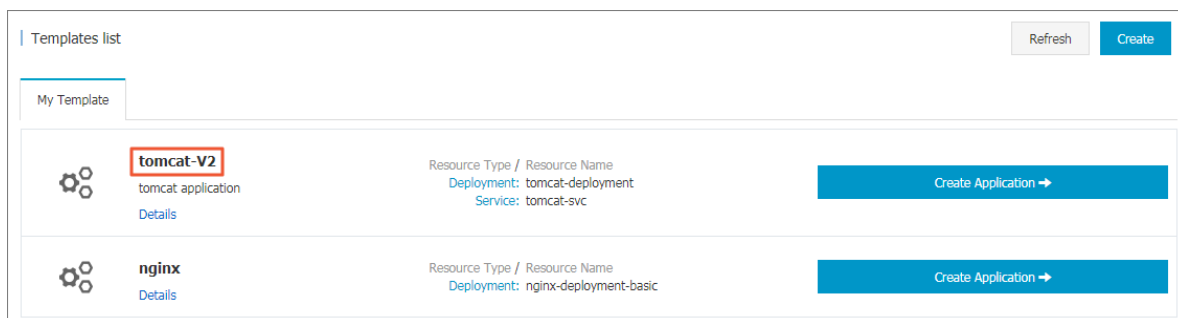
Template:

```
1 apiVersion: apps/v1beta2 # for versions
  before 1.8.0 use apps/v1beta1
2 kind: Deployment
3 metadata:
4   name: tomcat-deployment
5   labels:
6     app: tomcat
7 spec:
8   replicas: 1
9   selector:
10    matchLabels:
11      app: tomcat
12   template:
13     metadata:
14       labels:
15         app: tomcat
16     spec:
17       containers:
18         - name: tomcat
19           image: tomcat # replace it with your
  exactly <image_name:tags>
20           ports:
21             - containerPort: 8080
22
```

Save

Cancel

6. Back to the Template List page, under My Template, you can see the template is changed.



1.19.3 Save an existing orchestration template as a new one

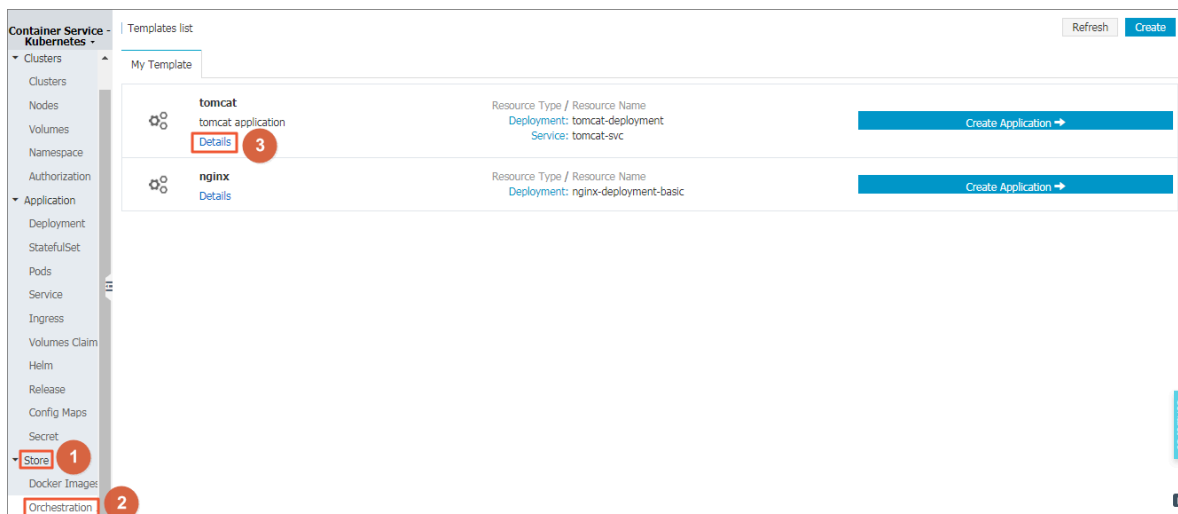
You can save an existing template as a new one.

Prerequisites

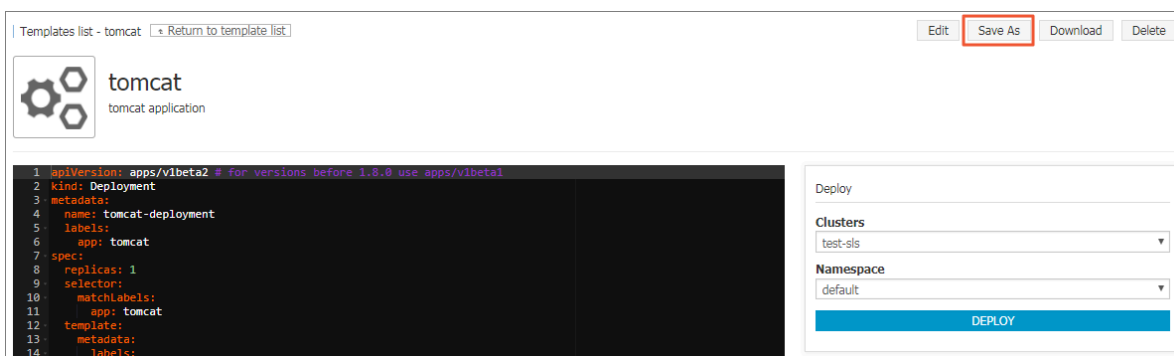
You have created an orchestration template, see [Create an orchestration template](#).

Procedure

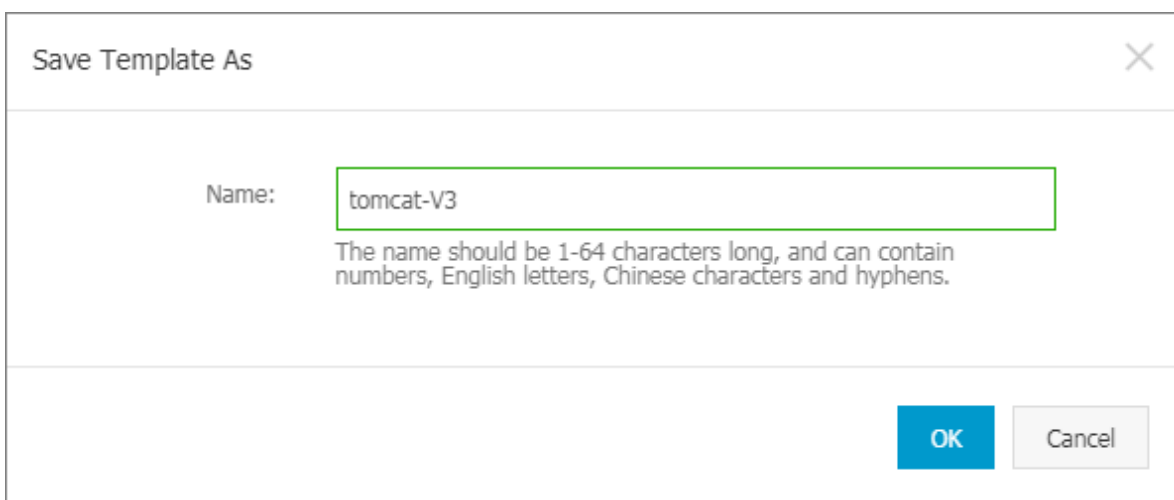
1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Store > Orchestration Templates. Existing orchestration templates are displayed under My Template.
3. Select a template and click Details.



4. You can modify the template and click Save as in the upper-right corner.



5. In the displayed dialog box, configure the template name and click OK.



6. Back to the Template List page, you can see that the saved template is displayed under My Template.



1.19.4 Download an orchestration template

You can download an existing orchestration template.

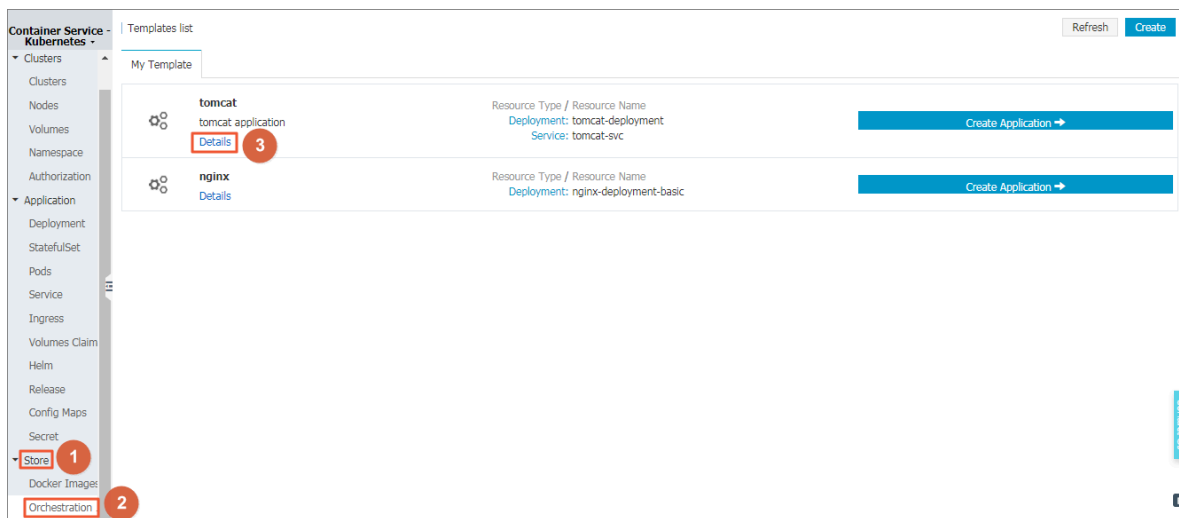
Prerequisites

You have created an orchestration template, see [Create an orchestration template](#).

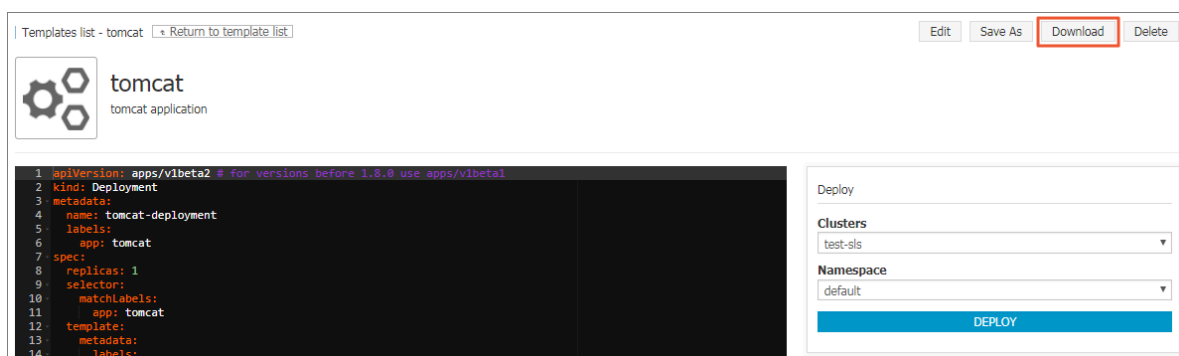
Procedure

1. Log on to the [Container Service console](#).

2. Under Kubernetes, click Store > Orchestration Templates. Existing orchestration templates are displayed under My Template.
3. Select a template and click Details.



4. Click Download in the upper-right corner, a template file with yml suffix is downloaded immediately.



1.19.5 Delete an orchestration template

You can delete an orchestration template that is no longer needed.

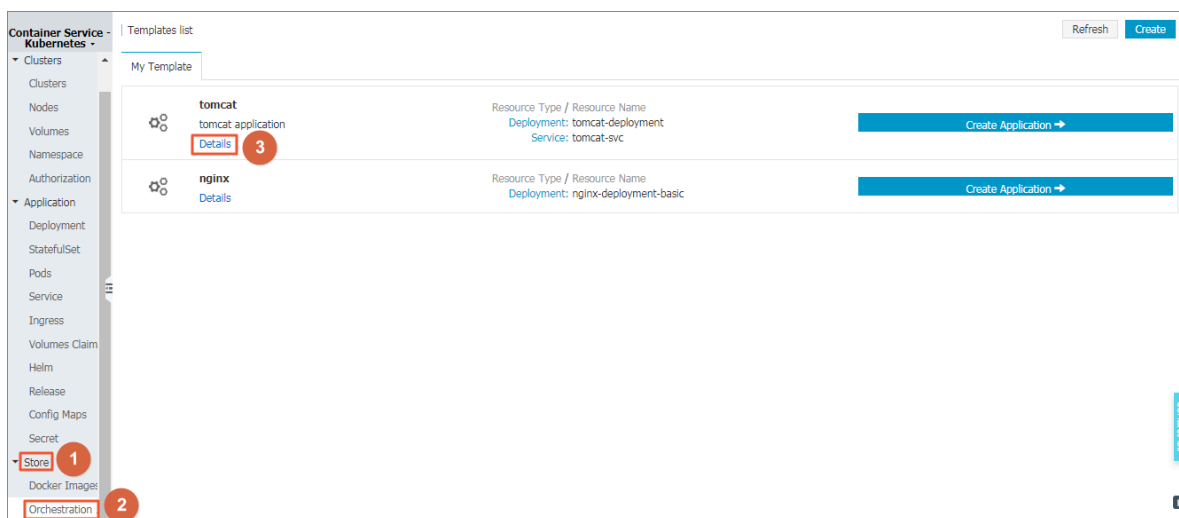
Prerequisites

You have created an orchestration template, see [Create an orchestration template](#).

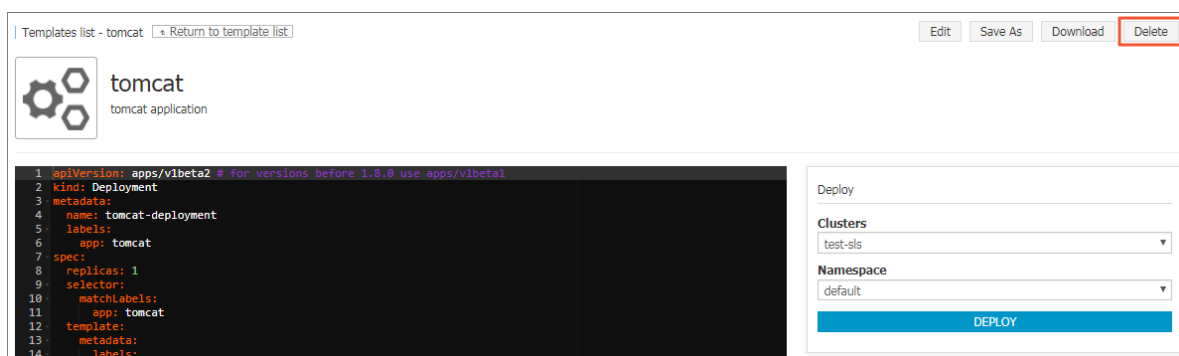
Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Store > Orchestration Template. Existing orchestration templates are displayed under My Template on the Template list page.

3. Select a template and click Detail.



4. On the detail page of the template, you can click Delete in the upper-right corner.



5. Click Confirm in the displayed dialog box.

1.20 App catalog management

1.20.1 App catalog overview

Microservice is the theme of container era. The application microservice brings great challenge to the deployment and management. By dividing a large single application into several microservices, the microservice can be independently deployed and extended so as to realize the agile development and fast iteration. Microservice brings great benefits to us. However, developers have to face the management issues of the microservices, such as the resource management, version management, and configuration management. The number of microservices is large because an application is divided into many components that correspond to many microservices.

For the microservice management issues under Kubernetes orchestration, Alibaba Cloud Container Service introduces and integrates with the Helm open-source project to help simplify the deployment and management of Kubernetes applications.

Helm is an open-source subproject in the Kubernetes service orchestration field and a package management tool for Kubernetes applications. Helm supports managing and controlling the published versions in the form of packaging softwares, which simplifies the complexity of deploying and managing Kubernetes applications.

Alibaba Cloud app catalog feature

Alibaba Cloud Container Service app catalog feature integrates with Helm, provides the Helm-related features, and extends the features, such as providing graphic interface and Alibaba Cloud official repository.

The chart list on the App Catalog page includes the following information:

- **Chart name:** A Helm package corresponding to an application, which contains the image, dependencies, and resource definition required to run an application.
- **Version:** The version of the chart.
- **Repository:** The repository used to publish and store charts, such as the official repository stable and incubator.

The information displayed on the details page of each chart may be different and include the following items:

- Chart introduction
- Chart details
- Prerequisites for installing chart to the cluster, such as pre-configuring the persistent storage volumes (pv)
- Chart installation commands
- Chart uninstallation commands
- Chart parameter configurations

Currently, you can deploy and manage the charts in the app catalog by using the Helm tool. For more information, see [Simplify Kubernetes application deployment by using Helm](#).

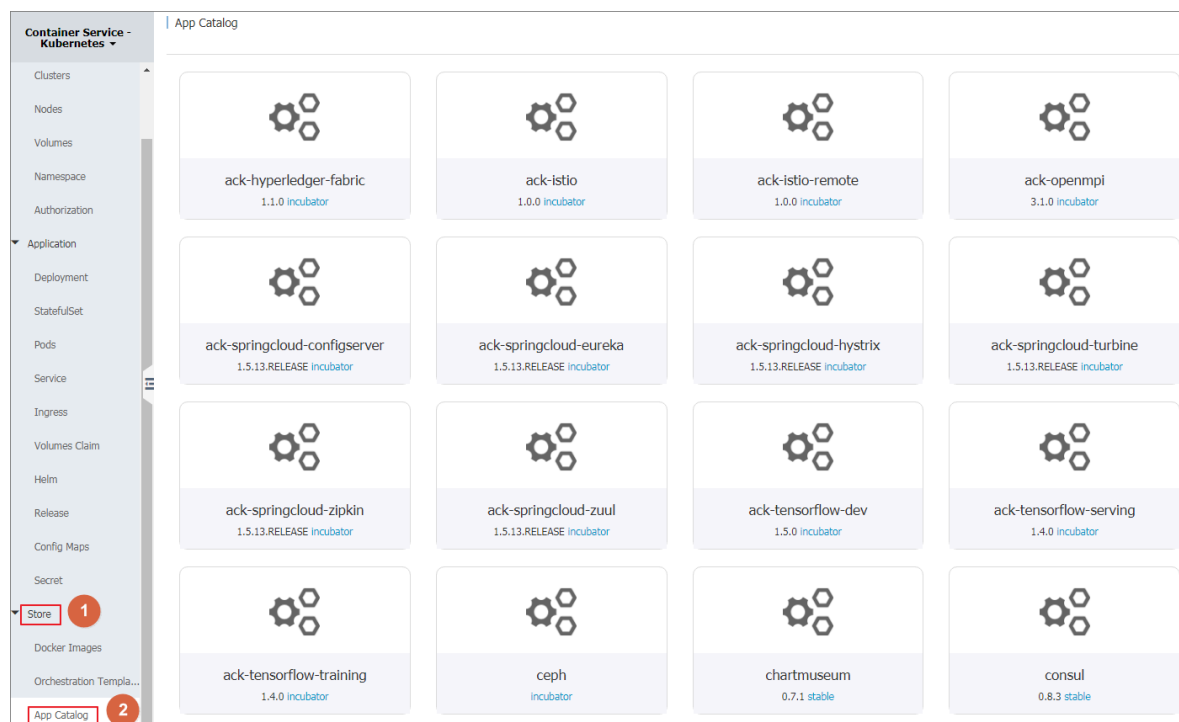
1.20.2 View app catalog list

Procedure

1. Log on to the [Container Service console](#).

2. Under Kubernetes, click Store > App Catalog in the left-side navigation pane.

View the charts on the App Catalog page, each of which corresponds to an application, containing some basic information such as the application name, version, and source repository.



What's next

You can click to enter a chart and get to know the detailed chart information. Deploy the application according to the corresponding information by using the Helm tool. For more information, see [Simplify Kubernetes application deployment by using Helm](#).

1.21 Auto Scaling

1.21.1 Autoscale a Kubernetes cluster

This topic describes how to autoscale a Kubernetes cluster to meet the requirements of your Kubernetes cluster workload. Alibaba Cloud Container Service for Kubernetes

(ACK) provides the capability to autoscale a Kubernetes cluster through using the cluster autoscaler program.

Background information

You can set the the cluster autoscaler to add different ECS instance types to your Kubernetes cluster, such as the general, GPU, and preemptive instance types. You can set multiple zones, instance specifications, and autoscaling modes.

Cluster autoscaler overview

The cluster autoscaler changes the size of a Kubernetes cluster based on the use of resource in the nodes of a pod in a Kubernetes cluster. Resource usage is calculated based on the pod resource requests.

When a pod requests more resources than what the associated node can provide, the pod enters the pending status. At which time, the autoscaler calculates the change to the size of cluster. It does so by calculating the number of nodes necessary to provide the requested resource with regard to the resource specification and threshold that you set for an autoscaling group.

For example, if you set a low threshold value for the number of nodes in an autoscaling group, the cluster autoscaler deletes a node, which reduces the amount of resources that the pod can request.

Notes

- By default, your account can use up to 30 Pay-As-You-Go ECS instances in all your clusters, and the route table of one VPC can contain up to 50 entries. To increase the number of available ECS instances or entries in a route table of one VPC, open a ticket.
- For a single type of ECS instances, the number of ECS instances of one specification that is permitted at one time varies frequently. Therefore, we recommend that you set multiple ECS instance types of one ECS specification.
- When a node for which you set the fast scaling mode is shut down and reclaimed, it is in the `NotReady` status. When the node is reused by the cluster autoscaler, the node enters the `Ready` status.
- When a node for which you set the fast scaling mode is shut down and reclaimed, only the disks attached to the node are charged (except that the node uses local disks, for example, `ecs.d1ne.2xlarge`).

Enable cluster autoscaling

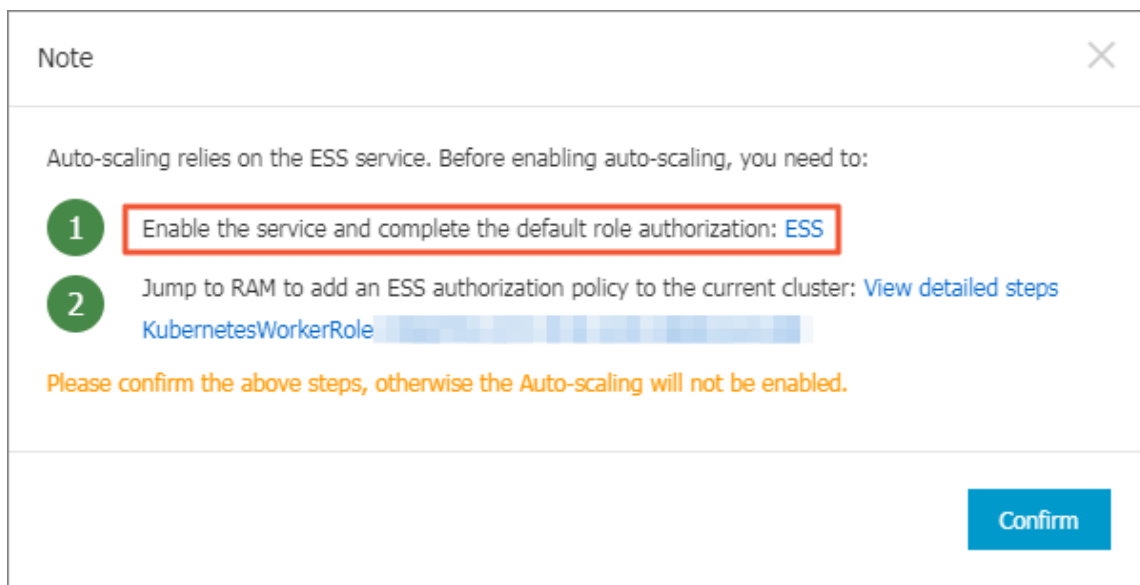
1. Log on to the [Container Service console](#).
2. In the left-side navigation pane under Container Service-Kubernetes, choose **Clusters > Clusters**.
3. Find the target cluster. Then, in the Action column, choose **More > Auto Scaling**.

Cluster Name/ID	Cluster Type	Region (All) ▾	Network Type	Cluster Status	Time Created	Kubernetes Version	Action
mymanagedk8s1	ManagedKubernetes	China North 2 (Beijing)	VPC vpc-zzetk0nc8jr...	Running	09/07/2018,16:01:46	1.10.4	Manage View Logs Dashboard Scale Cluster More ▾
myserverlessk8s1	Serverless Kubernetes	China East 1 (Hangzhou)	VPC vpc-bp1bw6x01cy...	Running	09/04/2018,20:13:20	1.9.7	Manage View Logs Delete
mytest1	Kubernetes	China North 2 (Beijing)	VPC vpc-zze7c50xju...	Running	09/03/2018,10:06:49	1.10.4	Manage View Logs Dashboard Scale Cluster More ▾
myk8s1	Kubernetes	China East 1 (Hangzhou)	VPC vpc-bp1faadn5o...	Running	08/26/2018,11:53:17	1.10.4	Manage Delete Add Existing Instance Upgrade Cluster Auto-scaling Addon Upgrade Upgrade monitoring service Deploy Istio

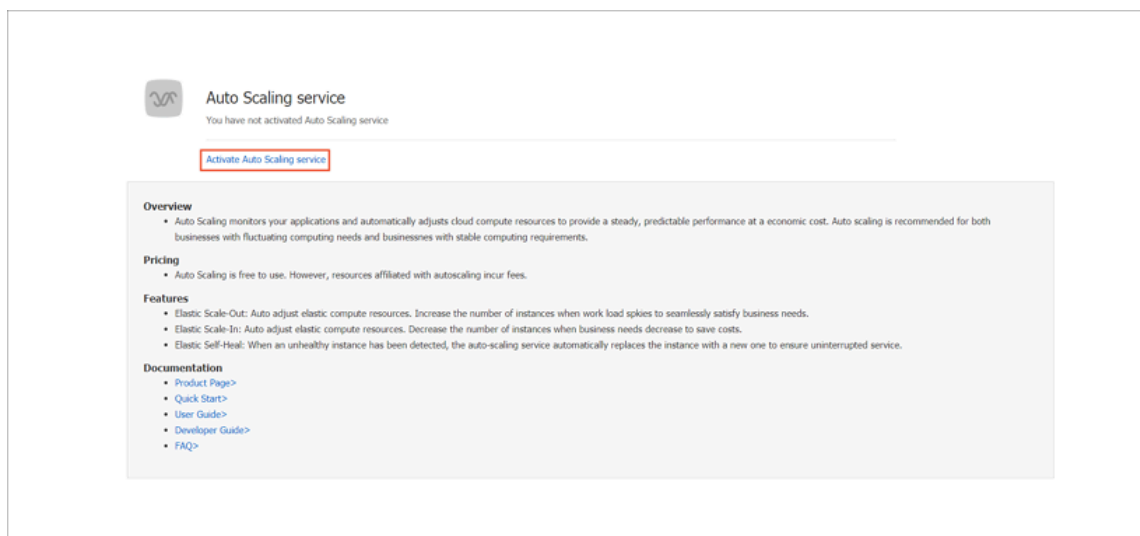
Grant required permissions for the Auto Scaling service and the cluster

- Activate the Auto Scaling service

1. Click ESS in the displayed dialog box.



2. Click Activate Auto Scaling service.



3. Read and confirm that you agree to the conditions by selecting the I agree with Auto Scaling Service Agreement of Service check box, and then click Enable Now.

Enable Service

Auto Scaling Service

Basic Configuration

Product


Auto Scaling Service

☒ I agree with [Auto Scaling Service Agreement of Service](#)

Enable Now

4. Click Console.

Confirm Order

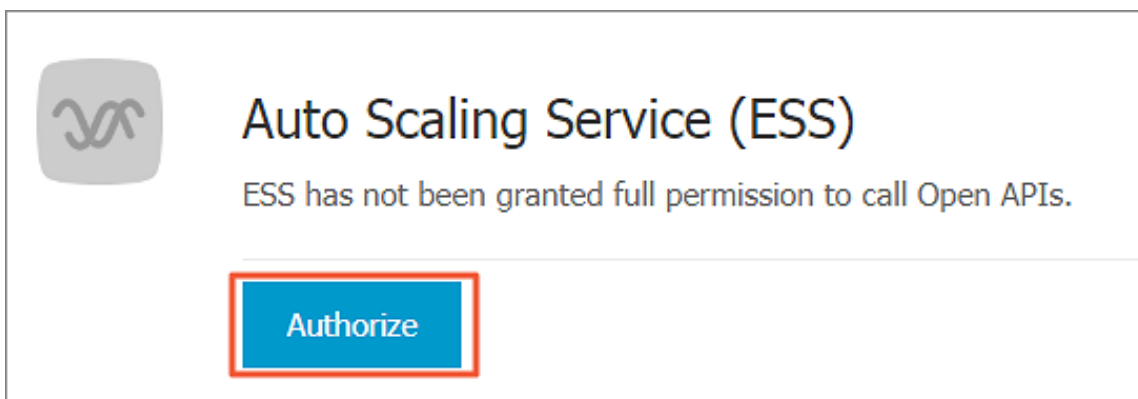
 Order completed

The service you ordered is being opened. Please wait for a few minutes.

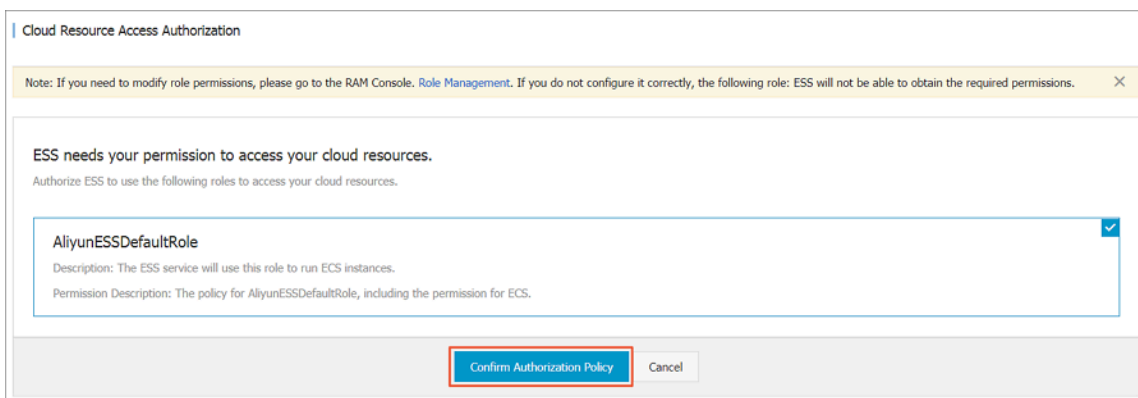
Console

Purchase history

5. Click Authorize.



6. Click Confirm Authorization Policy to grant ESS the permission to access your cloud resources.



Verify the result

If the page automatically redirects to the Auto Scaling console, the activation is successful.

· Add ESS authorization policies to the cluster

1. Click the Worker RAM role (`Kubernetes WorkerRole [xxx]`) in the following dialog box.



You need to use the primary account to log on to the console before performing this operation.

Note

Auto-scaling relies on the ESS service. Before enabling auto-scaling, you need to:

1

Enable the service and complete the default role authorization: [ESS](#)

2

Jump to RAM to add an ESS authorization policy to the current cluster: [View detailed steps](#)
[KubernetesWorkerRole](#)

Please confirm the above steps, otherwise the Auto-scaling will not be enabled.

Confirm

2. Click View Permissions on the right of the target authorization policy.

<	KubernetesWorkerRole-827d8853-3009-40c9-add1-3f378613c5d1	Edit Authorization Policy							
Role Details									
Role Authorization P...	<table><thead><tr><th>Authorization Policy Name</th><th>Description</th><th>Type</th><th>Actions</th></tr></thead><tbody><tr><td>kbsWorkerRolePolicy-827d8853-3009-40c9-add1-3f378613c5d1</td><td></td><td>Custom</td><td>View Permissions Revoke Authorization</td></tr></tbody></table>	Authorization Policy Name	Description	Type	Actions	kbsWorkerRolePolicy-827d8853-3009-40c9-add1-3f378613c5d1		Custom	View Permissions Revoke Authorization
Authorization Policy Name	Description	Type	Actions						
kbsWorkerRolePolicy-827d8853-3009-40c9-add1-3f378613c5d1		Custom	View Permissions Revoke Authorization						

3. In the upper-right corner of the page, click **Modify Authorization Policy**.

k8sWorkerRolePolicy-827d8853-3009-40c9-add1-3f378613c5d1 (Custom)

Authorization Policy...

Versions

References

Policy Details

Name	k8sWorkerRolePolicy-827d8853-3009-40c9-add1-3f378613c5d1	Type	Custom
------	----------------------------------------------------------	------	--------

Description

```

1 {
2   "Version": "1",
3   "Statement": [
4     {
5       "Action": [
6         "ecs:AttachDisk",
7         "ecs:DetachDisk",
8         "ecs:DescribeDisks",
9         "ecs>CreateDisk",
10        "ecs>CreateSnapshot",
11        "ecs>DeleteDisk",
12        "ecs>CreateNetworkInterface",
13        "ecs:DescribeNetworkInterfaces",
14        "ecs:AttachNetworkInterface",
15        "ecs:DetachNetworkInterface",
16        "ecs>DeleteNetworkInterface",

```

[More information about the authorization policy formulation.](#)

4. In the **Action** field of the Policy Document area, add the following policies:

```
" ess : Describe *",
" ess : CreateScalingRule ",
" ess : ModifyScalingGroup ",
```

```
" ess : RemoveInstances ",
" ess : ExecuteScalingRule ",
" ess : ModifyScalingRule ",
" ess : DeleteScalingRule ",
" ecs : DescribeInstancesType s ",
" ess : DetachInstances "
```

**Note:**

You must add a comma (,) to the end of the last line in the **Action** field before adding these policies.

5. Click Modify Authorization.

Set cluster auto scaling parameters

1. On the Automatic Scaling page, set the following parameters.


Configuration	Description
Cluster	Target cluster name.
Scale-in Threshold	The ratio of the amount of resources requested by the cluster workload to the amount of cluster resources. When the amount of resources requested by the cluster workload is less than or equal to the threshold, the system automatically scales in the cluster. The default value is 50%.
Defer Scale-in For	The number of minutes for which the system must wait to automatically scale in the cluster after the scale-in threshold is reached. The default value is 10 minutes.
Cool-Down Time	The period (in minutes) during which the system does not automatically scale in or scale out a cluster after the number of cluster nodes increases or decreases. The default is 10 minutes.


2. Click **Create** on the right of the target type of resource (which can be CPU or GPU) that you want to autoscale.

On the **Scaling Group Configuration** page, set the following parameters to create a scaling group:

Configuration	Description
Region	The region to which the scaling group is deployed . You must ensure that the scaling group and the cluster where it is located share the same region. This region cannot be modified.
Zone	The zone where the scaling group is created.
VPC	The network where the scaling group is created. You must ensure that the scaling group and the cluster where it is located are in the same region.

Set worker nodes.

Configuration	Description
Instance Type	Set the specifications of instances in the scaling group .
System Disk	Set the system disks of the scaling group.
Attach Data Disk	Mount a data disk when you create a scaling group. By default, no data disk is mounted.
Instance Quantity	Set the number of instances in the scaling group. <div>  Note: <ul style="list-style-type: none"> • The number does not include the existing instances in the cluster. • By default, this parameter value is 0, and the cluster adds instances to the scaling group and the Kubernetes cluster where the scaling group is located when this parameter exceeds 0. </div>

Configuration	Description
Key Pair	<p>Set the key pair used to log on to the node added through autoscaling. You can create a new key pair in the Elastic Compute Service (ECS) console.</p> <div> Note: Only key pair logon is supported.</div>
RDS Whitelist	Set the Relational Database Service (RDS) instances that can be accessed by the node added through autoscaling.

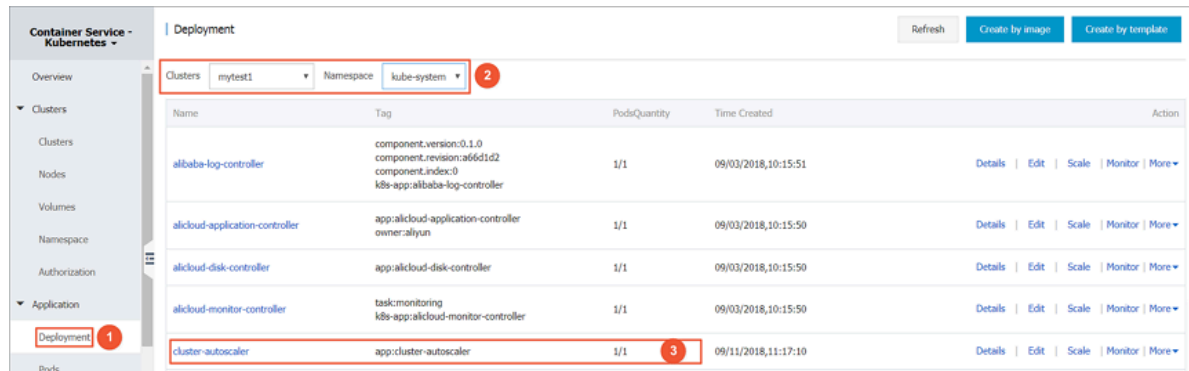
3. Click OK.

Verify the result

- You can directly verify that a scaling group under CPU is displayed on the Auto Scaling page.

Auto-scaling Back to Cluster List				
Overview				
Cluster ID: c68b3c5a4273c4fba8151f06276bcc224		Cluster Name: mytest1		
Shrinkage Threshold: 50%		Shrinkage Trigger Delay: 10Min.		
Scaling Group	Status	Instance Type	Total Instance Number	Min Instance Number
CPU asg-2ze9npi7u9efzddlmrzc	Active	ecs.sn1ne.xlarge	0	0
GPU				

- To verify the created autoscaling component, follow these steps:
 1. In the left-side navigation pane, choose Application > Deployment.
 2. Select the target cluster and the kube-system namespace to view the created component named cluster-autoscaler.



Troubleshooting

- If the cluster autoscaler cannot add nodes to a pod that requested more resources, you can perform the following checks:
 - Make sure that the amount of resources provided by the ECS instances that you set for the autoscaling group is greater than the amount of resources requested by the pod.
 - Make sure that you have granted the required permissions by following the preceding steps. You must grant the required permissions for each target cluster .
 - Make sure that the target Kubernetes cluster are connected to the Internet. The cluster autoscaler calls an API action from Alibaba Cloud, therefore you must ensure that the cluster nodes can be accessed through the Internet.

- If the cluster autoscaler cannot delete nodes from the autoscaling group, you can perform the following checks:
 - Make sure that the resource request threshold of pods on all nodes is not greater than that of used to scale in the cluster.
 - Make sure that no node runs the pods that belong to the `kube-system` namespace.
 - Make sure that no node runs the pod for which any constrained scheduling policies are set. This is because a constrained scheduling policy can limit a pod to a fixed node.
 - Make sure that no pod contains a `PodDisruptionBudget` object that has reached the minimum value allowed. For more information, see [How do Disruption Budgets work](#). For more information, see [How do Disruption Budgets work](#).

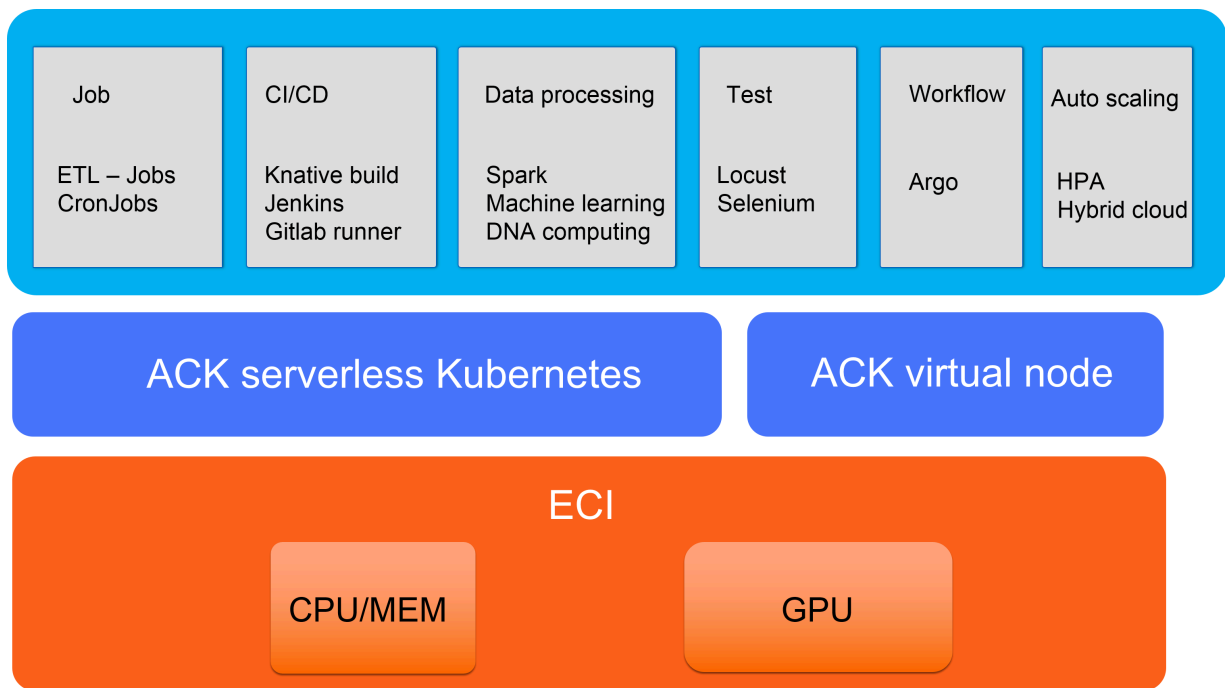
For more information, see [Cluster autoscaler](#).

1.21.2 Deploy a virtual node

This topic describes how to deploy a virtual node for a Kubernetes cluster by using the virtual node addon.

Background information

Virtual nodes are implemented by using Virtual Kubelet. Virtual nodes connect Kubernetes with Alibaba Cloud Elastic Container Instances (ECI), and provide Kubernetes clusters with a high level of elasticity. For more information about Virtual Kubelet, see [How does Virtual Kubelet work?](#)



Virtual nodes can help to enhance the elasticity of Kubernetes clusters. Specifically, virtual nodes based on Alibaba Cloud ECI can enhance the elasticity of clusters because they support GPU container instances, EIP addresses, and container instances of high specifications. Given the expansion capability provided by virtual nodes, you can easily manage multiple workloads for computing in a Kubernetes cluster.

In a cluster that contains virtual nodes, the pod in a physical node communicates with the ECI pod in the corresponding virtual node.



Note:

- The ECI pod in a virtual node is charged according to the specific amount of resources that you use. For information about ECI billing rules, see [Billing overview](#).
- ECS instances of the specifications range of 0.25 vCPU to 64 vCPU are supported. For more information, see [Limits](#).

Prerequisites

A managed Kubernetes cluster is created. For more information, see [Create a managed Kubernetes cluster](#).

Create a virtual node for a cluster Work node

1. Log on to the [Container Service console](#).

2. In the left-side navigation pane under Container Service-Kubernetes, choose Store > App Catalog. Then, click ack-virtual-node.
3. Click the Values tab, and then set these parameters.

- `ECI_VSWITCH_ID` : Set the VSwitch ID.

To obtain the ID of the VSwitch that is associated with a Worker node, follow these steps:

- a. In the left-side navigation pane under Container Service-Kubernetes, choose Clusters > Nodes.
- b. Select the target cluster, and then click a Worker node.

In the Configuration Information area, the VSwitch ID is displayed.

- `ECI_SECURITY_GROUP_ID` : Set the security group ID.

To obtain the ID of the security group that is associated with a Worker node, follow these steps:

- a. In the left-side navigation pane under Container Service-Kubernetes, choose Clusters > Nodes.
- b. Select the target cluster, and then click a Worker node.
- c. In the left-side navigation pane, click Security Groups.

On the Security Groups page, the security group ID is displayed.

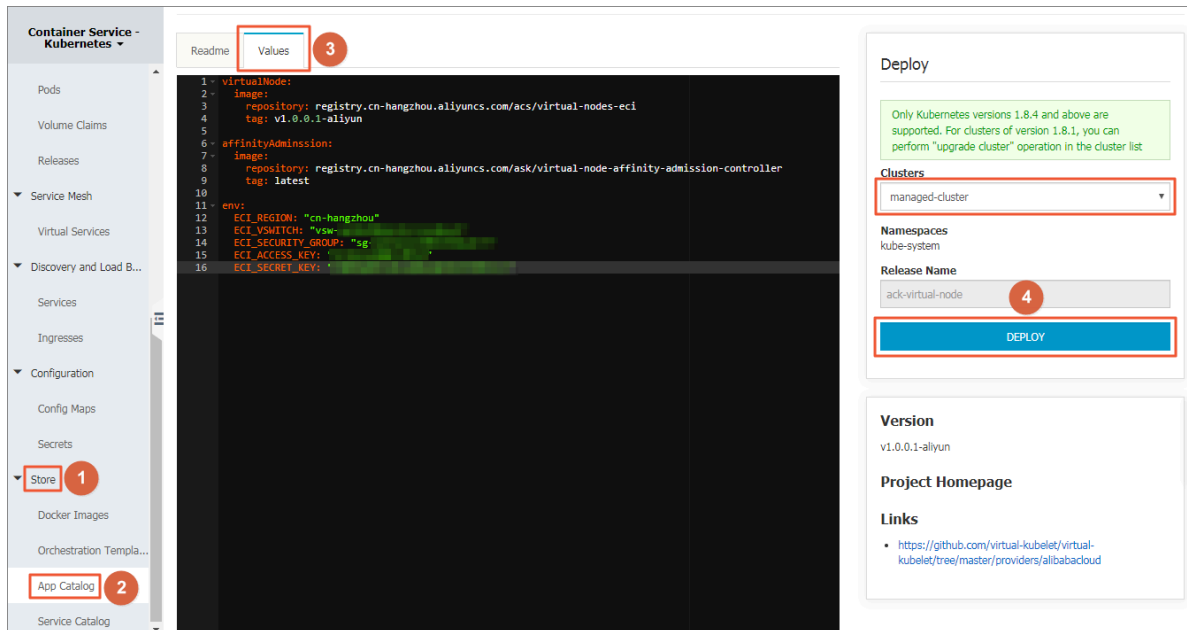
- `ECI_ACCESS_KEY_ID` : Set the AccessKey.
- `ECI_SECRET_KEY` : Set the SecretKey.

4. In the Deploy area on the right of the page, select the target cluster, and then click DEPLOY.

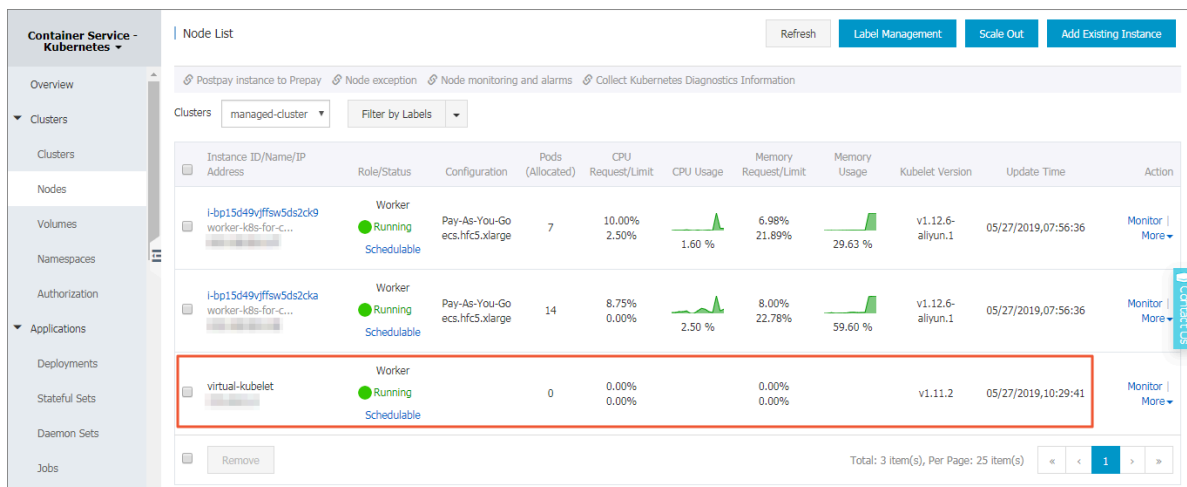


Note:

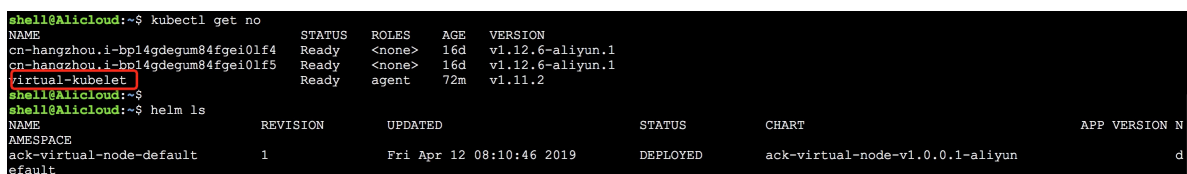
By default, the namespace is set to `kube-system`, and the release name is set to `ack-virtual-node`.



5. In the left-side navigation pane under Container Service-Kubernetes, choose **Clusters > Nodes** to verify that a node named `virtual-kubelet` is displayed.



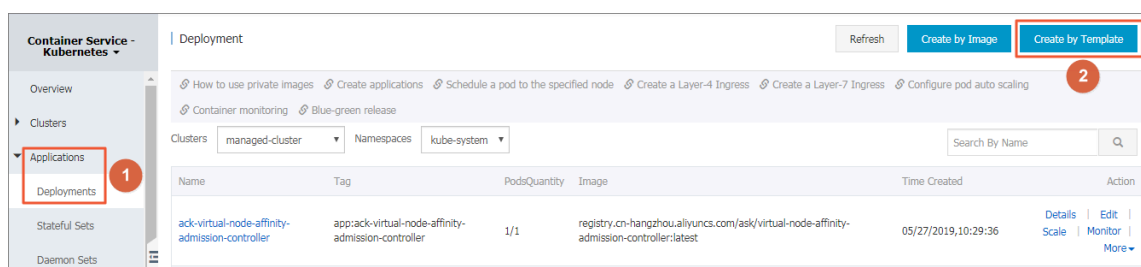
You can use `kubectl` commands to view cluster nodes, and the status of Helm deployment. You can also use Helm to upgrade and manage the deployed `ack-virtual-node`. For more information, see [Use kubectl commands in Cloud Shell to manage a Kubernetes cluster](#).



Create a pod in the virtual node

If a Kubernetes cluster contains a virtual node, you can create a pod in the virtual node. Then, the virtual kubelet will create a corresponding ECI pod. You can use one of the following three methods to create a pod for the virtual node:

- Set `nodeSelect` or `and` `toleration s` to create a pod.
- Log on to the [Container Service console](#).
 - In the left-side navigation pane under Container Service-Kubernetes, choose Applications > Deployments. Then, in the upper-right corner, click Create by Template.



- Select the target cluster and namespace, select an example template or customize a template, and then click DEPLOY.

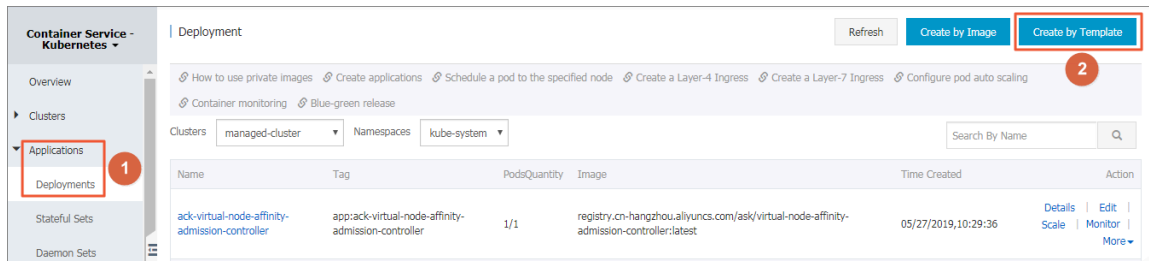
You can use the following template to customize a pod:

```
apiVersion : v1
kind : Pod
metadata :
  name : nginx
spec :
  containers :
  - image : nginx
    imagePullPolicy : Always
    name : nginx
    nodeSelector :
      type : virtual - kubelet
    tolerations :
    - key : virtual - kubelet . io / provider
```

operator : Exists

- Set `nodeName` to create a pod.

1. In the left-side navigation pane under Container Service-Kubernetes, choose Applications > Deployments. Then, in the upper-right corner, click Create by Template.



2. Select the target cluster and namespace, select an example template or customize a template, and then click DEPLOY.

You can use the following template to customize a pod:

```
apiVersion : v1
kind : Pod
metadata :
  name : nginx
spec :
  containers :
  - image : nginx
    imagePullPolicy : Always
    name : nginx
    nodeName : virtual - kubelet
```

- Set a namespace tag to create a pod.

1. Use kubectl on Cloud Shell to connect to the target Kubernetes cluster. For more information, see [Use kubectl commands in Cloud Shell to manage a Kubernetes cluster](#).
2. Run the following commands to create a pod:

```
kubectl create ns vk
kubectl label namespace vk virtual-node-affinity-injection=enabled
kubectl -n vk run nginx --image nginx
```

```
shell@aliyuncloud:~$ kubectl create ns vk
namespace/vk created
shell@aliyuncloud:~$ kubectl label namespace vk virtual-node-affinity-injection=enabled
namespace/vk labeled
shell@aliyuncloud:~$ kubectl -n vk run nginx --image nginx
kubectl run --generator=deployment/apps.v1 is DEPRECATED and will be removed in a future version. Use kubectl run --generator=run-pod/v1 or kubectl create instead.
deployment.apps/nginx created
shell@aliyuncloud:~$
```

In the left-side navigation pane under Container Service-Kubernetes, choose Applications > Pods to verify that the pod is created.

Container Service - Kubernetes Overview Clusters Applications Deployments Stateful Sets Daemon Sets Jobs Cron Jobs Pods	logtail-ds-8bqfc logtail-0.16.19.0-5be7ad3-aliyun									
	Running	0			05/27/2019,07:58:17		0.004		24.645 Mi	Details More
	logtail-ds-9l77c logtail-0.16.19.0-5be7ad3-aliyun									
	Running	0			05/27/2019,07:58:26		0.004		24.176 Mi	Details More
	metrics-server-748b5b5b-kdr4j metrics-server:v0.2.1-9dd9511-aliyun									
	Running	0			05/27/2019,07:58:13		0.002		25.563 Mi	Details More
	nginx nginx									
	Running	0		virtual-kubelet	05/27/2019,10:37:49	0	0			Details More
	nginx-ingress-controller-5f65d97854-sn26f aliyun-ingress-controller:v0.22.0.5-552e0db-aliyun									
	Running	0			05/27/2019,07:58:13		0.003		180.445 Mi	Details More

2 Serverless Kubernetes cluster

2.1 Overview

Alibaba Cloud Serverless Kubernetes allows you to quickly create Kubernetes container applications without having to manage and maintain clusters and servers. The Pay-As-You-Go billing method is applied, which is based on the amount of CPU and memory resources used by applications. With Serverless Kubernetes, you can focus on designing and building applications, rather than managing the infrastructure on which your applications run. Serverless Kubernetes is based on the Alibaba Cloud elastic computing architecture and is fully compatible with the Kubernetes API, combining the security, elasticity, and Kubernetes ecosystem of virtualized resources.

Regions available for public beta

Currently, Serverless Kubernetes clusters of Alibaba Cloud Container Service are in public beta stage. Now, only several regions are available for public beta. Other regions are going to be available soon.

Limits

Pods in a serverless Kubernetes cluster are created based on Elastic Container Instance (ECI). For more information about the pod specifications and pod usage limits, see [Limits](#).

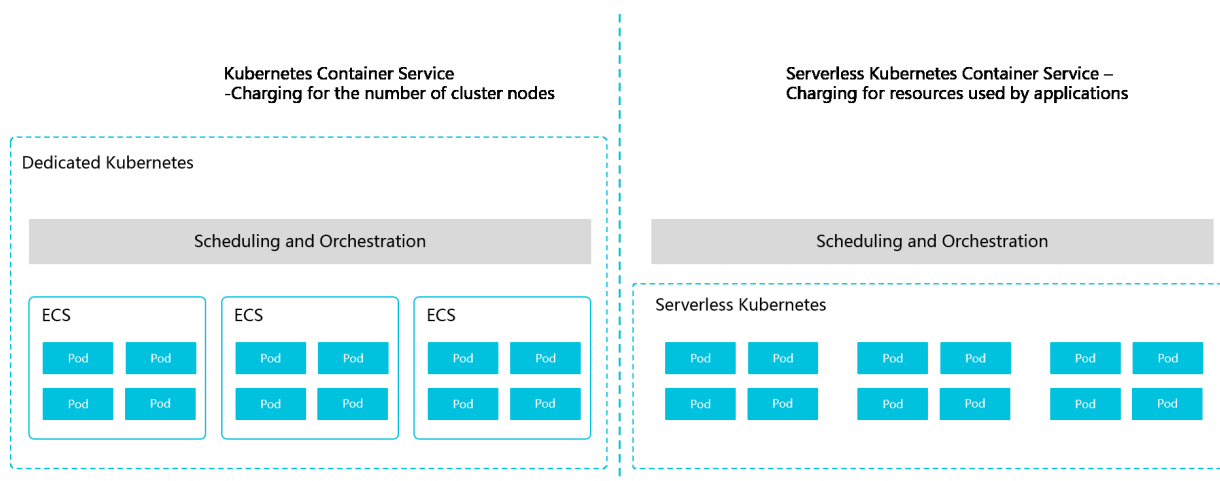
Pricing

Serverless Kubernetes clusters are free of charge.

For more information about ECI pricing, see [Pricing](#).

Each type of resource (such as an SLB instance and a private zone) used in a serverless Kubernetes cluster are charged according to the price specified by the corresponding product.

Comparison with Container Service



2.2 Kubernetes supported functions

API version

Kubernetes 1.9 API is supported.

Application load

- Deployment, StatefulSet, Job/CronJob, Bare Pod are supported.
- DaemonSet is not supported.

Pod definition

Supports starting multiple containers, setting environment variables, RestartPolicy, health check commands, and mounting volumes.

Load balancing

- Supports creation of Load Balancer type applications.
- Ingress is supported.
- NodePort type is not supported.

Configuration

Secret and ConfigMap are supported.

Storage

- EmptyDir and NFS volume types are supported.
- PersistentVolume and PersistentVolumeClaim are not supported.

Namespace

Only the default namespace can be viewed, and no namespaces can be added.

Node

Node information of Kubernetes cannot be viewed.

Events

Default namespace events can be viewed.

Containers logs

View the container logs in real time by using `kubectl logs`.

Container exec/attach

Enter the container to run the commands by using `kubectl exec`.

2.3 Cluster management

2.3.1 Create a serverless Kubernetes cluster

You can create a serverless Kubernetes cluster quickly and easily in the Container Service console.

Prerequisites

Log on to the [Container Service console](#), and the [RAM console](#) to activate the corresponding service.

Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Clusters in the left-side navigation pane to enter the Cluster List page.
3. Click Create a Serverless Kubernetes cluster in the upper-right corner.

4. Enter the cluster name.

Cluster name can only contain 1 - 63 characters, including numbers, Chinese characters, English characters, and hyphens (-).

5. Select the region and availability zone where the cluster resides.



Note:

Serverless Kubernetes cluster is currently in beta stage, and only East China 2 (Shanghai) region is available.

6. Set the cluster network type.

- Network type: Kubernetes clusters only support the VPC network type.
- VPC: Auto Create and Use existing options are supported.



Note:

- Auto Create: System automatically creates VPC, NAT gateway and configures SNAT rules in VPC when the cluster is created.
- Use existing: Select the target VPC and VSwitch from the list of existing VPCs. To access the public network, for example, to download container images, configure the NAT gateway. We recommend that container images be uploaded to the Alibaba Cloud image service in the region where the cluster resides, and pull the image by the VPC address.

7. Check whether to enable the PrivateZone-based service discovery feature, which allows you to access the service by the domain name within the cluster VPC.



Note:

Before using this feature, confirm that you have enabled the PrivateZone service, see [#unique_219](#).

8. Check the Serverless Kubernetes service protocol.

9. Click Create cluster to start the deployment.

What's next

After the cluster is created, you can view the cluster in the Kubernetes cluster list in the Container Service console.

You can also click Manage at the right of the cluster to view the basic and connection information of this cluster.

2.3.2 Connect to a Kubernetes cluster by using kubectl

To connect to a Kubernetes cluster from a client computer, use the Kubernetes command line client kubectl.

Procedure

1. Download the latest kubectl client from the [Kubernetes version](#) page.
2. Install and set the kubectl client.

For more information, see [Install and set kubectl](#).

3. Configure the cluster credentials.

You can view the cluster credentials on the cluster information page.

- a) Log on to the [Container Service console](#).
- b) Under Kubernetes, click Clusters in the left-side navigation pane.
- c) Click Manage at the right of the cluster.
- d) In the Connection Information section, view the master node SSH IP address.
- e) Copy the cluster credentials to a local file, and you can create and save the cluster credentials to `$ HOME /. kube / config` (location where kubectl credentials are to be stored). You can also name a new file, such as `/ tmp / kubeconfig`, and run the `export KUBECONFIG =/ tmp / kubeconfig` command.
- f) After the preceding operation is performed, you can confirm the cluster connectivity by running the following command.

```
# kubectl get pod
No resources found .
```

What's next

After the configuration is complete, you can use kubectl to access the Kubernetes cluster from a local computer.

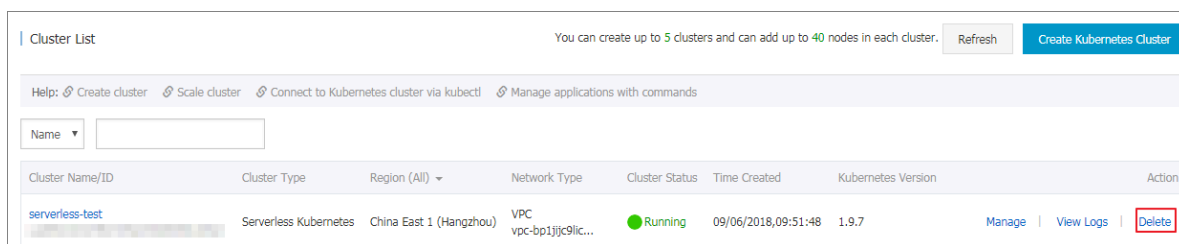
2.3.3 Delete a cluster

You can delete clusters that are no longer in use in the Container Service console.

Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Clusters in the left-side navigation pane.

3. Select the target cluster and click Delete on the right.



4. In the dialog box, click OK.

2.4 Application management

2.4.1 Manage applications by using commands

You can create applications or view containers in applications by using commands.

Prerequisites

Before using commands, configure [Connect to a Kubernetes cluster by using kubectl](#) first.

Create an application by using commands

Execute the following statements to run a simple container (a Nginx Web server in this example).

```
root @ master # kubectl run nginx -- image = registry . cn -
hangzhou . aliyuncs . com / spacexnice / netdia : latest
```

This command creates a service portal for this container. Specify `-- type =`

`LoadBalancer` and Alibaba Cloud Server Load Balancer route will be created to the Nginx container.

```
root @ master # kubectl expose deployment nginx -- port = 80
-- target - port = 80 -- type = LoadBalancer
```

View containers by using commands

Run the following command to list all the running containers in the default namespaces.

```
root @ master # kubectl get pods
NAME                                READY    STATUS
RESTARTS    AGE
```

nginx - 2721357637 - dvwq3	1 / 1	Running
1 9h		

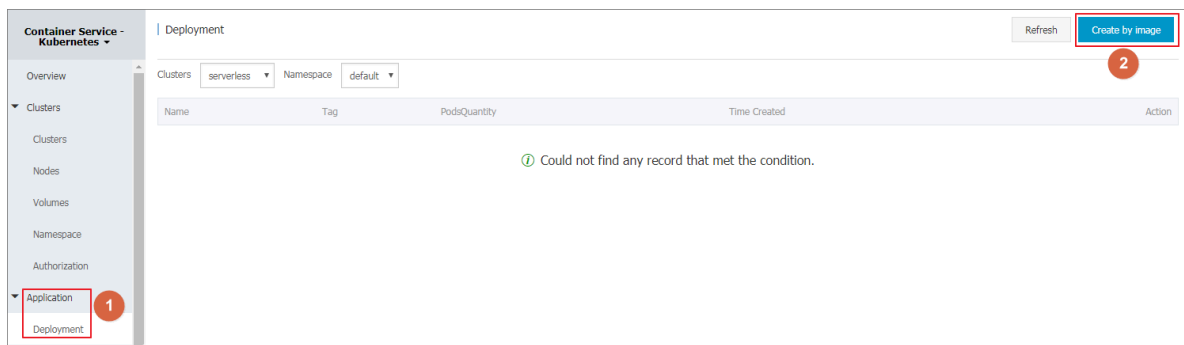
2.4.2 Create an application by using an image

Prerequisites

Create a Serverless Kubernetes cluster. For more information, see [Create a serverless Kubernetes cluster](#).

Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click **Application > Deployment** in the left-side navigation pane to enter the Deployment List page.
3. Click **Create by image** in the upper-right corner.



4. Set Application name, Cluster deployment, and Namespace. Then click **Next** to enter the application configuration page.

If the namespace is not set, the default namespace is used.

5. Select the image you want to use and the version of the image.

- **Image name:** In the displayed dialog box click **Select Image**, and then click **OK**. You can also enter the private registry. In the format of `domainname / namespace / imagename : tag`. In this example, the image name is `nginx`.
- **Image version:** Click **Select image version** to select the version. If not specified, the latest version is used by default.

6. Configure the number of containers (Scale).

This example is a single container pod, and if multiple containers are specified, the same number of pods is started.

7. Configure resource limits and required resources for the container.

Serverless Kubernetes is currently in beta stage and only supports the 2C4G specification.

- **Resource limits:** You can specify the maximum amount of resources that the application can use, including CPU and memory to prevent excessive use of resources.
- **Required resources:** The amount of resources reserved for the application, including CPU and memory. That is, container monopolizes the resource, so as to prevent other services or processes from competing for the resource due to insufficiency, resulting the application to be unavailable.

CPU is measured in millicores (one thousandth of one core). Memory is measured in bytes, which can be GB, MB, or KB.

8. Configure the environment variable.

You can configure the environment variable for the pod in the format of key-value pairs to add the environment label or pass the configurations for the pod. For more information, see [Pod variable](#).

9. Configure the container.

You can configure the Command and Arguments for the container running in the pod.

Command and Args: If not configured, the default settings for the image is used. If configured, the default settings are overwritten. If only arguments are configured, when the container starts, the default command executes the new arguments.

Command and arguments cannot be modified after pod is created.

10. After the application configuration is complete, click Next to enter the Access Settings page to set up a service that binds the backend pod.

You can select not to create a service, or select a service type. Currently, only load balancing type is supported.

- **Load Balancing:** Load Balancer is a load balancing service provided by Alibaba Cloud, public network access or intranet access can be used.
- **Name:** By default, a service name with the application name suffix `svc` is generated, in this example `serverless-app-svc`. Name of the service can be modified.
- **Port mapping:** Specify the port mapping between service and container, and select TCP or UDP as the protocol.

Create Service

Name:

serverless-app-svc

Type:

Server Load Balancer

public

Port Mapping:

+ Add

service port	Container Port	Protocol	
80	80	TCP	-

annotation:

+ Add Annotations for load balancer

Tag:

+ Add

Create

Cancel

11. After the access configuration is complete, click Create.

What's next

After the creation is successful, go to the creation completion page. The objects contained in the application are displayed. You can click View to view the deployment list.

You can view the new serverless-app-svc under the deployment list.

Clusters

serverless


 Namespace

default

Name	Tag	PodsQuantity	Time Created	Action
serverless-app-deployment	app:serverless-app	0/1	08/13/2018,11:55:14	Details Edit Monitor More

Click Application > Service in the left-side navigation pane to see the new service serverless-app-svc under the list of services.

Clusters serverless Namespace default

Name	Type	Time Created	ClustersIP	InternalEndpoint	ExternalEndpoint	Action
serverless-app-svc	LoadBalancer	08/13/2018,11:55:14		serverless-app-svc:80 TCP serverless-app-svc:31530 TCP	 :80	Details Update View YAML Delete

Access the external endpoint in the browser to access the Nginx welcome page.

2.4.3 Create a service

Kubernetes service, which is generally called a microservice, is an abstraction which defines a logical set of pods and a policy by which to access them. This set of pods can be accessed by the service, typically by using the Label Selector.

Kubernetes pods are created and deleted in a short time even if they have their own IP addresses. Therefore, using pods directly to provide services externally is not a solution of high availability. The service abstraction decouples the relationship between the frontend and the backend. Therefore, the loose-coupling microservice allows the frontend to not care about the implementations of the backend.

For more information, see [Kubernetes service](#).

Prerequisites

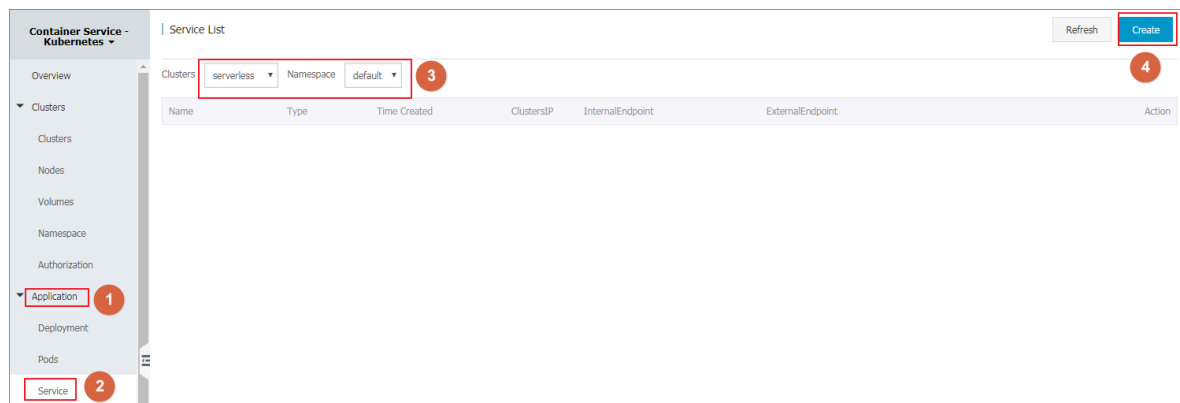
You have successfully created a Serverless Kubernetes cluster, see [Create a serverless Kubernetes cluster](#).

Step 1. Create a deployment

Create a deployment by using the image, in this example create serverless-app-deployment. For more information, see [Create an application by using an image](#).

Step 2. Create a service

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Application > Service in the left-side navigation pane to enter the Service List page.
3. Select the target cluster and namespace and click Create in the upper-right corner.



4. Complete the configurations in the displayed Create Service dialog box.

Create Service

Name:

serverless-service

Type:

Server Load Balancer

public

Related deployment:

serverless-app-deploymen

Port Mapping:

+ Add

service port	Container Port	Protocol	
80	8080	TCP	-

annotation:

+ Add Annotations for load balancer

Name	Value	
service.beta.kubernetes.io	20	-

Tag:

+ Add

Name	Value	
app	nginx	-

Create

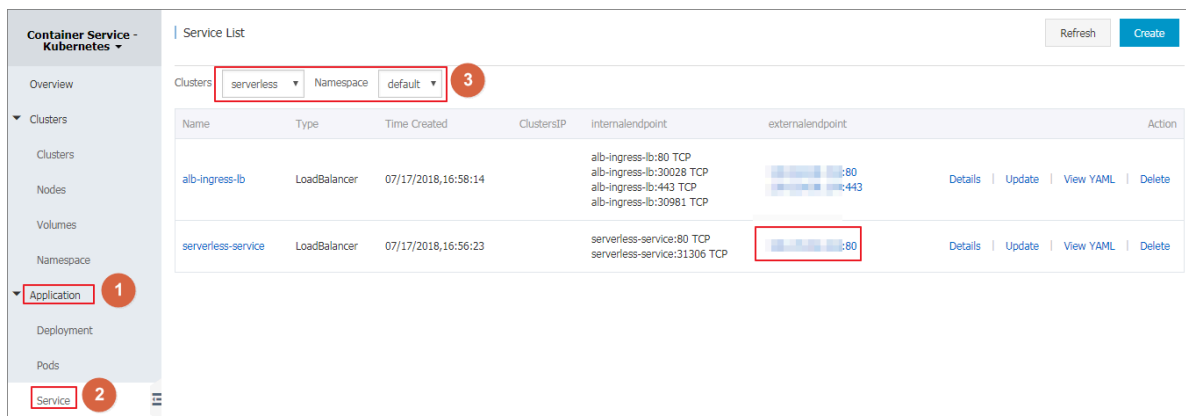
Cancel

- **Name:** Enter the name of the service, in this example serverless-service.
- **Type:** Select the service type, namely, the access method of the service.
Currently, only load balancing type is supported. Load Balancer is the load

balancing service provided by Alibaba Cloud, public network access or intranet access can be used.

- **Related deployment:** Select the backend object to bind with this service, in this example, select nginx-deployment-basic. The corresponding Endpoints object is not created if no deployment is selected here. You can manually map the service to your own endpoints. For more information, see [services-without-selectors](#).
- **Port Mapping:** Add the service port and container port. The container port must be the same as the one exposed in the backend pod.
- **Annotation:** Add an annotation to the service and configure load balancing parameters such as `service.beta.kubernetes.io/alibaba-cloud-loadbalancer-bandwidth: 20` indicating that the bandwidth of the service is set to 20 Mbit/s to control the traffic of the service. For more information, see [Server Load Balancer](#).
- **Label:** You can add a label to the service to identify the service.

5. Click Create, and the serverless-service appears in the list of services.



6. You can view the basic information of the service and access the external endpoint of serverless-service in your browser.



Then, you have created a service that is related to a backend deployment and accessed the Nginx welcome page successfully.

2.4.4 Delete a service

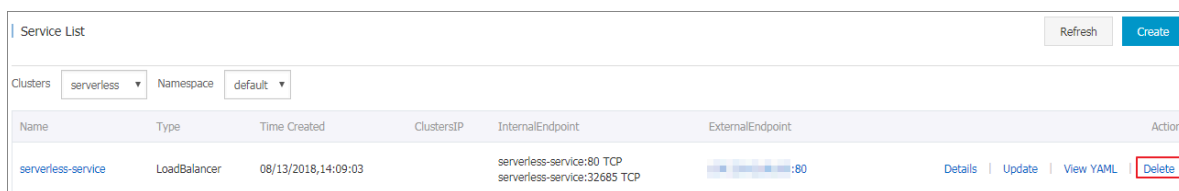
You can quickly delete a service in the Container Service console.

Prerequisites

- You have successfully created a Serverless Kubernetes cluster, see [Create a serverless Kubernetes cluster](#).
- You have successfully created a service, see [Create a service](#).

Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Application > Service in the left-side navigation pane to enter the Service List page.
3. Select the cluster and namespace, select the target service (serverless-service in this example), and click Delete on the right.



Name	Type	Time Created	ClustersIP	InternalEndpoint	ExternalEndpoint	Action
serverless-service	LoadBalancer	08/13/2018,14:09:03		serverless-service:80 TCP serverless-service:32685 TCP		Details Update View YAML Delete

4. In the displayed window, click OK to confirm the deletion, and the service disappears from the list of services.

2.4.5 View pods

You can view the pods of the Serverless Kubernetes cluster in the Container Service console.

Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Application > Pods to enter the Pods page.
3. Select the target cluster and namespace, the target pod, and click Details on the right.



Note:

You can update or delete a pod. For pods created by using deployments, we recommend that you manage these pods by using deployments.

Pods								Refresh
Clusters		serverless	Namespace		default			
Name	Status	Pod IP	Node	Time Created	CPU	Memory		
serverless	Pending			08/13/2018,14:08:11	0	0	Detail	More

4. View the pod details.

Pods - server

Refresh

Overview

Name : server

Status : Pending

Node : viki

Tag : app: serverless-app pod-template-hash: 2803040087

Namespace : default

Time Created : 08/13/2018,14:08:11

Pod IP :

Container

Events

Created by

Init Containers

Volumes

Logs

Name

serverless-app

Image

nginx:latest

Port

TCP 80

2.4.6 View services

If the external service is configured when you create the application, in addition to running containers, system creates the external services for pre-assigning the Server Load Balancer to bring traffic to the containers in the cluster.

Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Application > Service in the left-side navigation pane to enter the Service List page.
3. You can view the deployed services by selecting the required clusters and namespaces.

You can view information, such as the name, type, creation time, cluster IP, and external endpoints of the service. In this example, view the external endpoint (IP address) assigned to the service.

Service List							Refresh	Create
Clusters		serverless	Namespace		default			
Name	Type	Time Created	ClusterIP	InternalEndpoint	ExternalEndpoint			
serverless-service	LoadBalancer	08/13/2018,14:09:03		serverless-service:80 TCP serverless-service:32685 TCP		Details	Update	View YAML

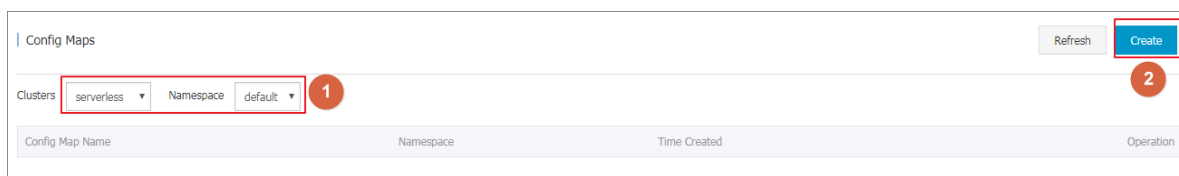
2.5 Config Map

2.5.1 Create a ConfigMap

This topic describes how to create a ConfigMap by using the Container Service console.

Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Application > Config Maps in the left-side navigation pane.
3. On the Config Maps page, select the cluster and namespace, and click Create.



4. Complete the configuration and then click OK.
 - **Namespace:** Select the namespace to which the ConfigMap belongs. ConfigMap is a Kubernetes resource object that must be applied to the namespace.
 - **Config Map Name:** Enter the ConfigMap name, which can contain lowercase letters, numbers, hyphens (-), and periods (.). The name cannot be empty.

Other resource objects must reference the ConfigMap name to obtain the configuration information.

- **Configuration:** Enter the Variable Name and the Variable Value. Then, click Add on the right. You can also click Edit, complete the configuration in the displayed dialog box, and click OK.

In this example, configure the variables enemies and lives to pass the parameters aliens and 3 respectively.

The screenshot shows a configuration dialog for a ConfigMap named 'test-config'. The 'Configuration' section is a table with two columns: 'Name' and 'Value'. It contains two entries: 'enemies' with value 'aliens' and 'lives' with value '3'. The 'OK' button is highlighted in blue.

Name	Value
enemies	aliens
lives	3

5. You can view the ConfigMap test-config on the Config Maps page after clicking OK.

2.5.2 Delete a config map

You can delete a config map that is no longer in use.

Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, Application > Config Maps in the left-side navigation pane.
3. Select the target cluster and config map, and click Delete on the right.

2.5.3 Modify a config map

You can modify the configurations of a config map.

Context

**Note:**

Modify a config map affects applications that use this file.

Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, Application > Config Maps in the left-side navigation pane.
3. Select the target cluster, namespace, and config map you want to modify. Then click Modify on the right.
4. Click Confirm in the displayed dialog box.
5. Modify the configurations.
 - Click Edit on the right of the configuration you want to modify. Update the configuration and then click Save.
 - You can also click Edit configuration file. Finish editing, and click OK.
6. After modifying the configurations, click OK.

2.6 Server Load Balancer management

2.6.1 Server Load Balancer

You can access services by using Alibaba Cloud Server Load Balancer.

Operate by using command line

1. Create an Nginx application by using command line.

```
root @ master # kubectl run nginx -- image = registry .
aliyuncs . com / acs / netdia : latest
root @ master # kubectl get po
NAME                                READY    STATUS
RESTARTS      AGE
nginx - 2721357637 - d ****          1 / 1    Running
1              6s
```

2. Create Alibaba Cloud Server Load Balancer service for the Nginx application and specify `type = LoadBalancer` to expose the Nginx service to the Internet.

```
root @ master # kubectl expose deployment nginx -- port =
80 -- target - port = 80 -- type = LoadBalancer
```

```

root @ master # kubectl get svc
NAME                                CLUSTER - IP                EXTERNAL - IP
PORT ( S )                          AGE
nginx                               172 .**.**.***               101 .**.**.***             80 : 3
****/ TCP                           4s

```

3. Visit `http://101.**.**.***` in the browser to access your Nginx service.

More information

Alibaba Cloud Server Load Balancer also supports parameter configurations such as health check, billing method, and load balancing. For more information, see [Server Load Balancer configuration parameters](#).

Annotations

Alibaba Cloud supports a lot of Server Load Balancer features by using annotations.

Use existing intranet Server Load Balancer instance

You must specify two annotations. Replace with your own Server Load Balancer instance ID.

```

apiVersion : v1
kind : Service
metadata :
  annotations :
    service.beta.kubernetes.io/alibaba-cloud-loadbalancer-address-type : intranet
    service.beta.kubernetes.io/alibaba-cloud-loadbalancer-id : your-loadbalancer-id
    service.beta.kubernetes.io/alibaba-cloud-loadbalancer-force-override-listeners : "true"
  labels :
    run : nginx
  name : nginx
  namespace : default
spec :
  ports :
    - name : web
      port : 80
      protocol : TCP
      targetPort : 80
  selector :
    run : nginx
  sessionAffinity : None
  type : LoadBalancer

```

Save the preceding contents as `slb.svc` and then run the command `kubectl apply -f slb.svc`.

Create an HTTPS type Server Load Balancer instance

Create a certificate in the Alibaba Cloud console and record the cert-id. Then, use the following annotation to create an HTTPS type Server Load Balancer instance.

```
apiVersion : v1
kind : Service
metadata :
  annotation s :
    service . beta . kubernetes . io / alicloud - loadbalanc er -
cert - id : your - cert - id
    service . beta . kubernetes . io / alicloud - loadbalanc er -
protocol - port : " https : 443 "
  labels :
    run : nginx
    name : nginx
    namespace : default
spec :
  ports :
    - name : web
      port : 443
      protocol : TCP
      targetPort : 443
  selector :
    run : nginx
  sessionAff inity : None
  type : LoadBalanc er
```



Note:

Annotations are case sensitive.

Annotation	Description	Default value
service.beta.kubernetes.io/alibabacloud-loadbalancer-protocol-port	Use commas (,) to separate multiple values. For example, https:443,http:80.	None
service.beta.kubernetes.io/alibabacloud-loadbalancer-addr-type	The value is Internet or intranet.	Internet
service.beta.kubernetes.io/alibabacloud-loadbalancer-slb-type	Server Load Balancer network type. The value is classic or VPC.	Classic
service.beta.kubernetes.io/alibabacloud-loadbalancer-charge-type	The value is paybytraffic or paybybandwidth.	paybybandwidth

Annotation	Description	Default value
service.beta.kubernetes.io/alibabacloud-loadbalancer-id	The Server Load Balancer instance ID. Specify an existing Server Load Balance with the loadbalancer-id, and the existing listener is overwritten. Server Load Balancer is not deleted when the service is deleted.	None
service.beta.kubernetes.io/alibabacloud-loadbalancer-backend-label	Use label to specify which nodes are mounted to the Server Load Balancer backend.	None
service.beta.kubernetes.io/alibabacloud-loadbalancer-region	The region in which Server Load Balancer resides.	None
service.beta.kubernetes.io/alibabacloud-loadbalancer-bandwidth	Server Load Balancer bandwidth.	50
service.beta.kubernetes.io/alibabacloud-loadbalancer-cert-id	ID of a certificate on Alibaba Cloud. You must have uploaded a certificate first.	“”
service.beta.kubernetes.io/alibabacloud-loadbalancer-health-check-flag	The value is on or off.	The default value is off. No need to modify the TCP parameters because TCP enables health check by default and you cannot configure it.
service.beta.kubernetes.io/alibabacloud-loadbalancer-health-check-type	See HealthCheck .	None
service.beta.kubernetes.io/alibabacloud-loadbalancer-health-check-uri	See HealthCheck .	None
service.beta.kubernetes.io/alibabacloud-loadbalancer-health-check-connect-port	See HealthCheck .	None

Annotation	Description	Default value
service.beta.kubernetes.io/alibabacloud-loadbalancer-healthy-threshold	See HealthCheck .	None
service.beta.kubernetes.io/alibabacloud-loadbalancer-unhealthy-threshold	See HealthCheck .	None
service.beta.kubernetes.io/alibabacloud-loadbalancer-health-check-interval	See HealthCheck .	None
service.beta.kubernetes.io/alibabacloud-loadbalancer-health-check-connect-timeout	See HealthCheck .	None
service.beta.kubernetes.io/alibabacloud-loadbalancer-health-check-timeout	See HealthCheck .	None

2.6.2 Use Ingress to provide Layer-7 service access

In the Alibaba Cloud Serverless Kubernetes cluster, Server Load Balancer provides Layer-4 service access. You can also use Layer-7 service access provided by Ingress. This document describes how to provide Layer-7 domain name service access in the Serverless Kubernetes cluster.

Prerequisites

- You have created a Serverless cluster. VPC cluster must be configured with a NAT gateway to access the external network and download the container image.
- You have connected to the cluster by using kubectl, see [Connect to a Kubernetes cluster by using kubectl](#).

Instructions

1. If Server Load Balancer is not specified, system automatically generates a public network Server Load Balancer instance.
2. The default front-end listening ports for SLB instances are 80 (HTTP Protocol) and 443 (HTTPS protocol).

3. By default, the HTTPS certificate of the SLB instance is initialized for the first created Ingress-configured TLS certificate. Otherwise, the system default certificate is initialized. You can modify it in the SLB console as needed.
4. When you specify to use an existing SLB instance, SLB instance specification must be of performance guarantee type (supports ENI). Also, make sure that ports 80 and 443 are not currently used by other services.

Annotations

Note	Description
service.beta.kubernetes.io/alibabacloud-loadbalancer-id	Specify an existing SLB ID.

Use the default generated SLB instance.

If an SLB instance is not specified, the system automatically generates a performance guaranteed public network SLB instance when the first Ingress is created.

1. Deploy test services

Deploy a coffee service and tea service first, the layout template is as follows:

```

apiVersion : extensions / v1beta1
kind : Deployment
metadata :
  name : coffee
spec :
  replicas : 2
  selector :
    matchLabels :
      app : coffee
  template :
    metadata :
      labels :
        app : coffee
    spec :
      containers :
        - name : coffee
          image : registry.cn-hangzhou.aliyuncs.com/acs-sample/nginxdemos : latest
          ports :
            - containerPort : 80
---
apiVersion : v1
kind : Service
metadata :
  name : coffee-svc
spec :
  ports :
    - port : 80
      targetPort : 80
      protocol : TCP
  selector :
    app : coffee

```

```

clusterIP : None
---
apiVersion : extensions / v1beta1
kind : Deployment
metadata :
  name : tea
spec :
  replicas : 1
  selector :
    matchLabels :
      app : tea
  template :
    metadata :
      labels :
        app : tea
    spec :
      containers :
        - name : tea
          image : registry.cn-hangzhou.aliyuncs.com/acs-sample/nginxdemos:latest
          ports :
            - containerPort : 80
---
apiVersion : v1
kind : Service
metadata :
  name : tea-svc
  labels :
spec :
  ports :
    - port : 80
      targetPort : 80
      protocol : TCP
  selector :
    app : tea
clusterIP : None

```

```

$ kubectl apply -f cafe-service.yaml #
Enter the template above in the cafe-service.yaml
file
deployment "coffee" created
service "coffee-svc" created
deployment "tea" created
service "tea-svc" created

# After deployment is complete

$ kubectl get svc,deploy
NAME                                TYPE                CLUSTER-IP          EXTERNAL-IP
PORT(S)      AGE
svc/coffee-svc    ClusterIP           <none>              <none>
80/TCP         1m
svc/tea-svc       ClusterIP           <none>              <none>
80/TCP         1m

NAME                                DESIRED    CURRENT    UP-TO-DATE
AVAILABLE      AGE
deploy/coffee    2          2          2          2
1m

```

deploy / tea 1m	1	1	1	1
--------------------	---	---	---	---

2. Configure Ingress

Configure the domain name and path exposed by coffee service and tea service by using Ingress:

```
apiVersion : extensions / v1beta1
kind : Ingress
metadata :
  name : cafe - ingress
spec :
  rules :
    # Configure Layer - 7 domain name
    - host : foo . bar . com
      http :
        paths :
          # Configure context path
          - path : / tea
            backend :
              serviceNam e : tea - svc
              servicePor t : 80
          # Configure context path
          - path : / coffee
            backend :
              serviceNam e : coffee - svc
              servicePor t : 80
```

```
apiVersion : extensions / v1beta1
kind : Ingress
metadata :
  name : cafe - ingress
spec :
  rules :
    # Configure Layer - 7 domain name
    - host : foo . bar . com
      http :
        paths :
          # Configure context path
          - path : / tea
            backend :
              serviceNam e : tea - svc
              servicePor t : 80
          # Configure context path
          - path : / coffee
            backend :
              serviceNam e : coffee - svc
              servicePor t : 80
```

```
$ kubectl apply -f cafe - ingress . yaml
ingress "cafe - ingress" created

# Address is automatically generated for SLB instance
IP after deployment is complete

$ kubectl get ing
NAME          HOSTS          ADDRESS          PORTS          AGE
```

```
cafe - ingress      foo . bar . com      139 .***.***.***      80
1m
```

3. Test service access



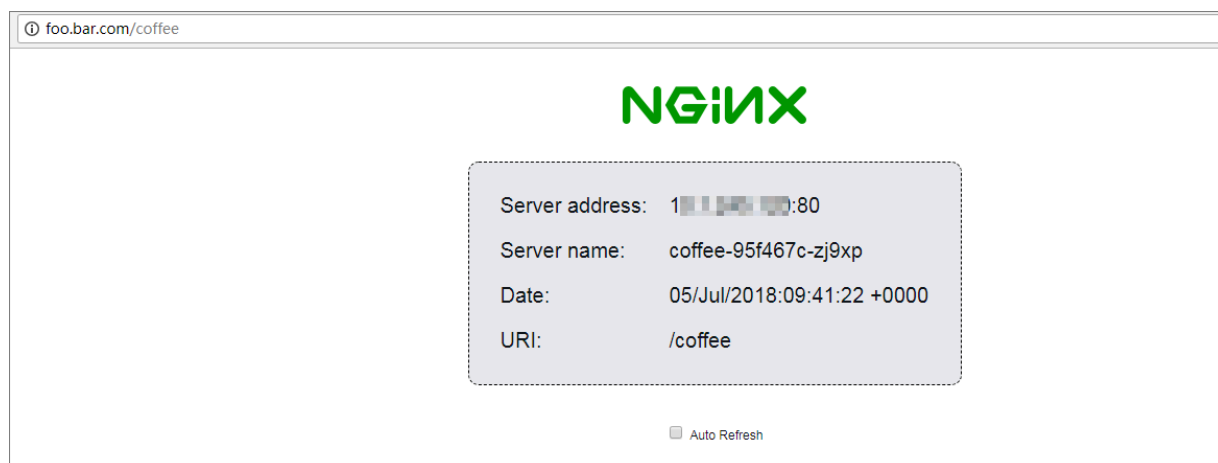
Note:

Currently, the domain name of the SLB instance IP must be resolved manually.

In this example, a DNS domain name resolution rule is added to `hosts` for testing service access. We recommend that you enter the domain name in your work environment.

```
139 .***.***.***      foo . bar . com
```

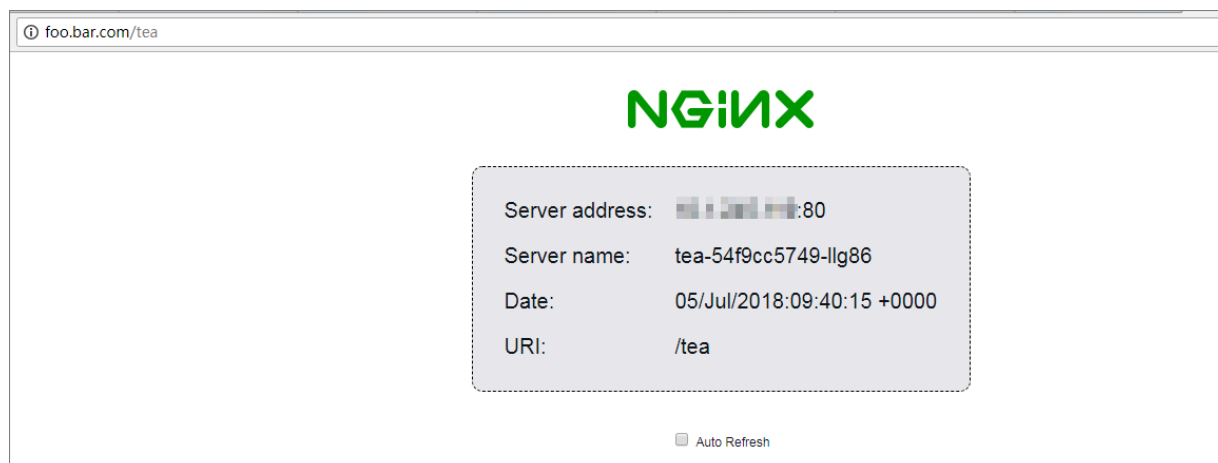
Test access to the coffee service by using browser.



Test access to the coffee service by using command line.

```
curl -H "Host : foo . bar . com " http :// 139 .***.***.***/
coffee
```

Test access to the tea service by using browser.



Test access to the tea service by using command line.

```
curl -H "Host : foo . bar . com " http :// 139 .***.***.***/ tea
```

Use the specified SLB instance

You can specify to use of an existing SLB instance by using the `service . beta . kubernetes . io / alicloud - loadbalanc er - id` annotation, but the SLB instance specification must be of performance guarantee type (supports ENI).



Note:

System automatically initializes ports 80 and 443 of the SLB instance, make sure that ports are not currently used by other services.

1. Deploy test services

Deploy a Tomcat test application first, the layout template is as follows:

```
apiVersion : extensions / v1beta1
kind : Deployment
metadata :
  name : tomcat
spec :
  replicas : 1
  selector :
    matchLabels :
      run : tomcat
  template :
    metadata :
      labels :
        run : tomcat
    spec :
      containers :
        - image : tomcat : 7 . 0
          imagePullPolicy : Always
          name : tomcat
          ports :
            - containerPort : 8080
              protocol : TCP
          restartPolicy : Always
---
apiVersion : v1
kind : Service
metadata :
  name : tomcat
spec :
  ports :
    - port : 8080
      protocol : TCP
      targetPort : 8080
  selector :
    run : tomcat
```

```
clusterIP : None
```

```
$ kubectl apply -f tomcat - service . yml #
Enter the template above in tomcat - service . yaml
deployment " tomcat " created
service " tomcat " created

# After deployment is complete

$ kubectl get svc , deploy tomcat
NAME          TYPE          CLUSTER - IP      EXTERNAL - IP      PORT
( S )         AGE
svc / tomcat   ClusterIP      < none >           < none >           8080 /
TCP           1m

NAME          DESIRED      CURRENT      UP - TO - DATE
AVAILABLE     AGE
deploy / tomcat 1            1            1            1
1m
```

2. Apply for SLB instance

You must apply for a performance guarantee type SLB instance (such as slb.s2.small) under the cluster and region. According to the specific needs, private or public network can be used. In this example, apply for a public network SLB instance and record the ID of the SLB instance.

Instances									
Server Load Balancer Name		Enter instance names separated by commas		Search	Tag	Configure Auto Renew		Create Server Load Balancer	
Server Load Balancer ID/Name	Zone	IP Address(All)	Status	Network(All)	Port/Health Check/Backend Server	Instance Specification	Bandwidth Billing Method(All)	Billing Method(All)	Actions
slb-serverless	China East 2 Zone B(Master) China East 2 Zone C(Slave)	(Public IP)	Running	Classic Network	HTTPS : 443 Normal Not ConfiguredConfigure HTTP : 80 Normal Not ConfiguredConfigure	Guaranteed-Performance Instance slb.s2.small	Pay by Traffic	Pay-As-You-Go 2018-07-05 17:52:12 Created	Manage More

3. Configure the TLS certificate

You must configure the TLS certificate for HTTPS access.



Note:

System automatically initializes the SLB HTTPS default certificate according to the first created Ingress TLS certificate. If you want to modify the HTTPS default certificate, you can modify it in the SLB console. If you want to configure multiple certificates, you can manually add them in the SLB console HTTPS listener domain name extension.

```
# Generate test TLS certificate
$ openssl req -x509 -nodes -days 365 -newkey rsa : 2048
- keyout tls . key - out tls . crt - subj "/ CN = bar . foo
. com / 0 = bar . foo . com "# Create a TLS certificate
secret
$ kubectl create secret tls cert - example -- key tls .
key -- cert tls . crt
secret " cert - example " created
```

```
# View new TLS certificate

$ kubectl get secret cert - example
NAME          TYPE          DATA          AGE
cert - example  kubernetes . io / tls      2             12s
```

4. Configure Ingress

Configure the domain name and path exposed by Tomcat service by using Ingress.

The layout template is as follows:

```
apiVersion : extensions / v1beta1
kind : Ingress
metadata :
  name : tomcat - ingress
  annotations :
    # Configure to use the specified SLB instance ( SLB ID )
    service . beta . kubernetes . io / alicloud - loadbalancer - id : lb - xxxxxxxxxx ## Replace with your SLB ID
    service . beta . kubernetes . io / alicloud - loadbalancer - force - override - listeners : " true "
spec :
  tls :
    - hosts :
      - bar . foo . com
      # Configure the TLS certificate
      secretName : cert - example
  rules :
    # Configure Layer - 7 domain name
    - host : bar . foo . com
      http :
        paths :
          # Configure context path
          - path : /
            backend :
              serviceName : tomcat
              servicePort : 8080
```

```
apiVersion : extensions / v1beta1
kind : Ingress
metadata :
  name : tomcat - ingress
  annotations :
    # Configure to use the specified SLB instance ( SLB ID )
    service . beta . kubernetes . io / alicloud - loadbalancer - id : lb - xxxxxxxxxx ## Replace with your SLB ID
spec :
  tls :
    - hosts :
      - bar . foo . com
      # Configure the TLS certificate
      secretName : cert - example
  rules :
    # Configure Layer - 7 domain name
    - host : bar . foo . com
      http :
        paths :
          # Configure context path
```

```
- path : /
  backend :
    serviceNam e : tomcat
    servicePor t : 8080
```

```
$ kubectl apply -f tomcat - ingress . yml
# Enter the template above in tomcat - ingress . yml
ingress " tomcat - ingress " created

# After deployment is complete , ADDRESS is the
specified SLB IP address

$ kubectl get ing tomcat - ingress
NAME          HOSTS          ADDRESS          PORTS          AGE
tomcat - ingress  bar . foo . com  47 .***.***.**  80 ,  443
1m
```

5. Test service access



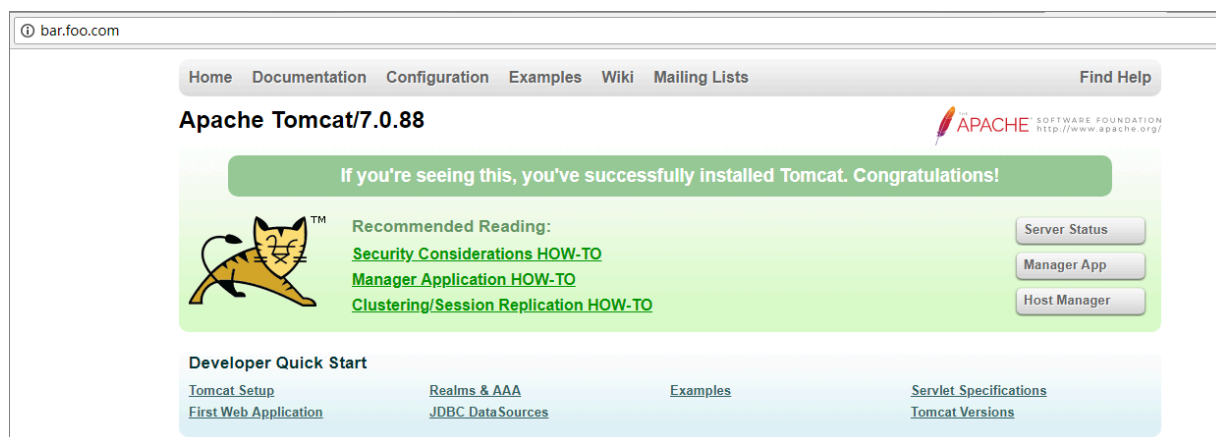
Note:

Currently, the domain name of the SLB instance IP must be resolved manually.

In this example, a DNS domain name resolution rule is added to `hosts` for testing service access. We recommend that you enter the domain name in your work environment.

```
47 .***.***.**  bar . foo . com
```

Test access to Tomcat service by using browser:



Test access to Tomcat service by using command line:

```
curl -k -H "Host : bar . foo . com " https :// 47 .***.***.**
```

2.7 Log management

2.7.1 Overview

You can view the logs of the Serverless Kubernetes cluster in the following ways.

- View the container running logs by using `kubectl logs` command.

For more information, see [kubectl logs](#).



Note:

Before using the `kubectl logs` command to view the container running logs, see [Connect to a Kubernetes cluster by using kubectl](#).

- [View cluster logs](#)
- [Collect logs by using Alibaba Cloud Log Service](#)

2.7.2 View cluster logs

You can view the cluster operation logs by using the simple log service of Container Service.

Procedure

1. Log on to the [Container Service console](#).
2. Under Kubernetes, click Clusters in the left-side navigation pane.
3. Select the target cluster and click View Logs on the right.

Cluster List

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

Refresh

Create Serverless Kubernetes Cluster

Create Kubernetes Cluster

Help: [Create cluster](#) [Scale cluster](#) [Connect to Kubernetes cluster via kubectl](#) [Manage applications with commands](#)

Name

Cluster Name/ID	Cluster Type	Region (All)	Network Type	Cluster Status	Time Created	Kubernetes Version	Action
serverless	Serverless Kubernetes	China East 2 (Shanghai)	VPC vpc-uf6in78kva3...	Running	08/13/2018,11:32:47	1.9.3	Manage View Logs Delete

View the cluster operation information.

Cluster Logs: serverless

Back to Cluster List

Refresh

Time	Information
08/13/2018,11:32:55	c13f665b4ccb744828cf18cc32801ea10 Allocate controller resource successfully
08/13/2018,11:32:55	c13f665b4ccb744828cf18cc32801ea10 Create Serverless Kubernetes cluster successfully
08/13/2018,11:32:52	c13f665b4ccb744828cf18cc32801ea10 Create security group sg-uf6f8e7wo9fmywlv3bv4 Successfully
08/13/2018,11:32:52	c13f665b4ccb744828cf18cc32801ea10 Start to allocate controller resource...
08/13/2018,11:32:52	c13f665b4ccb744828cf18cc32801ea10 Prepare aliyun resource successfully
08/13/2018,11:32:52	c13f665b4ccb744828cf18cc32801ea10 Start to create cluster resource
08/13/2018,11:32:52	c13f665b4ccb744828cf18cc32801ea10 Successfully creating cluster resource
08/13/2018,11:32:51	c13f665b4ccb744828cf18cc32801ea10 Wait for cluster order finish successfully
08/13/2018,11:32:51	c13f665b4ccb744828cf18cc32801ea10 Start to create Serverless Kubernetes cluster...
08/13/2018,11:32:51	c13f665b4ccb744828cf18cc32801ea10 Start to prepare aliyun resource

2.7.3 Collect logs by using Alibaba Cloud Log Service

This document uses the sample Tomcat application logs as an example. The log output path is `/usr/local/tomcat/logs/catalina.*.log`.

Step 1. Create a Log Service project

For more information, see [Manage a project](#).

Step 2. Create a Logstore

For more information, see [Manage a Logstore](#).

Step 3. Create a machine group



Note:

We recommend that the machine group ID be set to user-defined ID. User-defined ID is configured to yaml environment variable `ALIYUN_LOG TAIL_USER_DEFINED_ID`.

Step 4. Create a Logtail configuration

`/ecilogs` is the `mountPath` path to the `ilogtail` container in yaml.

* Configuration Name:

* Log Path:

All files under the specified folder (including all directory levels) that conform to the file name convention will be monitored. The file name can be a complete name or a name that contains wildcards. The Linux file path must start with "/"; for example, `/apsara/nuwa/.../app.Log`. The Windows file path must start with a drive; for example, `C:\Program Files\Intel\...*.Log`.

Docker File: ☐

If the file is in the docker container, you can directly configure the internal path and container label, Logtail will automatically monitor the create and destroy of the container, and collect the log of the specified container according to specified label

Mode:

Reminder : In Simple Mode, each line is treated as one log. The system will not extract the fields in the logs, and the parse time is used as the log time.

Advanced Options: Open ▾

Previous Next

Step 5. Create a Tomcat application

Configure the corresponding environment variable (ENV) based on the following template example and create the pod.

```
apiVersion : apps / v1beta2
kind : Deployment
metadata :
  name : tomcat - app
```

```



labels :
  k8s - app : tomcat - app
spec :
  selector :
    matchLabels :
      k8s - app : tomcat - app
  template :
    metadata :
      labels :
        k8s - app : tomcat - app
    spec :
      containers :
        - name : ilogtail
          image : registry - vpc .${ your_region_id }. aliyuncs .
com / acs / logtail : 0 . 16 . 16 . 0 - 1cf247c - aliyun
          env :
            - name : " ALIYUN_LOG_TAIL_CONFIG "
              value : "/ etc / ilogtail / conf / ${ your_region_id } /
ilogtail_config . json "
            - name : " ALIYUN_LOG_TAIL_USER_ID "
              value : "${ your_aliyun_user_id }"
            - name : " ALIYUN_LOG_TAIL_USER_DEFINED_ID "
              value : "${ your_machine_group_name }"
          volumeMounts :
            - name : tomcat - log
              # The following root path will be added to
the Logtail configuration .
              mountPath : / ecilogs
              readOnly : true
            - name : tomcat
              image : tomcat : 7 . 0
              volumeMounts :
                - name : tomcat - log
                  # Set the directory to which the application
logs are output .
                  mountPath : / usr / local / tomcat / logs
          volumes :
            - name : tomcat - log
              emptyDir : {}

```

**Note:**

- You need to replace `${ your_aliyun_user_id }` with your primary account ID that can be obtained by running the `echo $ ALIBABA_CLOUD_ACCOUNT_ID` command in [CloudShell](#).
- You need to replace `${ your_machine_group_name }` with the machine group ID that you set in the corresponding project in Log Service. For more information, see [Create an ID to identify a machine group](#).
- You need to replace `${ your_region_id }` with your region. For example, `cn-hangzhou` and `ap-southeast-1`.
- We recommend that you use the image tag of the latest version. For more information, see [Tag list](#).

Step 6. View the Tomcat logs in Log Service

1		03-14 21:04:29	__source__: 172.16.3.4 __tag__:__hostname__: tomcat-app-644fd86875-2bkr9 __tag__:__path__: /ecilog/catalina.2018-03-14.log __tag__:__user_defined_id__: eci-app __topic__: content: INFO: Server startup in 1631 ms
2		03-14 21:04:29	__source__: 172.16.3.4 __tag__:__hostname__: tomcat-app-644fd86875-2bkr9 __tag__:__path__: /ecilog/catalina.2018-03-14.log __tag__:__user_defined_id__: eci-app __topic__: content: Mar 14, 2018 1:04:29 PM org.apache.catalina.startup.Catalina start

2.8 Service discovery based on Alibaba Cloud DNS Private Zone

Alibaba Cloud Serverless Kubernetes supports service discovery. Currently, service discovery for intranet SLB and headless service is supported.

Alibaba Cloud DNS Private Zone

Alibaba Cloud DNS Private Zone is a private domain name resolution and management service based on VPC (virtual private cloud) environment of Alibaba Cloud. You can map a private domain name to an IP resource address in one or more of the custom private networks, while your private domain names are not accessible in other network environments.

Prerequisites

1. You have activated Alibaba Cloud DNS Private Zone in the Alibaba Cloud DNS console.
2. You have successfully created a Serverless Kubernetes cluster, see [Create a serverless Kubernetes cluster](#).
3. You have connected to a Serverless Kubernetes cluster, see [Connect to a Kubernetes cluster by using kubectl](#).

Procedure

1. Connect to the Kubernetes cluster by using kubectl and run the following command to confirm the connection to the cluster.

```
kubectl cluster - info
```

```
Kubernetes master is running at https://xxxxxx.
serverless-1.kubernetes.cn-shanghai.aliyuncs.com:6443
```

2. Deploy and create a service. Currently, only intranet service and headless service are supported.

Take the intranet service as an example. Create a sample file: `nginx -`

`deployment - basic . yaml .`

```
vim nginx - deployment - basic . yaml
```

The sample template is as follows, copy the following yaml code in the yaml file.

Then run the `kubectl create -f nginx - deployment - basic . yaml` command to create it.

```
apiVersion : apps / v1beta2 # for versions before 1 . 8 .
0 use apps / v1beta1
kind : Deployment
metadata :
  name : nginx - deployment - basic
  labels :
    app : nginx
spec :
  replicas : 2
  selector :
    matchLabels :
      app : nginx
  template :
    metadata :
      labels :
        app : nginx
    spec :
      containers :
        - name : nginx
          image : nginx : 1 . 7 . 9 # replace it
          ports :
            - containerPort : 80
---
apiVersion : v1
kind : Service
metadata :
  name : nginx - service - intranet # Access by
  annotations :
    service.beta.kubernetes.io/alibaba-cloud-loadbalancer-address-type : intranet ## Add annotation
spec :
  ports :
    - port : 80
      protocol : TCP
  selector :
    app : nginx
```

```
type : LoadBalancer
```

You can also create services of the headless service type, as shown in the following sample template.

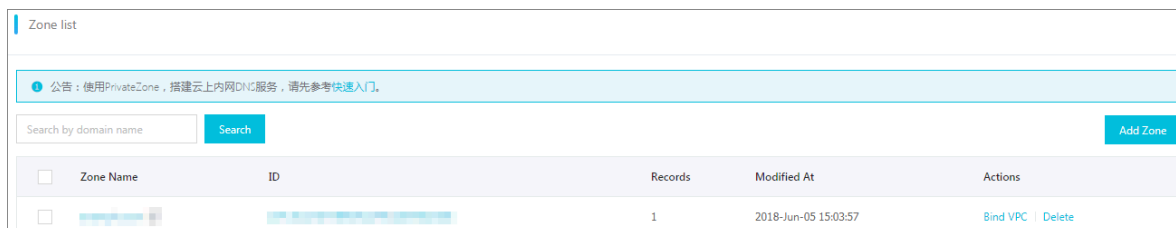
```
apiVersion : v1
kind : Service
metadata :
  name : nginx - service - headless
spec :
  ports :
    - port : 80
      protocol : TCP
  selector :
    app : nginx
  clusterIP : None
```

3. Execute the following command to view the health status of the application.

```
kubectl get svc , pod , deployment
```

4. Log on to the [Alibaba Cloud DNS console](#).

5. Click Private Zone > Zone List in the left-side navigation pane to see that the record is automatically generated under this list.



Zone Name	ID	Records	Modified At	Actions
		1	2018-Jun-05 15:03:57	Bind VPC Delete

You can access Service (long domain or short domain) by the private domain name in VPC network environment.

- Long domain access: In this example, `nginx - service - intranet . $NAMESPACE . svc . cluster . local` , where `$NAMESPACE` is the Serverless cluster ID, which can be seen in the console, or you can also use environmental variables in the yaml file of the pod.

```
env :
  - name : NAMESPACE
    valueFrom :
      fieldRef :
        fieldPath : metadata . namespace
```

- Short domain access: `nginx-service-intranet` or `nginx-service-headless`, the name of the service defined in the yaml template.

For more information, see [serverless-k8s-examples](#).