

Alibaba Cloud Aliyun Container for Kubernetes

Product Introduction

Issue: 20190419

Legal disclaimer

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








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

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

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

Generic conventions

Table -1: Style conventions

Style	Description	Example
	This warning information indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	 Danger: Resetting will result in the loss of user configuration data.
	This warning information indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	 Warning: Restarting will cause business interruption. About 10 minutes are required to restore business.
	This indicates warning information, supplementary instructions, and other content that the user must understand.	 Notice: Take the necessary precautions to save exported data containing sensitive information.
	This indicates supplemental instructions, best practices, tips, and other content that is good to know for the user.	 Note: You can use Ctrl + A to select all files.
>	Multi-level menu cascade.	Settings > Network > Set network type
Bold	It is used for buttons, menus, page names, and other UI elements.	Click OK.
Courier font	It is used for commands.	Run the <code>cd / d C :/ windows</code> command to enter the Windows system folder.
<i>Italics</i>	It is used for parameters and variables.	<code>bae log list --instanceid Instance_ID</code>
[] or [a b]	It indicates that it is an optional value, and only one item can be selected.	<code>ipconfig [-all -t]</code>

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

Contents

Legal disclaimer.....	I
Generic conventions.....	I
1 What is Container Service for Kubernetes.....	1
2 Architecture.....	4
3 Advantages.....	6
4 Scenarios.....	9
5 Limits.....	14
6 Terms.....	16
7 Serverless Kubernetes cluster.....	20
8 Types of Kubernetes clusters supported by Container Service.....	22
9 Update history.....	25
9.1 Changes to Alibaba Cloud Container Service for Kubernetes.....	26
9.2 Changes to Cloud Controller Manager.....	44

1 What is Container Service for Kubernetes

Container Service for Kubernetes provides the high-performance and scalable container application management service, which enables you to manage the lifecycle of enterprise-class containerized applications by using Kubernetes. By simplifying the setup and capability expansion of cluster and integrating with the Alibaba Cloud abilities of virtualization, storage, network, and security, Container Service for Kubernetes makes an ideal running cloud environment for Kubernetes containers. As a service provider certified by Kubernetes and the world's first platform service that has passed Kubernetes consistency certification, Container Service provides you with professional support and services.

Container Service for Kubernetes contains the classic dedicated Kubernetes mode and the serverless Kubernetes mode for you to choose on demand.

- **Classic dedicated Kubernetes mode:** You can get more fine-grained control over cluster infrastructure and container applications, for example, select the host instance specification and the operating system, specify Kubernetes version, custom Kubernetes attribute switch settings, and more. Alibaba Cloud Container Service for Kubernetes is responsible for creating the underlying cloud resources for the cluster, upgrading and other automated operations for the cluster. You need to plan, maintain, and upgrade the server cluster. You can add servers to or remove servers from the cluster manually or automatically.
- **Serverless Kubernetes mode:** You do not need to create the underlying virtualization resource. To launch the application directly, use Kubernetes commands to specify the application container image, CPU and memory requirements as well as external service methods.

Features

Cluster management

- You can create a classic dedicated Kubernetes cluster through the console within ten minutes. GPU instances and ECS Bare Metal (EBM) Instance are supported. You can also create a highly available cluster that is across zones.
- Provides OS images for optimizing containers, and the Kubernetes and Docker versions of stability testing and security hardening.

- Supports multi-cluster management, highly available clusters across zones, and cluster federation management.

One-stop container lifecycle management

- Network

Provides the high performance Virtual Private Cloud (VPC) and elastic network interfaces (ENI) network plug-in optimized for Alibaba Cloud, which is 20% better than the average network solution.

Supports container access policies and flow control restrictions.

- Storage

Container Service integrates with Alibaba Cloud cloud disk, Network Attached Storage (NAS), and Object Storage Service (OSS), and provides the standard FlexVolume drive.

Supports dynamic creation and migration of storage volumes.

- Logs

Supports high-performance automatic log collection and integrating with Alibaba Cloud Log Service.

You can also integrate Container Service with third-party open-source log solutions

.

- Monitoring

Supports the monitoring at the level of containers and virtual machines (VMs).

You can also integrate Container Service with third-party open-source monitoring solutions.

- Permission

Supports Resource Access Management (RAM) authorization and management at the level of clusters.

Supports permission configuration management at the level of applications.

- Application management

Support gray release and blue-green release.

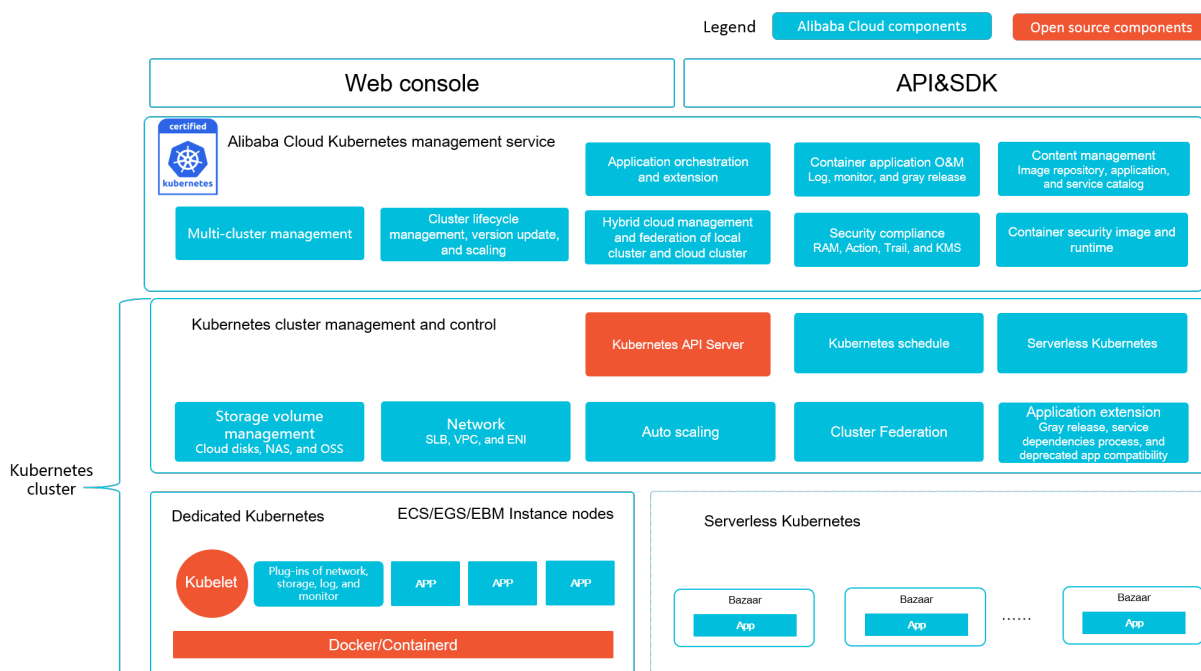
Support application monitoring and application elastic scaling.

Supports one-click deployment of the Helm application by the built-in application store. Supports the directory of services to simplify cloud service integration.

Easily deal with upstream and downstream delivery process by using high-availability scheduling policy

- **Supports affinity policy and horizontal scaling of services.**
- **Supports high availability across zones and disaster recovery.**
- **Supports the APIs for cluster and application management to easily interconnect with the continuous integration and private deployment system.**

2 Architecture



Description

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

Features	Description
Multiple Kubernetes clusters	<p>Integrated with Alibaba Cloud virtualization technology, Container Service for Kubernetes supports the Dedicated cluster and Serverless cluster.</p> <ul style="list-style-type: none">· Dedicated: Elastic Compute Service (ECS), Elastic GPU Service (EGS), and ECS Bare Metal Instance can be used as cluster nodes. The instance specifications can be configured flexibly and support a wide range of plug-ins.· Serverless: Provides Kubernetes Serverless services which are free of server management, simplifies the underlying resource management, supports flexible capacity expansion, and reduces resource overhead.
Alibaba Cloud Kubernetes cluster management and control service	Support powerful network, storage, hybrid cluster management, horizontal capacity expansion, application extension, and other features.
Alibaba Cloud Kubernetes management service	Supports secure images and Helm tools , and is highly integrated with Alibaba Cloud Resource Access Management (RAM), Key Management Service (KMS), log, monitoring and other products to provide a secure and compliant Kubernetes solution that provides you with enhanced capabilities such as hybrid cloud, container security, CI/CD, DevOps and more.
Convenient and efficient use	Container Service for Kubernetes provides Web console, API&SDK.

3 Advantages

Advantages of Container Service for Kubernetes

Ease of use

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

You may have to deal with self-built Kubernetes clusters of different versions at the same time, including version 1.8.6, 1.9.4, 1.10, and later. 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 functions

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>

Function	Description
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>
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 Alibaba Cloud 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 hangs, 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.

4 Scenarios

DevOps continuous delivery

Optimal continuous delivery process

Working with Jenkins, Container Service encompasses the complete process of DevOps from code submission to application deployment in an automated fashion. It makes sure that only codes passing the automated test can be delivered and deployed, and efficiently replaces the traditional method of complicated deployment and slow iteration in the industry.

Container Service can implement:

- Automation of DevOps.

The automation of the full process from code changes to code building, image building, and application deployment.

- Consistency of environment.

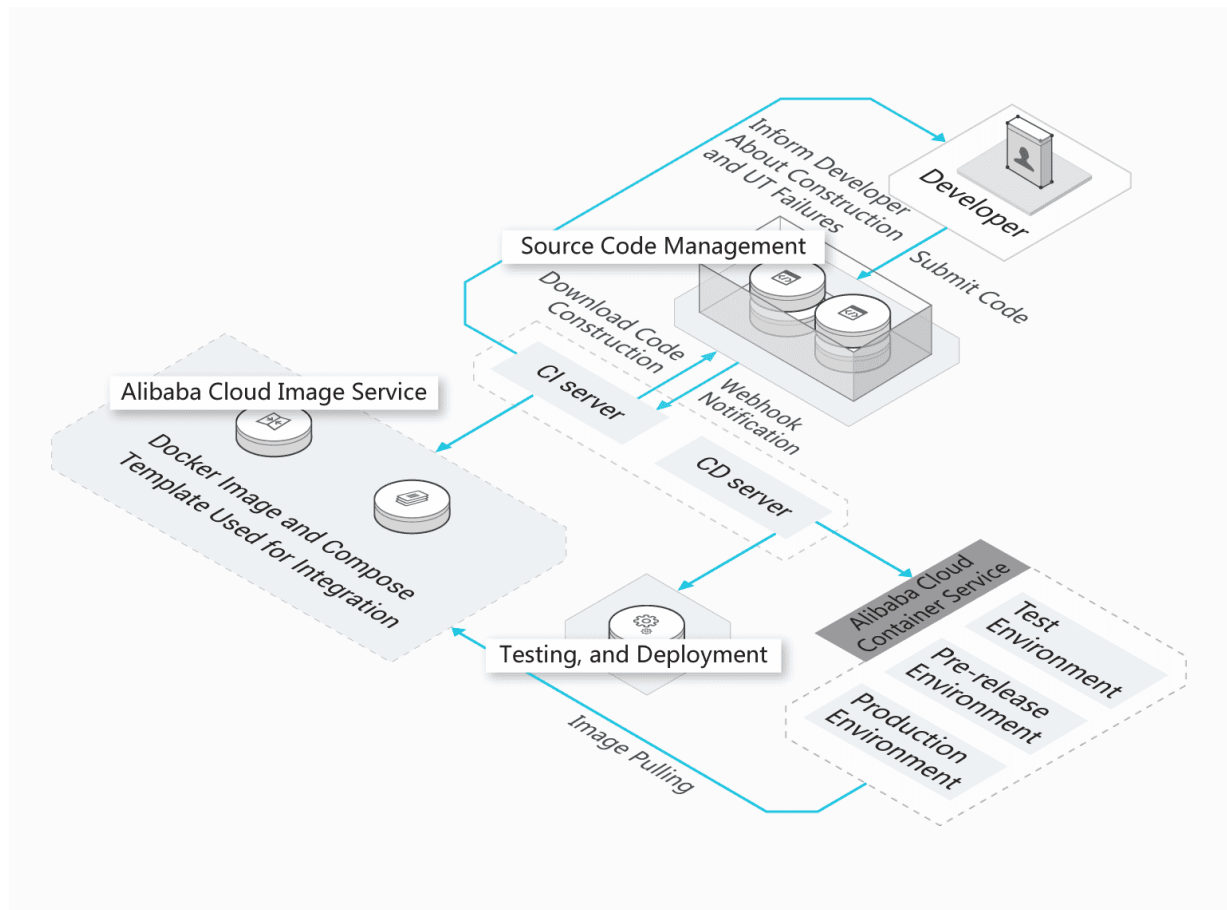
Container Service allows you to deliver not only code but also a running environment based on the immutable architecture.

- Continuous feedback

Results of each integration or delivery are fed back in real time.

We recommend that you use

Elastic Compute Service (ECS) and Container Service together.



Microservice architecture

Implement agile development and deployment to accelerate business iteration of enterprises

In the production environment of enterprises, microservices are divided reasonably and each microservice application is stored in the Alibaba Cloud image repository. You only have to iterate each microservice application, and Alibaba Cloud provides the capabilities of scheduling, orchestration, deployment, and gated launch.

Container Service can implement:

- Server Load Balancer and service discovery.

Supports Layer-4 and Layer-7 request forwarding and backend binding.

- Many policies of scheduling and exception recovery.

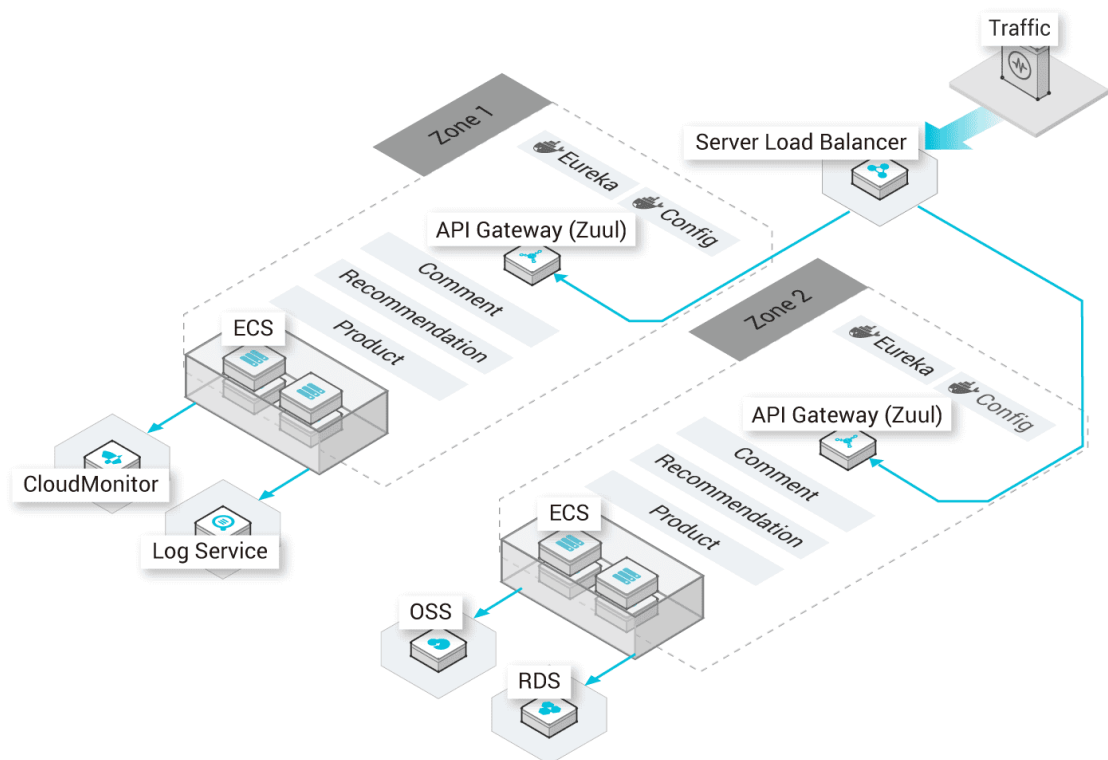
Supports affinity scheduling at the level of services. Supports cross-zone high-availability and disaster recovery.

- Microservice monitoring and auto scaling.

Supports the monitoring at the level of microservices and containers. Supports auto scaling of microservices.

We recommend that you use

ECS, Relational Database Service (RDS), Object Storage Service (OSS), and Container Service together.



Hybrid cloud architecture

Unified Operation and Maintenance of multiple cloud resources

Manage resources on and off the cloud at the same time in the Container Service console, without switching between multiple cloud consoles. Deploy applications on and off the cloud at the same time by using the same image and orchestration based on the characteristics unrelated to the container infrastructure.

Container Service supports:

- Scaling in and out applications on the cloud.

Expand the capacity rapidly on the cloud at the business peak period to bring some business traffic to the cloud.

- Disaster recovery on the cloud.

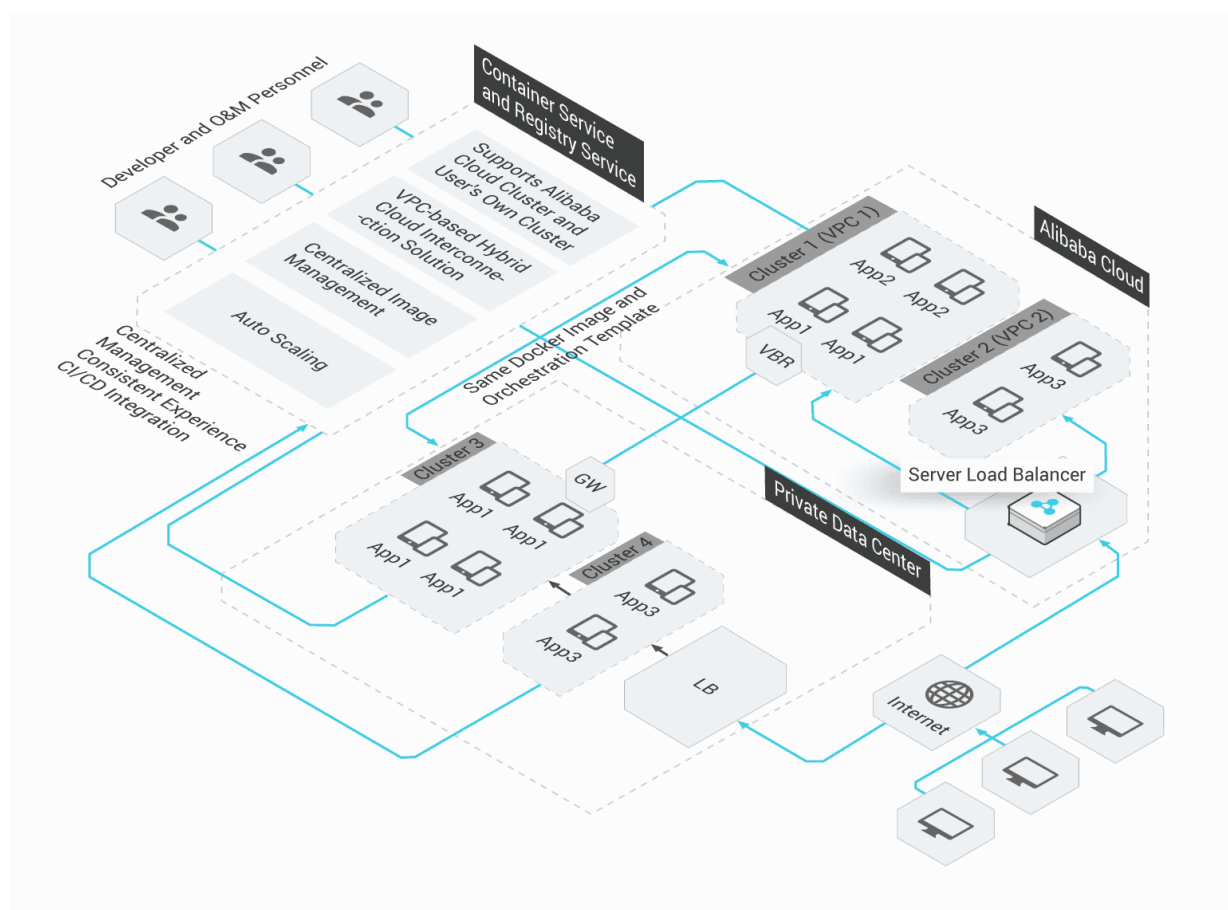
Deploy business systems on and off the cloud at the same time to provide services off the cloud and provide disaster recovery on the cloud.

- Development and test off the cloud.

Release the applications seamlessly on the cloud after the development and test off the cloud.

We recommend that you use

ECS, Virtual Private Cloud (VPC), and Express Connect together.



Auto scaling architecture

Automatic expansion/contraction for the business according to the business traffic

Container Service can automatically expand or contract the business according to the business traffic, without manual intervention. In this way, the system is not

down because of traffic surge and not timely expansion, and the waste due to a large number of idle resources is avoided.

Container Service can implement:

- Rapid response.

Trigger the container expansion in seconds when the business traffic reaches the expansion indicator.

- Full automation.

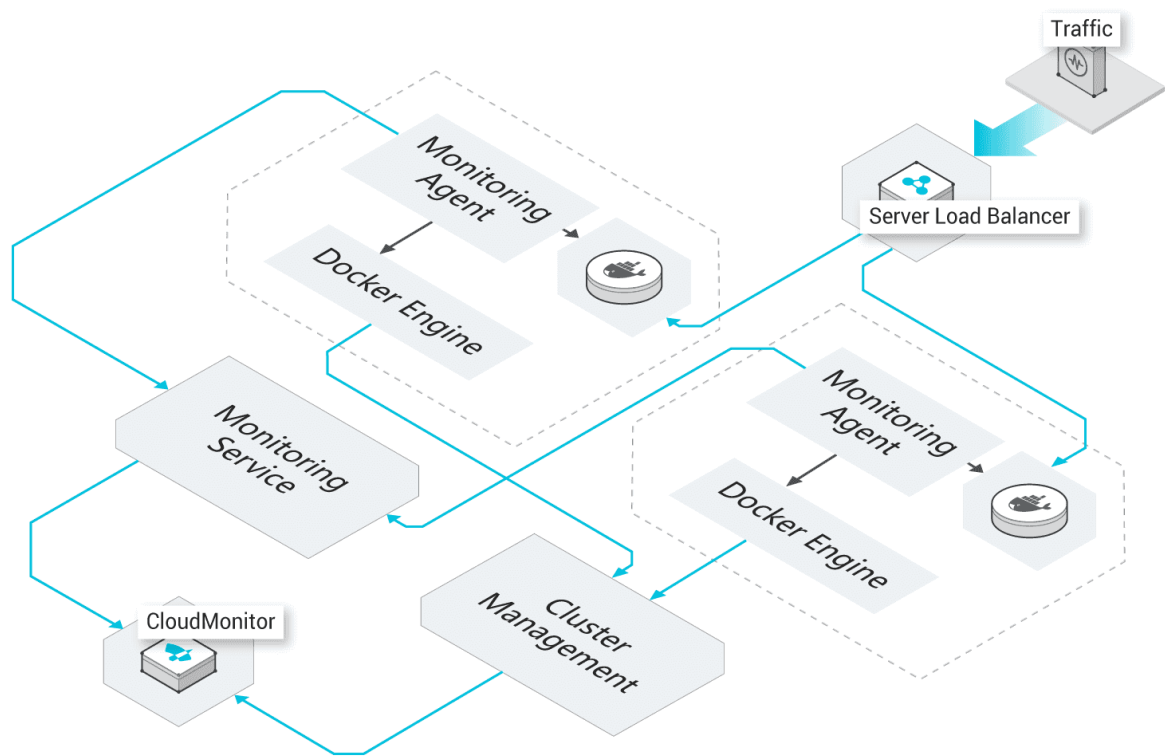
The expansion or contraction process is fully automated, without manual intervention.

- Low cost.

Contract the capacity automatically when the traffic is reduced to avoid the waste of resources.

We recommend that you use

ECS and CloudMonitor together.



5 Limits

Limits for Alibaba Cloud Container Service Kubernetes clusters are as follows.

Clusters

- By default, you can create up to five clusters in all regions and add up to 40 nodes to each cluster. To create more clusters or add more nodes to a cluster, please open a ticket.
- Currently, Kubernetes clusters only support Linux containers. The support for Windows containers is planned.
- Currently, Kubernetes clusters only support Virtual Private Cloud (VPC). You can create a VPC or use an existing VPC when creating a Kubernetes cluster.
- Currently, Elastic Compute Service (ECS) instances support the following two billing methods: Pay-As-You-Go and Subscription. You can change the billing method of your instance from Pay-As-You-Go to Subscription in the ECS console. Other resources such as Server Load Balancer (SLB) instances support only the Pay-As-You-Go billing method.

ECS instances

- Only support the CentOS operating system.
- Limits for adding an existing node:
 - The ECS instance to be added must be in the same region as the cluster and use VPC.
 - When adding an existing ECS instance, make sure that your ECS instance has an Elastic IP (EIP) for the network type VPC. Otherwise, the ECS instance will fail to be added.
 - The ECS instance to be added must be under the same account as the cluster.

Cluster expansion and contraction

- The number of worker nodes must be in the range of 1 to 37.
- Manual and automatic cluster elastic scaling are supported.
- Currently, master nodes in a Kubernetes cluster cannot be expanded automatically.
- According to the Resource Orchestration Service (ROS) rules, elastic contraction does not contract the nodes created when the cluster is created or the nodes

manually added to the cluster, but only contracts the nodes manually expanded by you. The contraction rule is based on the creation time. When you contract a cluster, the worker nodes are removed from the cluster in the order that they are added when you expand the cluster.

- Nodes added through cluster expansion are Pay-As-You-Go nodes.

6 Terms

Basic terms

Clusters

A collection of cloud resources that are required to run containers. It associates with several server nodes, Server Load Balancer instances, Virtual Private Cloud (VPC), and other cloud resources.

Nodes

A server (either an Elastic Compute Service (ECS) instance or a physical server) that has a Docker Engine installed and is used to deploy and manage containers. The Agent program of Container Service is installed in a node and registered to a cluster. The number of nodes in a cluster is scalable.

Container

A running instance created by using a Docker image. A single node can run multiple containers.

Image

A standard packaging format of a container application in Docker. You can specify an image to deploy containerized applications. The image can be from the Docker Hub, Alibaba Cloud image service, or your private registry. An image ID is uniquely identified by the URI of the image repository and the image tag (latest by default).

Kubernetes terms

Node

Nodes in Kubernetes clusters provide the computing power and are the active servers where all the pods are running. The active servers can be physical machines or virtual machines. Containers that are running on the kubelet management nodes must run on the active servers.

Namespace

The namespace provides the virtual isolation for Kubernetes clusters. By default, Kubernetes clusters have three namespaces: the default namespace default, the

system namespaces kube-system, and kube-public. The administrator can also create new namespaces to meet the requirements.

Pod

The minimum basic unit of Kubernetes that is used to deploy applications or services.

A pod can encapsulate one or multiple containers, storage resources, an independent network IP address, and policy options of managing and controlling the running method of containers.

Replication Controller

Replication Controller (RC) makes sure that a specified number of pod replicas are running in a Kubernetes cluster at any time. Replication Controller (RC) makes sure that a specified number of pod replicas are running in a Kubernetes cluster at any time by monitoring the running pod. One or more pod replicas can be specified. If the number of pod replicas is less than the specified number, RC starts to run new pod replicas. If the number of pod replicas exceeds the specified number, RC starts to stop the redundant pod replicas.

Replica Set

The upgraded version of RC. The only difference between Replica Set (RS) and RC is the support for selector. RS supports more types of matching modes. Generally, the RS objects are not used independently, but are used as the deployment parameters in the ideal status.

Deployment

The deployment indicates an update operation for a Kubernetes cluster by users, has a wider application range than RS, and can create a service, update a service, and perform a rolling update of a service. Performing a rolling update of a service actually creates a new RS, gradually adds the number of replicas in the new RS to the ideal status, and reduces the number of replicas in the old RS to zero. Such a compound operation cannot be described well by an RS, but can be described by a more common deployment. We do not recommend that you manually manage and use the RS created by the deployment.

Service

The basic operation unit of Kubernetes. As the abstract of real application service, each service has many containers to provide the support. The port of Kube-Proxy and

service selector determine the service request to pass to the backend container, and a single access interface is displayed externally. The external is not required to know how the backend works, which is good for expanding or maintaining the backend.

Labels

Essentially a collection of key-value pairs that are attached to the resource objects. Labels are used to specify the attributes of objects that are meaningful for users, but do not have any direct significance for kernel systems. You can add a label directly when creating an object or modify the label at any time. Each object can have more labels, but the key value must be unique.

Volume

The volumes in Kubernetes clusters are similar to the Docker volumes. The only difference is that the range of Docker volumes is a container, while the lifecycle and range of Kubernetes volumes are a pod. The volumes declared in each pod are shared by all the containers in the pod. You can use the Persistent Volume Claim (PVC) logical storage, and ignore the actual storage technology in the backend. The specific configurations about Persistent Volume (PV) are completed by the storage administrator.

PV and PVC

PV and PVC allow the Kubernetes clusters to have the abstract logical capabilities of storage, and you can ignore the configurations of actual backend storage technology in the pod configuration logic, leaving the configurations to the PV configurator. The relationship between PV and PVC of storage is similar to that between node and pod of computing. PV and node are the resource provider, changed according to the cluster infrastructure, and configured by the Kubernetes cluster administrator. PVC and pod are the resource user, changed according to the business and service requirements, and configured by the Kubernetes cluster user, namely, the service administrator.

Ingress

A collection of rules that authorize the inbound access to the cluster. You can provide the externally accessible URL, Server Load Balancer, SSL, and name-based virtual host by using the Ingress configurations. You can request the Ingress by posting Ingress resources to API servers. The Ingress Controller is used to implement Ingress

by using Server Load Balancer generally, and can configure the edge router and other frontends, which helps you handle the traffic in the HA method.

Related documents.

- [Docker glossary](#)
- [Kubernetes concepts](#)

7 Serverless Kubernetes cluster

Introduction

Alibaba Cloud serverless Kubernetes allows you to quickly create Kubernetes container applications without having to manage and maintain clusters and servers. Based on the amount of CPU and memory resources actually used by the applications, Pay-as-you-go is applied. By using serverless Kubernetes, you can focus on designing and building applications, rather than managing the infrastructure on which your applications run. It is based on the Alibaba Cloud flexible computing infrastructure and is fully compatible with the Kubernetes API, combining the security, elasticity, and Kubernetes ecosystem of virtualized resources.

Alibaba Cloud Kubernetes service has two modes: one is the classic cluster approach, and the other is the serverless mode.

- In the classic Kubernetes cluster mode, you can get more fine-grained control over cluster infrastructure and container applications. For example, select the host instance specification (ECS/EGS/EBM) and the operating system, and specify the Kubernetes version, custom Kubernetes attribute switch settings, use different network models, and more. The Alibaba Cloud Kubernetes service is responsible for creating the underlying cloud resources for the cluster, upgrading the cluster and other automated operations. It monitors CPU, memory, and other resources in the cluster, and finds the server that best fits the container according to the resource requirements you specify. You need to plan, maintain, and upgrade the server cluster. You can also add or remove servers in the cluster manually or automatically by monitoring the resource usage of the application container.
- In serverless mode, you do not need to create the underlying virtualization resource. You can launch the application directly by using the Kubernetes command to indicate application container image, CPU and memory requirements, and external service methods.

Benefits

- **Easy to use:** Second-level application deployment. No need to manage the Kubernetes cluster infrastructure for high availability and security.
- **Compatibility:** You can deploy container applications by using the familiar Kubernetes command line and API that support Kubernetes typical application

scenarios without modification. You can also easily migrate your applications to a running platform that supports Kubernetes, such as the Alibaba Cloud Kubernetes cluster.

- **Secure isolation:** Based on the flexible computing infrastructure of the Alibaba Cloud, application containers are isolated from each other to prevent interference.
- **On-demand capacity expansion:** You do not need to worry about horizontal expansion of cluster nodes. Based on application load, you can easily and flexibility to scale up the resources needed for your application.
- **Connectivity:** Supports seamless integration of container applications with Alibaba Cloud basic services. Supports direct interaction with existing applications and databases in your Virtual Private Cloud (VPC); supports the connection of containers with virtual machine applications.

Architecture

The comparison architecture of serverless Kubernetes Container Service and Kubernetes Container Service is shown below.

Scenarios

Serverless Kubernetes has a wide range of application scenarios. The high portability and flexibility based on containers, and the high elasticity and isolation of Infrastructure as a Service (IaaS) provide a powerful base operating facility for serverless applications. Serverless applications include multimedia processing, captured data modification, IoT sensor data processing, streaming computing, chat robotics, batch computing, web applications, mobile application backend services, business logic processing, and continuous integration.

8 Types of Kubernetes clusters supported by Container Service

This topic describes three types of Kubernetes clusters (dedicated, managed, and serverless) in terms of their features, the resources that you pay for, their applicable scenarios, the user features, and the creation procedures.

Alibaba Cloud Container Service for Kubernetes (ACK for short) provides high-performance and scalable cloud native application management capabilities, and supports full life cycle management of enterprise-level Kubernetes containerized applications.

Features

- **Dedicated Kubernetes cluster**

You must create three Master nodes and one or multiple Worker nodes for the cluster. In addition, you need to plan, maintain, and upgrade the cluster as needed. With such a Kubernetes cluster, you can control cluster infrastructures in a more fine-grained manner.

- **Managed Kubernetes cluster**

You only need to create Worker nodes for the cluster, and Alibaba Cloud Container Service for Kubernetes creates and manages Master nodes for the cluster. This type of Kubernetes cluster is easy to use, low-cost, and highly available. You can focus on the services supported by the cluster without needing to operate and maintain the Kubernetes cluster Master nodes.

- **Serverless Kubernetes cluster**

You do not need to create or manage any Master nodes or Worker nodes for the cluster. You can directly use the Container Service console or the command line interface to set container resources, specify container images for applications, set methods to provide services, and start applications.

Resources that you pay for

- **Dedicated Kubernetes cluster**

The Kubernetes cluster can be used free of charge. You are charged only for Master node and Worker node resources.

- **Managed Kubernetes cluster**

The Kubernetes cluster can be used free of charge. You are charged only for Worker node resources.

- **Serverless Kubernetes cluster**

You are charged for the resources used by container instances. The amount of used resources is measured according to resource usage duration (in seconds).

Applicable scenarios

- **Dedicated Kubernetes cluster**

Applies to all scenarios.

- **Managed Kubernetes cluster**

Applies to all scenarios.

- **Serverless Kubernetes cluster**

Applies to batch tasks, burst expansion, and CI/CD testing.

User features

- **Dedicated Kubernetes cluster users**

- Have little to no cost constraints regarding their requirements.
- Are capable of operating and maintaining Kubernetes clusters.
- Have a deep understanding of Kubernetes.
- Have clear cluster resource plans.
- Require tailored Master nodes.
- Require completely self-managed Kubernetes clusters.

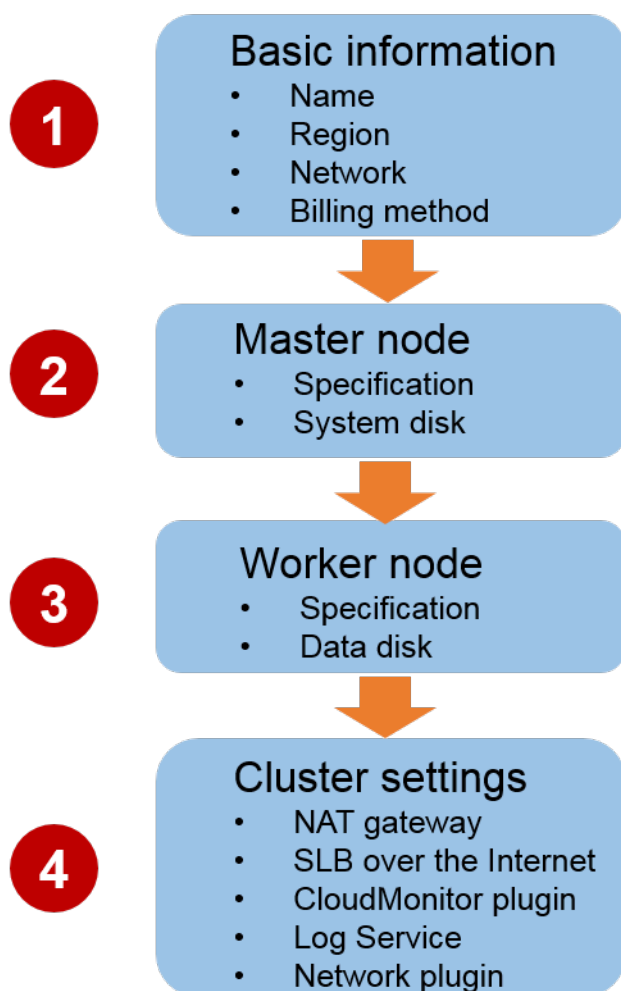
- **Managed Kubernetes cluster users**

- Have some cost constraints regarding their requirements.
- Want to focus on application deployment and operation without needing to worry about cluster management.
- Have a moderate understanding of Kubernetes.
- Require lightweight Kubernetes cluster O&M.
- Do not want to worry about Master node management.

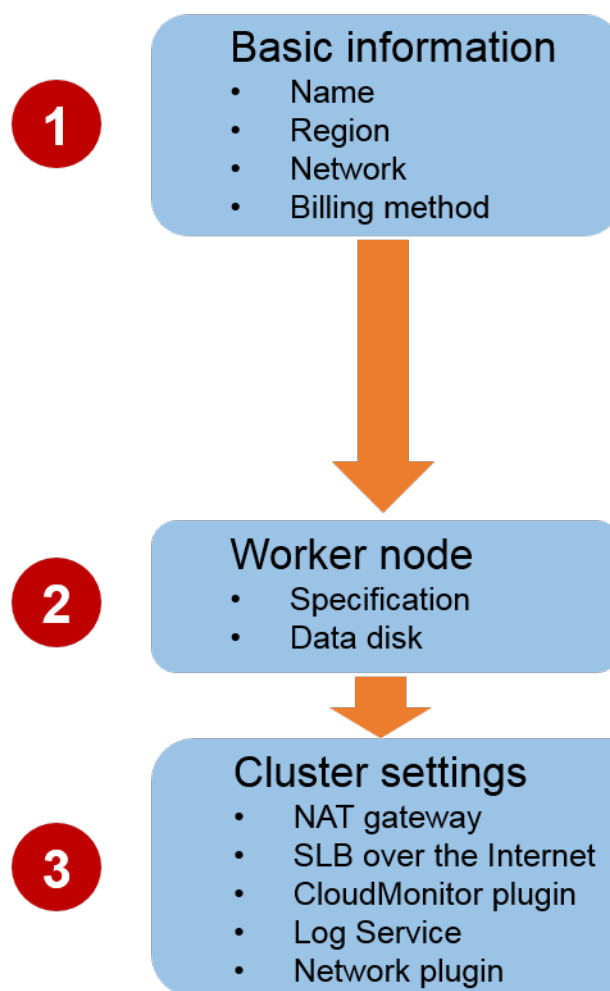
- Serverless Kubernetes cluster users
 - Need to process batch tasks.
 - Need to process bursts of service expansion.
 - Desire out-of-the-box functionality.
 - Have zero O&M requirements.
 - Prefer to pay for resources as they are used (Pay-As-You-Go billing method).
 - Do not want to worry about infrastructure setup.

Creation procedures

Dedicated Kubernetes cluster



Managed Kubernetes cluster



9 Update history

9.1 Changes to Alibaba Cloud Container Service for Kubernetes

This topic describes the changes to Alibaba Cloud Container Service for Kubernetes (ACK).

Date	Type	Change	Relevant topic
February 2019	New region	<p>For Alibaba Cloud users based in Mainland China and outside Mainland China, the cluster type of managed Kubernetes is released in the China South 1 (Shenzhen) region. This cluster type offers the following advantages:</p> <ul style="list-style-type: none">· Resources are reserved by default. Specifically, each managed Kubernetes cluster reserves you three Master nodes.· Cluster O&M is easy. Alibaba Cloud Container Service for Kubernetes (ACK) creates and manages Master nodes for each cluster.· Highly secure clusters. ACK secures your clusters according to your requirements.	N/A

Date	Type	Change	Relevant topic
	New feature	ACK released a Knative add-on to help you quickly build a Knative Serving environment in an ACK supported Kubernetes cluster. Knative provides a request-driven computing model in which workloads can be scaled back to zero. Based on Kubernetes and Istio, it provides deployment and service features for serverless applications and functions.	N/A
	New feature	For Alibaba Cloud users based in Mainland China, the intelligent cluster O&M feature is released across all regions.. With this feature, you can check your cluster resources , components, configurations, and other information to quickly troubleshoot your cluster problems.	N/A

Date	Type	Change	Relevant topic
January 2019	New feature	Containerized applications in Windows is released (beta version). To use this feature , you need to add Windows nodes to your dedicated Kubernetes clusters or managed Kubernetes clusters , which run in ACK. To request access to this feature, open a ticket.	N/A

Date	Type	Change	Relevant topic
	New feature	<p>Container Registry (Enterprise edition) is released (beta version). This edition allows you to:</p> <ul style="list-style-type: none">· Deploy a container image repository that is built on exclusive resources.· Build, distribute , and host your images in a fully secure environment with ease. <p>We recommend that you use Container Registry (Enterprise edition) if you have a large number of cluster nodes and require a higher level of security.</p> <p>To request access to this feature, open a ticket.</p>	N/A

Date	Type	Change	Relevant topic
	New feature	<p>The intelligent cluster O&M function is released in the China East 1 (Hangzhou) region. This feature provides the best solutions for container cluster users in different scenarios.</p> <p>With this feature, you can check your cluster resources , components, configurations, and other information to quickly find out your cluster problems.</p>	N/A

Date	Type	Change	Relevant topic
	New feature	<p>ACK is interoperable with Application Real-Time Monitoring Service (ARMS). After you install the ARMS plugin in your cluster, you can easily monitor the performance of the applications deployed in the cluster.</p> <p>ARMS is a Java application performance management (APM) monitoring product developed by Alibaba Cloud . If you use ARMS to monitor a Java application, you only need to mount a probe in the application startup script without modifying any code . By comprehensively monitoring the application , the probe helps you quickly locate faulty and slow interfaces , reproduce parameter calling , detect memory leaks, and discover system bottlenecks, more efficiently diagnosing problems online.</p>	N/A

Date	Type	Change	Relevant topic
	Commercialized product	Charges for Elastic Container Instance (ECI) used by serverless Kubernetes is introduced. From January 22, 2019 on, you will be charged if you use a serverless Kubernetes cluster to create container instances. However, serverless Kubernetes can be still used free of charge.	For more information, see Billing method .
	New regions	For Alibaba Cloud users in Mainland China, the cluster type of serverless Kubernetes is released for the following regions : China North 2 (Beijing) and China South 1 (Shenzhen).	N/A
December 2018	New region	The UK (London) region is added as a region for ACK services for Alibaba Cloud users in Mainland China and outside Mainland China.	N/A

Date	Type	Change	Relevant topic
	New region	The cluster type of managed Kubernetes is released for the following regions : China East 2 (Shanghai), Asia Pacific SE 3 (Kuala Lumpur), and Asia Pacific SOU 1 (Mumbai).	N/A
	New feature	Nodes can be removed from a Kubernetes cluster. Furthermore, you can determine whether to release the ECS instance used by the node when you remove a node.	Remove a node.
	New feature	Creating an application of the DaemonSet type is supported. DaemonSet (a type of daemon) guarantees that each node of a cluster has only one pod.	N/A
	New feature	Support for customizing Istio gateway parameters is added. You can set different parameters to customize Istio ingress and egress gateways.	N/A

Date	Type	Change	Relevant topic
	New feature	Istio CoreDNS is supported.	N/A
	New feature	Support for using existing ECS instances to create a managed Kubernetes cluster is added.	Create a managed Kubernetes cluster.
November 2018	New region	The cluster type of managed Kubernetes is released for the Asia Pacific SE 5 (Jakarta) region.	N/A
	Commercialized product	The Terway network plugin is released. This type of network plugin has higher performance than the Flannel network plugin . If you select Terway as your cluster network plugin, containers within the cluster can directly communicate with each other through the Elastic Network Interface (ENI).	Terway network plugin.
	New feature	Added a performance metric thumbnail of each Worker node to easily check the node status.	N/A

Date	Type	Change	Relevant topic
	New feature	Support for multiple existing nodes to be added to a cluster at one time is provided.	N/A
	New feature	Support for updating Kubernetes cluster certificates that are about to expire is provided.	N/A
October 2018	New region	The Shenzhen Financial Cloud region is released. Alibaba Cloud Container Service for Kubernetes is available in Shenzhen Financial Cloud and is guaranteed by the security and compliance of the Financial Cloud.	N/A
	New feature	The managed Kubernetes is released.	N/A
	New feature	Application management is optimized. You can manage deployment versions of applications. For example, you can roll back to an earlier version.	N/A

Date	Type	Change	Relevant topic
	New feature	Istio can be integrated with Alibaba Cloud Container Service for Kubernetes.	For more information, see Overview .
September 2018	New feature	Alibaba Cloud Container Service for Kubernetes 1.11 is available. <ul style="list-style-type: none"> · New feature includes CRD upgrade, CoreDNS GA, pod priority setting, and preemptive scheduling. · You can select v1.10 or v1.11 as needed. · You can configure multiple containers and create StatefulSet applications in the Container Service console. 	Create a StatefulSet application by using an image
	New feature	With Alibaba Cloud Container Service for Kubernetes, you can pull an image without a password from Container Registry.	N/A
	New feature	Automatic node scaling is available.	Cluster auto scaling
	New feature	Preemptible instances are available.	N/A

Date	Type	Change	Relevant topic
August 2018	New feature	The managed Kubernetes is released for beta.	Create a managed Kubernetes cluster
	New feature	The Istio addon plugin is available.	N/A
July 2018	New region	The Australia region is available for Alibaba Cloud Container Service for Kubernetes.	N/A
	New feature	Gray releases and batch releases are available.	For more information, see Gray releases and blue/green deployment and Use batch release on Alibaba Cloud Container Service for Kubernetes .
June 2018	New region	Asia Pacific NE 1 (Tokyo) and China North 5 (Hohhot) are available.	N/A
	New feature	Alibaba Cloud Container Service for Kubernetes 1.10 supports FPGA and HugePages.	N/A
	New feature	The monitoring function is upgraded. Specifically, application monitoring and alarm notifications can be configured.	Integration and usage with CloudMonitor
	New feature	You can create Subscription clusters.	Create a Kubernetes cluster

Date	Type	Change	Relevant topic
	New feature	In Serverless Kubernetes clusters , new kubectl commands, namely , kubectl exec and kubectl attach can be run, and the Ingress function is available.	Kubernetes supported functions
May 2018	New region	Shanghai Financial Cloud is available. Alibaba Cloud Container Service for Kubernetes is available in Shanghai Financial Cloud and is guaranteed by the security and compliance of the Financial Cloud.	N/A
	New feature	Serverless Kubernetes is released in Alibaba Cloud Container Service for Kubernetes.	Serverless Kubernetes cluster
	New feature	Blue/green deployment, gray releases, and A/B testing are available .	For more information, see Gray releases and blue/green deployment .

Date	Type	Change	Relevant topic
April 2018	New region	India (Mumbai), Indonesia (Jakarta), and Dubai are available. Five new regions covering the Southeast Asia, the Middle East, and India are available , and can use v1.9 that is the latest and stable version.	N/A
	New feature	Multi-zone Kubernetes is released for Alibaba Cloud Container Service for Kubernetes . You can create highly available Kubernetes clusters . You can deploy Master nodes in three zones, and deploy Worker nodes in multiple zones. With this new feature, you can make the best use of Kubernetes clusters for production.	Create a multi-zone Kubernetes cluster
	New feature	The service catalog function is updated to support MySQL , RDS, RabbitMQ, and Spark.	For more information, see Overview of service catalog management.

Date	Type	Change	Relevant topic
	New feature	You can manage applications released by using Helm.	Manage a Helm-based release
March 2018	New feature	Alibaba Cloud Container Service for Kubernetes 1.9 is available. With this version, you can customize ECS images. Alibaba Cloud Container Service for Kubernetes supports native Kubernetes 1.9.3. Workloads API is released. CRD is enabled by default . GPU scheduling is supported. In addition, you can customize ECS images when creating a cluster. You can add nodes to the cluster by automatically reset images.	N/A
	New feature	Alibaba Cloud Container Service for Kubernetes has released app catalog with which you can implement one-click application deployment through Helm.	Manage a Helm-based release

Date	Type	Change	Relevant topic
	New feature	Alibaba Cloud Container Service for Kubernetes has released service catalog which supports ServiceBroker.	For more information, see Overview of service catalog management.
	New feature	Alibaba Cloud Container Service for Kubernetes has upgraded resource monitoring. Specifically, you can monitor cluster nodes through CloudMonitor.	Group-based monitoring and alarms
January 2018	New region	The international Websites of both Alibaba Cloud Container Registry and Alibaba Cloud Container Service for Kubernetes are available.	N/A
	New feature	Alibaba Cloud Container Service for Kubernetes 1.8.4 is released to provide security upgrade and auto scaling.	Cluster auto scaling
	New feature	FlexVolume is released and supports cloud disks, NAS, and OSS	Use Alibaba Cloud cloud disk volumes Use NAS file systems of Alibaba Cloud Use Alibaba Cloud OSS volumes

Date	Type	Change	Relevant topic
	New feature	Network functions is upgraded to support network policies and bandwidth limits.	Access services by using Server Load Balancer
	New feature	ECS Bare Metal (EBM) instances are available.	N/A
October 2017	New feature	The native Kubernetes version is now available. Alibaba Cloud Container Service for Kubernetes 1.8.1 is available.	What is Container Service for Kubernetes
	New feature	Alibaba Cloud Container Service blockchain solution is in beta.	For more information, see Overview of blockchain solution.
August 2017	New feature	Alibaba Cloud Container Service for Kubernetes 1.7.2 is available.	N/A


9.2 Changes to Cloud Controller Manager

This topic describes the latest changes to Cloud Controller Manager.

Date	Image address	Change	Relevant topic
April 15, 2019	registry.cn-hangzhou.aliyuncs.com/acs/cloud-controller-manager-amd64:v1.9.3.105-gfd4e547-aliyun	<ul style="list-style-type: none">· Added the feature that multiple route tables can be set for a Kubernetes cluster in a VPC by using a configuration file· Fixed the bug that caused update configurations using HTTP to fail.	Access services by using Server Load Balancer

Date	Image address	Change	Relevant topic
March 20, 2019	registry.cn-hangzhou.aliyuncs.com/acs/cloud-controller-manager-amd64:v1.9.3.81-gca19cd4-aliyun	<ul style="list-style-type: none">· An SLB instance created in a non-Kubernetes cluster can be reused in a managed or dedicated Kubernetes cluster.· A Kubernetes node name can be customized by using Cloud Controller Manager (CCM) in addition to setting the Kubernetes <code>nodeName</code> parameter manually.· The incompatibility between CCM 1.8.4 and Kubernetes 1.11.5 was fixed. We recommend that you upgrade your CCM to the latest version.	Access services by using Server Load Balancer

Date	Image address	Change	Relevant topic
December 26, 2018	registry.cn-hangzhou.aliyuncs.com/acs/cloud-controller-manager-amd64:v1.9.3.59-ge3bc999-aliyun	<p>Multiple Kubernetes services can reuse the same SLB instance.</p> <ul style="list-style-type: none"> The SLB instances created by Kubernetes through a service cannot be reused. Otherwise, the reused SLB instances may be removed by accident. 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, you must use the listener name and the virtual server group name to mark the SLB instance . Do not modify the listener 	Access services by using Server Load Balancer
46		<p>names or virtual server group names.</p> <p>You can modify</p>	Issue: 20190419

Date	Image address	Change	Relevant topic
August 15, 2018	registry.cn-hangzhou.aliyuncs.com/acs/cloud-controller-manager-amd64:v1.9.3.10-gfb99107-aliyun	<ul style="list-style-type: none"> You can use annotation : service . beta . kubernetes . io / alicloud - loadbalancer - master - zoneid to specify the primary zone where an automatically created SLB is located. You can use annotation : service . beta . kubernetes . io / alicloud - loadbalancer - slave - zoneid to specify the secondary zone where an automatically created SLB is located. <div>  Note: This parameter does not take effect in regions where a primary zone or secondary zone SLB cannot be created. </div>	Access services by using Server Load Balancer
Issue: 20190419		<ul style="list-style-type: none"> When you specify an existing SLB, you can use 	

Date	Image address	Change	Relevant topic
June 25, 2018	registry.cn-hangzhou.aliyuncs.com/acs/cloud-controller-manager-amd64:v1.9.3	<ul style="list-style-type: none"> You can set annotation : service . beta . kubernetes . io / alicloud - loadbalancer - backend - label to use the Worker nodes specified by labels as backend servers. You can use annotation : service . beta . kubernetes . io / alicloud - loadbalancer - spec to specify an SLB type. For example, you can set a shared-performance SLB or an exclusive SLB. The service mode externalTraffic: Local is available. Only the node where a pod is located can be added as an SLB backend server. When you add or remove a cluster node, the system automatically processes the SLB backend server. That is, the system adds 	Access services by using Server Load Balancer
48			Issue: 20190419