Alibaba Cloud Alibaba Cloud Container Service for Kubernetes

Product Introduction

Issue: 20190910

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

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

Contents

1 What is ACK?

Container Service for Kubernetes (ACK) is a high-performance and scalable containerized application management service, which enables you to manage the entire lifecycle of enterprise-level containerized applications. ACK simplifies the deployment and scale-out operations of Kubernetes clusters and integrates Alibaba Cloud capabilities of virtualization, storage, networking, and security. Based on these capabilities, ACK provides an ideal runtime environment for Kubernetes containeri zed applications. Container Service is one of the first certified service platforms that obtain the certificate of conformance from Kubernetes in the world. It provides professional support and services related to containers.

ACK provides the following three types of Kubernetes clusters: Dedicated Kubernetes clusters, Managed Kubernetes clusters, and Serverless Kubernetes clusters.

- Dedicated Kubernetes clusters: You must create three master nodes and several worker nodes for a Dedicated Kubernetes cluster to achieve high availability. This type of Kubernetes cluster allows you to manage the cluster infrastructure in a more fine-grained manner. It requires you to plan, maintain, and upgrade the Kubernetes clusters on your own.
- Managed Kubernetes: You only need to create worker nodes. The master nodes are created and managed by Alibaba Cloud Managed Kubernetes. This type of Kubernetes cluster is easy to use with low cost and high availability. You can focus on the business without the need to operate and maintain the Kubernetes cluster master nodes.
- Serverless Kubernetes: You do not need to create and manage any master nodes or worker nodes. You can use the Container Service console or command line interface to configure resources for containers, specify container images for applications, provide methods for external access, and start applications.

Features

- Cluster management
  - You can create a Dedicated Kubernetes cluster, Managed Kubernetes cluster, and Serverless Kubernetes cluster in an easy way through the Container Service console. GPU instances and ECS Bare Metal (EMB) instances can be created in

the Kubernetes cluster. You can also create clusters across multiple zones to achieve high availability.

- ACK provides optimized OS images of Kubernetes containers and offers Kubernetes clusters and Docker versions with high stability and reinforced security.
- ACK supports multi-cluster management, high-availability clusters across zones , and cluster federation management.

- All-in-one container management
  - Network

    Provides high-performance VPC and elastic network interface (ENI) plug-ins . The performance of these plug-ins is 20% better than the performance of regular network solutions.

    Supports container access policies and traffic controls.
  - Storage

    Supports Alibaba Cloud disks, Apsara File Storage NAS, and Object Storage Service (OSS) buckets. FlexVolume drivers are also provided.

    Supports dynamic creation and migration of storage volumes.
  - Logs

    Supports log collection and integration into Log Service.

    Supports the integration with third-party open source logging solutions.
  - Monitoring

    Supports monitoring containers at both the container level and virtual machine (VM) level. You can also integrate Container Service with third-party open source monitoring solutions.
  - Access control

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

    Supports permission configuration and management at the level of applications.
  - Applications

    Supports canary release and blue-green release.

    Supports application monitoring and scaling.

    Supports the built-in application store and deploys applications in an easy way by using Helm.

    Supports Service Catalog to simplify cloud service integration.

- **High-availability scheduling policies and simple delivery processes**
  - **Supports affinity policies and horizontal scaling of services.**
  - **Provides high availability and disaster recovery across zones.**
  - **Provides API operations for cluster and application management to easily implement continuous integration and private system deployment.**

2 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	· High-performance Virtual Private Cloud (VPC) network plug-in. · Supports network policy and flow control. Container Service provides you with continuous network integration and the best network optimization.

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

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

3 Architecture

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.

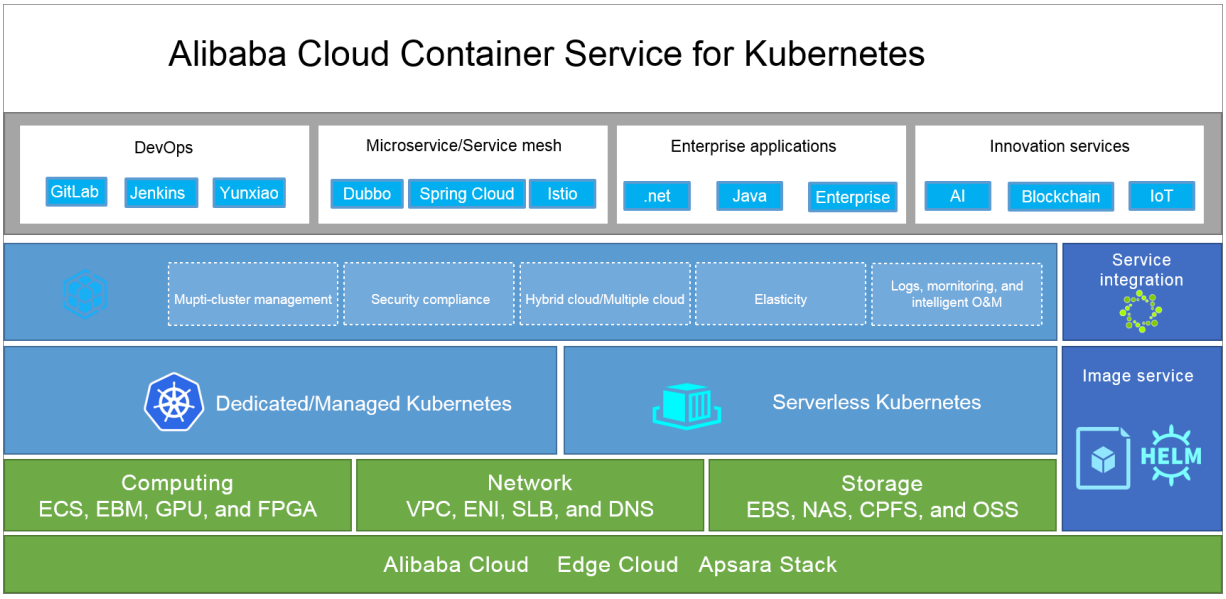

Feature	Description
Multiple Kubernetes clusters	Integrated with Alibaba Cloud virtualization technology, Container Service for Kubernetes supports three types of Kubernetes clusters: dedicated Kubernetes cluster, managed Kubernetes cluster, and serverless Kubernetes cluster. · Dedicated Kubernetes cluster: 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. · Managed Kubernetes cluster: Alibaba Cloud Container Service for Kubernetes hosts the Master nodes of a managed Kubernetes cluster. You do not need to to operate and maintain the Master nodes of the Kubernetes cluster. This type of Kubernetes cluster is easy to use, low-cost, and highly available. · Serverless Kubernetes cluster: 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.

Feature	Description
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.

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.

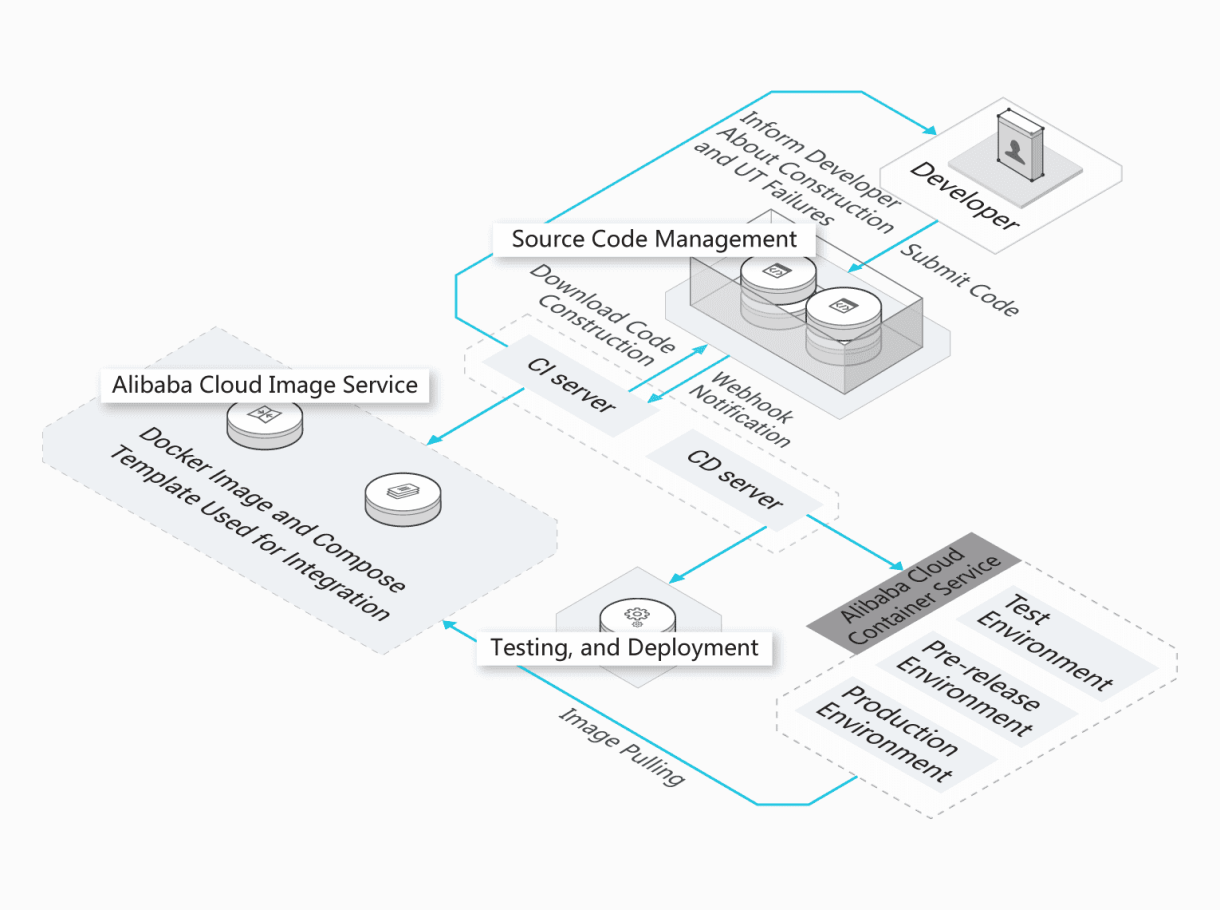

Microservice architecture

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.

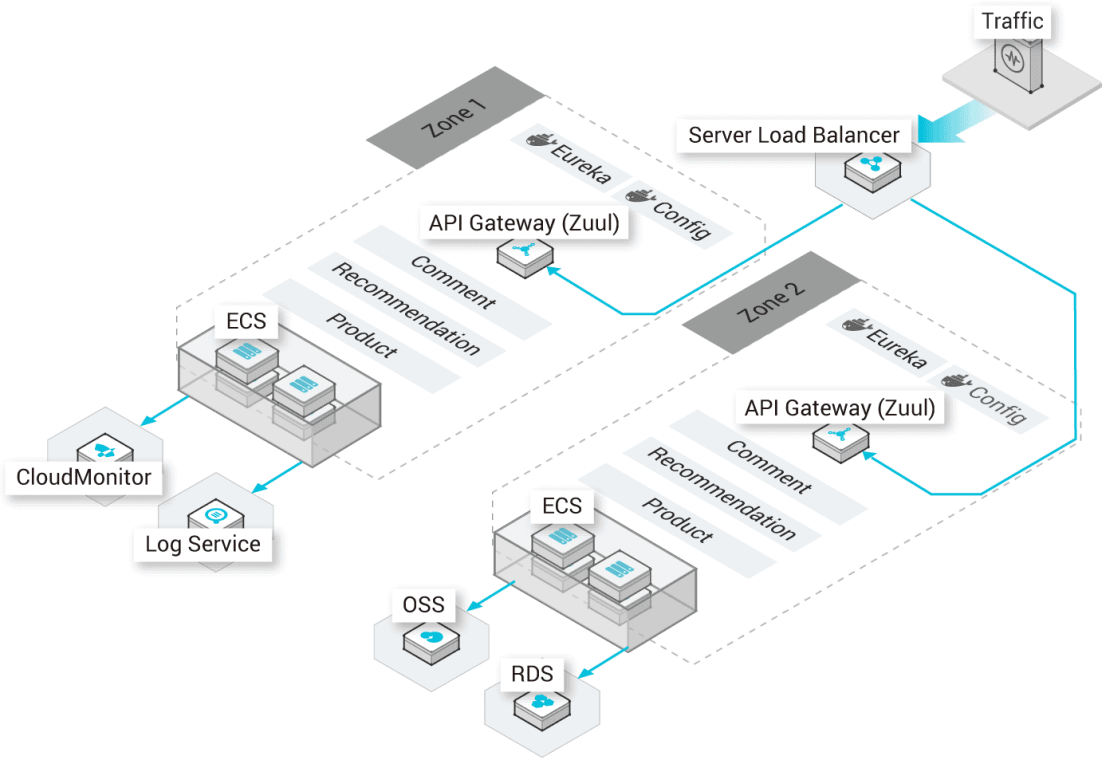

Hybrid cloud architecture

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.

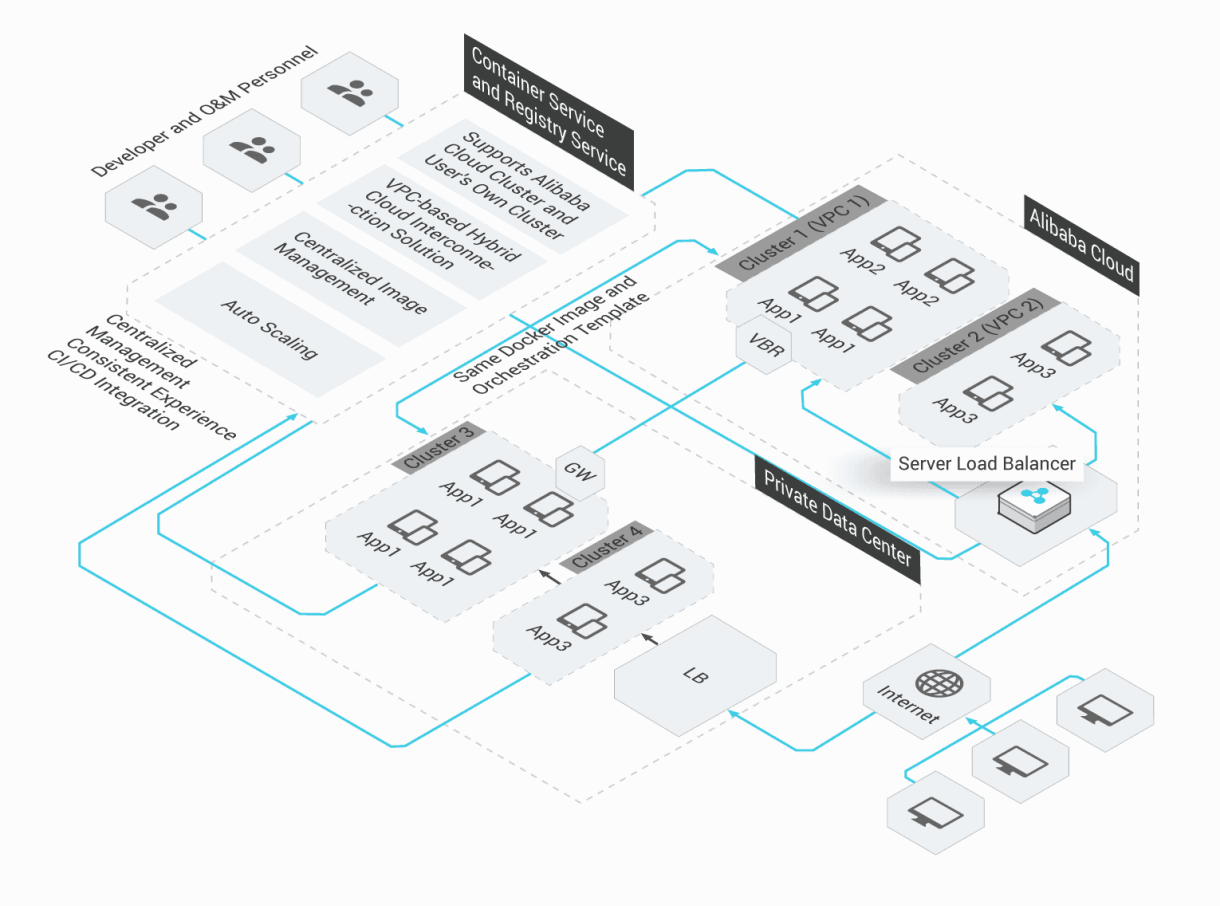

Auto scaling architecture

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.

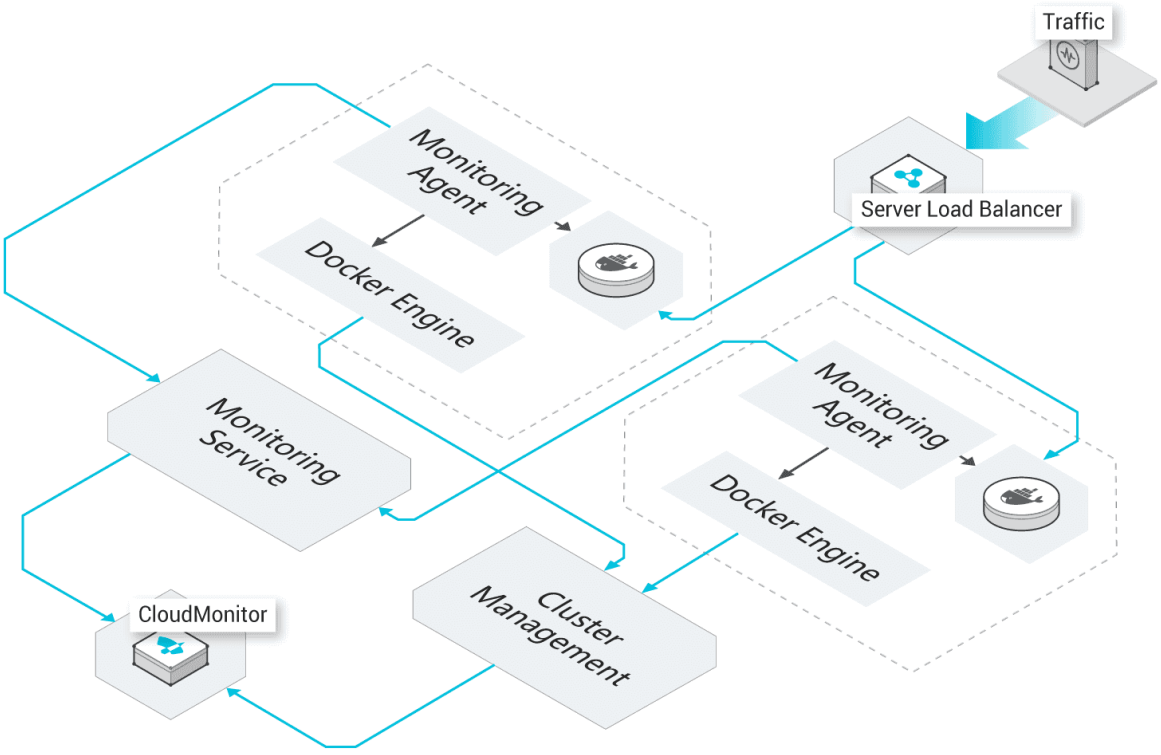

5 Limits

This topic describes the limits on Alibaba Cloud Container Service for Kubernetes product features and service performance, and also provides the method to apply for a higher quota.

After you create a Kubernetes cluster, the following limits apply:

- After a VPC is set for the cluster, it cannot be replaced with another VPC.
- The Kubernetes cluster cannot be converted from the singe-zone type to the multi-zone type.
- The Kubernetes cluster cannot be converted from a dedicated to a managed Kubernetes cluster type.
- No Internet Server Load Balancer (SLB) instance can be set for the cluster.

   Note:
  To expose the API server of a Kubernetes cluster to the Internet, you must set an Internet SLB instance when you set parameters for the cluster creation.
- No CloudMonitor plugin can be installed for the cluster.

   Note:
  A CloudMonitor plugin can be installed for the cluster only when you set parameters for the cluster creation.
- The network plugins of containers cannot be changed, otherwise the Kubernetes cluster cannot function properly.

In addition to the preceding limits, the following limits apply to resources related to the Kubernetes clusters:

- An application that is created in a Kubernetes cluster cannot be migrated across different namespaces.
- The Alibaba Cloud Container Service for Kubernetes accesses the Internet through an API when a Kubernetes cluster is created. For this to function properly, you must set a NAT gateway to create a Kubernetes cluster. This is applicable to both dedicated and managed Kubernetes clusters.
- ECS instances created for a Kubernetes cluster support both the Pay-As-You-Go and Subscription billing methods. You can change the billing method of an instance in

the ECS console. However, other cluster resources (for example, an SLB instance) only support the Pay-As-You-Go billing method.

- According to the ROS rule, automatic cluster scaling only removes nodes that you added through scaling out the cluster. Automatic cluster scaling does not remove the nodes that were added at the time which you created the cluster, or the nodes that you added to the cluster after you created the cluster. Nodes are removed from the cluster in the sequence that they were added when the cluster was scaled out.

Kubernetes cluster quotas

The following table details further limits of the Container Service for Kubternetes and states whether you can open a ticket to request changing the limits.

Note:

If you are a general user as opposed to an administrator, you must complete real-name authentication when you apply for more resource quotas.

Item	Limit	Can I open a ticket to change the limit?
The maximum number of dedicated Kubernetes clusters and managed Kubernetes clusters that one Alibaba Cloud account can create.	10	Yes
The maximum number of nodes (including both Master nodes and Worker nodes) in one Kubernetes cluster.	40	Yes
The maximum number of container instances that can be created on one Worker node.	256	No
The maximum number of serverless Kubernetes clusters that one Alibaba Cloud account can create.	2	Yes

Item	Limit	Can I open a ticket to change the limit?
The maximum number of container instances that can be created in one serverless Kubernetes cluster.	300	Yes

6 Kubernetes versions

This topic describes which Kubernetes versions are supported, the release cycle of Kubernetes versions, and the validity period of Kubernetes versions on Alibaba Cloud.

Supported Kubernetes versions

Alibaba Cloud Container Service for Kubernetes (ACK) supports four Kubernetes versions: V1.12, V1.11, V1.10, and V1.9. With Kubernetes V1.12 or V1.11, you can create clusters in the ACK console, and perform the corresponding O&M operations. However, with Kubernetes V1.10 and V1.9, you can only perform O&M operations for the clusters that were created in these two versions, and no new clusters can be created in these two versions.

Release cycle

- The Kubernetes Community releases a major Kubernetes version every three months.
- After the Kubernetes Community releases a major Kubernetes version, ACK examines and tests the Kubernetes version within two or three months later, and then releases the corresponding Kubernetes version to the Container Service console where you can upgrade your Kubernetes version.

Validity periods

- Each Kubernetes version released by ACK is valid for one year (starting from the date when it was released). For example, Kubernetes V1.12 was released on April 1, 2019, therefore this version will expire on March 31, 2020. You can view the validity period of a Kubernetes version on the page of upgrading a Kubernetes cluster.
- After a Kubernetes version is released to all the regions of Alibaba Cloud, ACK provides the function for you to upgrade your current Kubernetes version to the newly released version within two or three weeks. Then, you can upgrade the version of your Kubernetes on the page for upgrading Kubernetes clusters. However, if your current Kubernetes version expires, it cannot be upgraded through the page for upgrade Kubernetes clusters.