

Alibaba Cloud Elastic Compute Service

Product Introduction

Issue: 20190718

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>switch {stand slave}</code>

Contents

Legal disclaimer.....	I
Generic conventions.....	I
1 What is ECS?.....	1
2 Benefits of ECS.....	6
3 Scenarios.....	11
4 Regions and zones.....	12
5 Limits.....	16
6 Usage notes.....	26

1 What is ECS?

Elastic Compute Service (ECS) is a high-performance, stable, reliable, and scalable IaaS-level service provided by Alibaba Cloud. It is more convenient and efficient compared with physical servers. You can immediately acquire ECS instances and scale computing resources on-demand. ECS provides a variety of instance types that suit various business needs and help boost business growth.

Why ECS?

ECS provides the following benefits:

- You do not have to purchase hardware or construct data centers up front.
- Instances are delivered within minutes, enabling rapid deployment and reducing time to market.
- You can make use of resources in data centers and BGP grids around the world.
- You can scale and release resources based on actual business needs.
- Provides heterogenous computing servers such as GPU, FPGA, bare metal servers, and general x86 servers.
- You can use ECS to access other Alibaba Cloud services over the internal network, reducing traffic costs.
- Provides a host of security solutions, such as virtual firewalls, permission control, internal network isolation, virus protection, and traffic throttling.
- Provides a performance monitoring framework and proactive O&M systems.
- Provides standardized APIs to improve ease of use and applicability.

For more benefits, see [Benefits of ECS](#) and [Scenarios](#).

Architecture

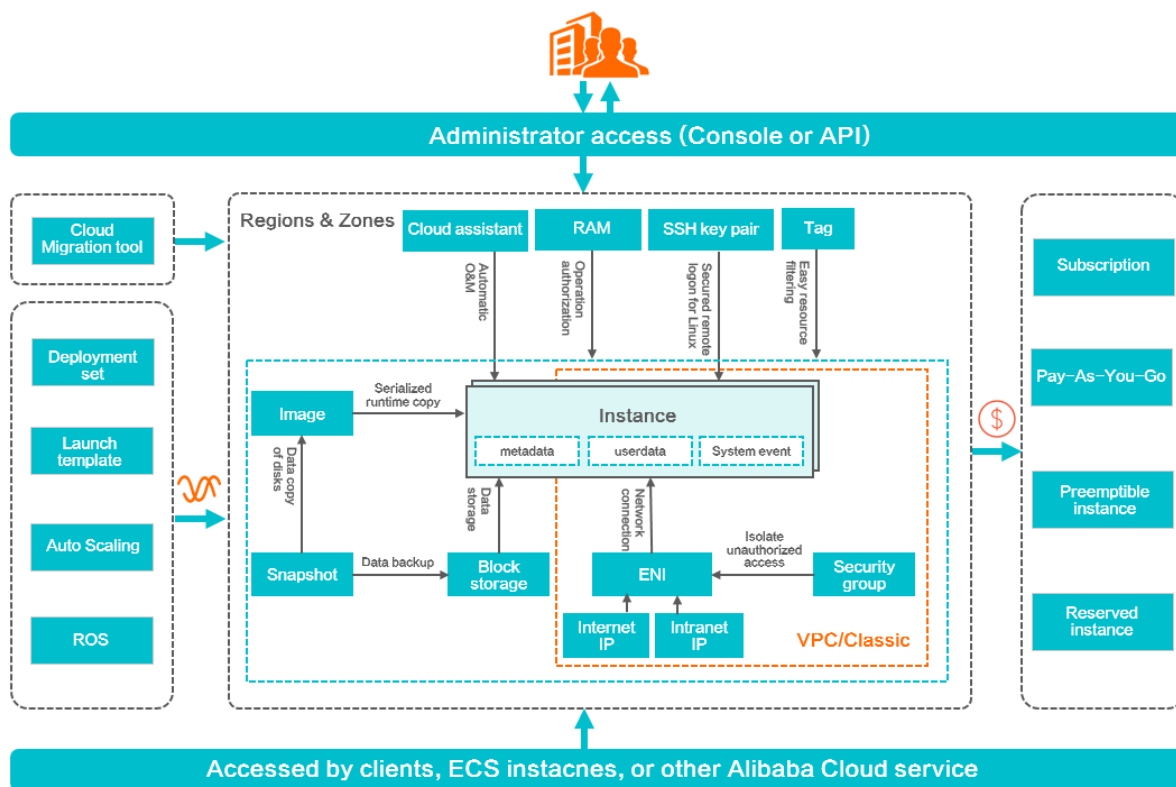
ECS comprises the following major components:

- **Instance**: a virtual computing environment that includes basic computing components such as CPU, memory, operating system, bandwidth, and disks. The computing performance, memory specifications, and applicable scenarios of an instance are determined by its [instance type](#). The performance of an instance is determined by its vCPUs, memory capacity, and network performance.

- **Image**: provides the operating system, initial application data, and pre-installed software for instances. Multiple Linux operating systems and Windows Server operating systems are supported.
- **Block storage**: a block storage device which features high performance and low latency. Distributed storage architecture-based **cloud disks**, and physical storage-based **local disks** can be used.
- **Snapshot**: a stateful data file of a cloud disk or shared block storage at a certain point in time. It is often used to back up and restore data, and to create custom images.
- **Security group**: a logical group of instances located in the same region that have the same security requirements and require access to each other. A security group works as a virtual firewall for the ECS instances inside it.
- **Network**:
 - **VPC**: a logically isolated private cloud network. You can configure a private IP address range, a route table, and a gateway for a VPC.
 - **Classic network**: all classic instances are built on a shared infrastructure network, which is planned and managed in a centralized manner.

For more information, see [Elastic Compute Service](#).

The following figure shows the architecture of ECS components. For more information about the functional components in the figure, see the related documentation.



Pricing

ECS supports multiple billing methods, such as subscription, pay-as-you-go, reserved instances (beta phase), and preemptible instances. For more information, see [Pricing overview](#) and [Product pricing](#).

Management tools

After registering an Alibaba Cloud account, you can create, use, or release an ECS instance in any region by using the following methods provided by Alibaba Cloud:

- **ECS console:** a Web service page used to manage ECS instances. For more information, see [Quick reference](#).
- **ECS APIs:** RPC APIs that support GET and POST requests. For more information, see [API Reference](#). The following developer tools can be used to call ECS APIs:
 - **CLI:** a flexible and scalable management tool based on Alibaba Cloud APIs. You can use CLI to encapsulate Alibaba Cloud native APIs to develop custom features.
 - **OpenAPI Explorer:** allows you to retrieve APIs, call APIs, and dynamically generate SDK example code.
 - **Alibaba Cloud SDK:** provides SDKs for multiple programming languages such as Java, Python, and PHP.

- **Resource Orchestration Service**: automatically creates and configures Alibaba Cloud resources based user-defined templates.
- **Terraform**: an open-source tool that uses configuration files to call computing resources on Alibaba Cloud and other platforms that support Terraform. Terraform also implements version control.

Deployment tips

Before you purchase an ECS instance, consider the following factors:

- Regions and zones

A region represents an Alibaba Cloud data center. The region and zone determine the physical location of an ECS instance. After an instance is created, its **metadata** is established and its region cannot be changed. Select a region and zone based on the target geographical location, availability of Alibaba Cloud services, application availability requirements, and whether internal network communication is required. For example, if you want to access both ECS and **ApsaraDB for RDS**, the RDS instance and ECS instance must be in the same region. For more information, see **Regions and zones**.

- High availability

To ensure business consistency and continuity, we recommend that you use **snapshots** to back up data, and use multi-zone deployment, **deployment sets**, and **Server Load Balancer (SLB)** for disaster recovery.

- Network planning

We recommend that you use **VPC** to plan your own private IP addresses. VPC supports all new features and new instance types. It also supports business system isolation and multi-region system deployment.

- Security solutions

You can use ECS security groups to control inbound and outbound access policies and the port listening status of ECS instances. For applications deployed on ECS, you can use free **Anti-DDoS Basic**. Alibaba Cloud Security is also available.

- For example, you can use Anti-DDoS Pro to ensure the stability and reliability of source sites. For more information, see **Anti-DDoS Pro documentation**.
- Server Guard can be used to safeguard the security of ECS instances. For more information, see **Server Guard documentation**.

Learning path

[ECS Learning Path](#) provides a sequence of courses to help you master the basics of ECS and gain advanced knowledge.

Related services

Together with ECS, you can select the following Alibaba Cloud services:

- [Auto Scaling](#) automatically adjusts the number of ECS instances based on business and policy changes.
- [Container Service](#) manages application lifecycles on groups of ECS instances.
- [Server Load Balancer \(SLB\)](#) distributes traffic among multiple ECS instances.
- [CloudMonitor](#) develops monitoring solutions for instances, system disks, and public network bandwidth.
- [ApsaraDB for RDS](#) provides database services accessible over internal networks to ECS instances, reduces network latency, access fees, and delivers top-notch performance. ApsaraDB for RDS supports multiple database engines, including MySQL, SQL Server, PostgreSQL, PPAS, and MariaDB.
- [Alibaba Cloud Marketplace](#): a platform to purchase software infrastructure, business software, website construction, hosted O&M, security, data & APIs, and solutions provided by third-party partners. You can also provide software programs and services as a service provider. For more information, see [Alibaba Cloud Marketplace documentation](#).

For more solutions, see [Alibaba Cloud Solutions](#).

2 Benefits of ECS

ECS provides better availability, security, and elasticity than Internet Data Centers (IDCs) and server vendors.

Availability

Alibaba Cloud adopts more stringent IDC standards, server access standards, and O & M standards to guarantee data reliability and high availability of cloud computing infrastructure and cloud servers.

In addition, each Alibaba Cloud region consists of multiple zones. For greater fault tolerance, you can build active/standby or active/active services in multiple zones. For a finance-oriented solution with three IDCs in two regions, you can build fault tolerant systems in multiple regions and zones. Those services include disaster tolerance and backup, which are supported by the mature solutions built by Alibaba Cloud.

Switching between services is smooth within the Alibaba Cloud framework. For more information, see [E-Commerce Solutions](#). Alibaba Cloud industry solutions support a variety of services, such as finance, E-commerce, and video services.

Alibaba Cloud provides you with the following support services:

- Products and services for availability improvement, including cloud servers, Server Load Balancer, multi-backup databases, and Data Transmission Services (DTS).
- Industry partners and ecosystem partners that help you build a more advanced and stable architecture and guarantee service continuity.
- Diverse training services that enable you to connect with high availability from the business end to the underlying basic service end.

Security

For cloud computing users, security and stability are priorities. Alibaba Cloud has passed a host of international information security certifications, including ISO 27001 and MTCS, which demand strict confidentiality of user data and user information, as well as user privacy protection. We recommend that you use ECS in an [Alibaba Cloud Virtual Private Cloud \(VPC\)](#).

- Alibaba Cloud VPC offers more business possibilities.

You only need to perform a simple configuration to connect your business environment to global IDCs, making your business more flexible, stable, and extensible.

- Alibaba Cloud VPC can connect to your IDC through a leased line to build a hybrid cloud architecture.

You can build a more flexible business with robust networking derived from Alibaba Cloud's various hybrid cloud solutions and network products. A superior business ecosystem is made possible with Alibaba Cloud's ecosystem.

- Alibaba Cloud VPC is more stable and secure.

Stable: After you build your business on VPC, you can update your network architecture and obtain new network functions on a daily basis as the network infrastructure evolves constantly, allowing your business to run steadily. You can divide, configure, and manage your network on VPC according to your needs.

Secure: VPC features traffic isolation and attack isolation to protect your services from attack traffic on the Internet. By building your business on VPC, the first line of defense is established.

VPC provides a stable, secure, fast-deliverable, self-managed, and controllable network environment. VPC hybrid cloud brings the technical advantages of cloud computing to traditional industries, in addition to industries and enterprises that are not engaged in cloud computing.

Elasticity

Elasticity is a key benefit of cloud computing. By using Alibaba Cloud, you can have all the IT resources necessary to get an IT company of medium size provisioned within minutes. The available resources and capacities can meet the requirements of most companies, allowing their applications built on the cloud to handle a huge volume of transactions without problems.

- Elastic computing
 - Vertical scaling involves modifying the configurations of a server.

After you purchase ECS or storage capacity of Alibaba Cloud, you can configure your server with great flexibility based on your actual transaction volume,

whereas it may be difficult to change configurations in the traditional IDC model. For more information about vertical scaling, see [Change configurations](#).

- Horizontal scaling

Horizontal scaling allows the re-division of resources between applications. For example, at peak hours for game or live video streaming apps, in the traditional IDC model, your hands may be tied if additional resources are required when you are already at full capacity. Cloud computing uses elasticity to provide additional resources to you over that period. When the period ends, you can release unnecessary resources to reduce your business costs. By using both horizontal scaling and auto-scaling that Alibaba Cloud provides, you can determine how and when you scale your resources or apply your scaling based on business loads. For more information about horizontal scaling, see [Auto Scaling](#).

- Elastic storage

Alibaba Cloud provides elastic storage. In the traditional IDC model, if more storage space is required, you can only add servers, but the number of servers that you can add is limited. In the cloud computing model, however, the sky is the limit. Order as much storage space as you need to meet business demand. For more information about elastic storage, see [Resize a disk](#).

- Elastic network

Alibaba Cloud features elastic network as well. When you purchase Alibaba Virtual Private Cloud (VPC), you can set network configurations to be the same as those of data centers. In addition, VPC has the following benefits: interconnection between data centers, separate secure domains in data centers, and flexible network configurations and planning within the VPC. For more information about elastic networks, see [Virtual Private Cloud](#).

Alibaba Cloud incorporates elasticity in computing, storage, network, and business architecture design. By using Alibaba Cloud, you can build your business portfolio in any way you want.

Comparison between ECS and traditional IDCs

The following table provides the benefits of ECS compared with traditional IDCs.

Item	ECS	Traditional IDC
Equipment rooms	Provides independently developed DC-powered servers with low PUE.	Provides traditional AC-powered servers with high PUE.
	Provides backbone equipment rooms with high outbound bandwidth and dedicated bandwidth.	Provides equipment rooms with various quality levels and shared bandwidth primarily, difficult for users to choose.
	Provides multiline BGP equipment rooms, enabling smooth and balanced access throughout a country.	Provides equipment rooms with single or dual line primarily.
Ease of operation	Provides built-in mainstream OSs, including activated Windows OS.	Purchases and installs OS manually.
	Switches OS online.	Reinstalls OS manually.
	Provides a Web-based console for online management.	Users must manage and maintain network manually.
	Provides mobile phone verification for password setting, increasing data security.	Has difficulty in resetting passwords, and exposes high risk of password cracking.
Disaster recovery and backup	Each data segment has multiple copies. When one copy is corrupted, the data can be quickly restored.	Users must build disaster recovery environment by themselves, and use traditional storage devices.
	Users can customize automatic snapshot policies to create automatic snapshots for data recovery.	Users must restore all corrupted data manually.
	Faults can be recovered fast and automatically.	Faults cannot be recovered automatically.
Security and reliability	Effectively prevents MAC spoofing and ARP attacks.	Fails to prevent MAC spoofing and ARP attacks.

Item	ECS	Traditional IDC
	Effectively defends against DDoS attacks by using black holes and cleaning traffic.	Needs additional costs for devices for traffic cleaning and black hole shielding systems.
	Provides additional services, such as port scanning, Trojan scanning, and vulnerability scanning.	Typically encounters problems such as vulnerability, Trojan, and port scanning.
Flexible scalability	Activates cloud servers on demand and upgrades configurations online.	Needs a long time for server delivery.
	Adjusts outbound bandwidth as required.	One-off purchase of outbound bandwidth, unable to adjust.
	Combines with Server Load Balancer online, enabling scaling up applications quickly and easily.	Uses hardware-based server load balancing, which is expensive and extremely difficult to set up.
Cost effectiveness	Low cost.	High cost.
	Small up-front investment.	Large up-front investment, possible waste of resources.
	Purchases on demand and pay as you go, meeting requirements for constant business changes.	Purchases up front to meet configuration requirements for peak hours.

3 Scenarios

ECS is a highly flexible solution. It can be used independently as a simple web server, or used with other Alibaba Cloud products, such as Object Storage Service (OSS) and Content Delivery Network (CDN), to provide advanced solutions.

ECS can be used in the following applications.

Official corporate websites and simple web applications

During the initial stage, corporate websites have low traffic volumes and require only low-configuration ECS instances to run applications, databases, storage files, and other resources. As your business expands, you can upgrade the ECS configuration and increase the number of ECS instances at any time. You no longer need to worry about insufficient resources during peak traffic.

Multimedia and large-traffic apps or websites

ECS can be used with OSS to store static images, videos, and downloaded packages, reducing storage fees. In addition, ECS can be used with CDN or Server Load Balancer to greatly reduce user access waiting time, reduce bandwidth fees, and improve availability.

Databases

A high-configuration I/O-optimized ECS instance can be used with an SSD cloud disk to support high I/O concurrency with higher data reliability. Alternatively, multiple lower-configuration I/O-optimized ECS instances can be used with Server Load Balancer to deliver a highly available architecture.

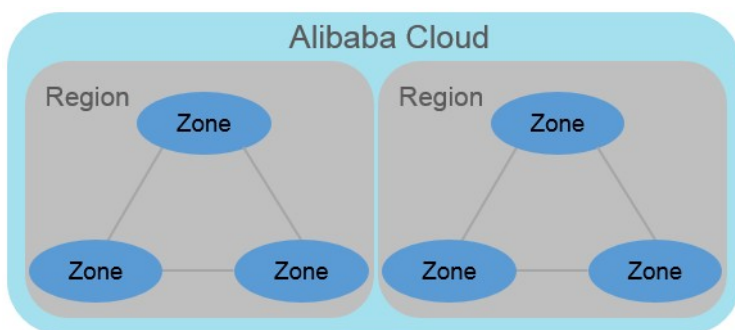
Apps or websites with large traffic fluctuations

Some applications may encounter large traffic fluctuations within a short period. When ECS is used with Auto Scaling, the number of ECS instances is automatically adjusted based on traffic. This feature allows you to meet resource requirements while maintaining a low cost. ECS can be used with Server Load Balancer to deliver a high availability architecture.

4 Regions and zones

This topic provides a complete list of Alibaba Cloud regions and zones.

Regions in Alibaba Cloud are independent. Zones in Alibaba Cloud are completely isolated. However, zones in the same region are connected through links with low latency. The following figure shows the relationship between regions and zones in Alibaba Cloud.



Regions

Alibaba Cloud regions are physical locations (data centers) that spread all over the world to reduce the network latency. The region cannot be changed once a resource is created in it. The following table lists all Alibaba Cloud regions, corresponding cities, and Region IDs.



Note:

The available regions vary according to the product. To view the list of regions available for each product, see [Alibaba Cloud Global Infrastructure](#).

• Regions in Mainland China

Region name	City	Region ID	Number of zones
China North 1	Qingdao	cn-qingdao	2
China North 2	Beijing	cn-beijing	8
China North 3	Zhangjiakou	cn-zhangjiakou	2
China North 5	Hohhot	cn-huhehaote	2
China East 1	Hangzhou	cn-hangzhou	8
China East 2	Shanghai	cn-shanghai	6
China South 1	Shenzhen	cn-shenzhen	5

Region name	City	Region ID	Number of zones
China SW 1	Chengdu	cn-chengdu	1

- International regions

Region name	City	Region ID	Number of zones
Hong Kong	Hong Kong	cn-hongkong	2
Asia Pacific SE 1	Singapore	ap-southeast-1	3
Asia Pacific SE 2	Sydney	ap-southeast-2	2
Asia Pacific SE 3	Kuala Lumpur	ap-southeast-3	2
Asia Pacific SE 5	Jakarta	ap-southeast-5	2
Asia Pacific SOU 1	Mumbai	ap-south-1	2
Asia Pacific NE 1	Tokyo	ap-northeast-1	2
US West 1	Silicon Valley	us-west-1	2
US East 1	Virginia	us-east-1	2
EU Central 1	Frankfurt	eu-central-1	2
UK (London)	London	eu-west-1	2
Middle East 1	Dubai	me-east-1	1

When choosing a region, consider the following factors:

- Geographic location

Choose a region based on the geographic locations of you and your target users.

- Regions in mainland China

In general, we recommend that you choose the data center closest to your target users so that they have fast access. Alibaba Cloud's data centers of the regions in mainland China are similar to each other in terms of infrastructure, BGP network quality, service quality, ECS operation and configuration. The regions offer BGP backbone network lines covering all provinces, autonomous regions, and municipalities to guarantee fast, stable access within mainland China.

- International regions

Alibaba Cloud's international regions are data centers outside mainland China. They offer access at international bandwidth, targeting areas outside mainland China. Users in mainland China that access services deployed in these regions

may experience high latency. Therefore, we do not recommend deploying services in international regions if users from mainland China require access.

- If you have business operations in Hong Kong or Southeast Asia, choose the region Hong Kong, Asia Pacific SE 1, Asia Pacific SE 3, or Asia Pacific SE 5.
- If you have business operations in Japan and Korea, choose the Asia Pacific NE 1.
- If you have business operations in India, choose the region Asia Pacific SOU 1.
- If you have business operations in Australia, choose the region Asia Pacific SE 2.
- If you have business operations in America, choose the region US East 1 or US West 1.
- If you have business operations in Continental Europe, choose the region EU Central 1.
- If you have business operations in the Middle East, choose the Eastern Middle East 1 region.

- Intranet communication between Alibaba Cloud products across regions

Intranet communication between Alibaba Cloud products that are not in the same region is not supported, which means:

- ECS instances and other products, such as ApsaraDB for RDS and OSS instances , cannot communicate with one another over the intranet if they are in different regions.
 - Server Load Balancer cannot be deployed for ECS instances in different regions.
- Resource pricing

The price of resources may vary according to the region. For details, see the [product pricing page](#).

- Business license filing

When you choose a region, consider the special requirements of some areas. If you have purchased an ECS instance in mainland China and used it for a web server, complete business license filing first.

To complete business license filing, pay attention to the following points:

- If your company is located in Beijing, choose the China North 2 region for the ECS instance you have purchased.
- If your company is located in Guangdong, choose the China South 1 region for the ECS instance you have purchased.



Note:

The communications administrations in different provinces, autonomous regions, and municipalities have different approval requirements for business license filing. In case of any change, the information published on the business license filing website of your local administration prevails.

Zones

Zones are physical areas with independent power grids and networks within one region. The network latency for instances within the same zone is lower.

Intranet communication can take place between instances in different zones of the same region, and fault isolation can be achieved between zones. Whether or not instances can be deployed in the same zone depends on the requirements for disaster recovery and network latency.

- If your applications require high disaster recovery capabilities, we recommend that you deploy your instances in different zones of the same region.
- If your applications require low network latency between instances, we recommend that you create your instances in the same zone.

For more information about regions and zones, see the [Alibaba Cloud Global Infrastructure](#).

5 Limits

This topic describes the limits on ECS product features, service performance, and the method for applying for a higher quota.

Overview

The limits of ECS instances are as follows:

- ECS does not support virtual application installation or revirtualization (such as installation of VMware Workstation). Currently, only [ECS Bare Metal Instances](#) and [super computing clusters](#) support revirtualization.
- ECS does not support sound card applications.
- ECS does not support external hardware devices (such as dongles, USB drives, or external hard drives). Instead, it supports a software protection dongle or two-step verification with dynamic passwords.
- ECS does not support IP address translation services such as SNAT. Instead, it supports a VPN or proxy.
- ECS does not support multicast protocols. We recommend that you use unicast instead.
- Currently, Log Service does not support 32-bit Linux ECS instances. For information about regions that support Log Service, see [Service endpoint](#). For information about operating systems that support Log Service, see [Overview](#).

In addition to the preceding limits, the following table details further limits of ECS and whether you can open a ticket to request changing the limits.

ECS instances

Item	Limit	Can I open a ticket to change the limit?
Permission to create instances	Complete real-name registration before creating ECS instances in any Mainland China regions.	No
Instance types for which you can create Pay-As-You-Go instances	Instance types with less than 16 vCPUs	Yes

Item	Limit	Can I open a ticket to change the limit?
vCPU quota of Pay-As-You-Go instances in each region for a single account	50 vCPUs You can click Privileges & Quota Management on the Overview page of the ECS console to view the quota.	Yes
vCPU quota of preemptible instances in each region for a single account	Up to 50 vCPUs are allowed after you open a ticket to obtain the permission. You can click Privileges & Quota Management on the Overview page of the ECS console to view the quota.	Yes
Quota of launch templates in each region for a single account	30	No
Quota of versions of a single launch template	30	No
Switch from Pay-As-You-Go to Subscription	The following instance types (families) are not supported: t1, s1, s2, s3, c1, c2, m1, m2, n1, n2, and e3.	No
Switch from Subscription to Pay-As-You-Go	<ul style="list-style-type: none"> Whether this feature is supported is based on your ECS instance resource usage. 5,000 vCPUs*hours per month. For more information, see Conversion notes. The maximum refund per month is limited, and is displayed on the switch page. 	No

Block storage

Item	Limit	Can I open a ticket to change the limit?
Create Pay-As-You-Go cloud disks	Complete real-name registration before creating cloud disks in any Mainland China regions.	No
Quota of Pay-As-You-Go cloud disks in all regions for a single account	Number of Pay-As-You-Go instances in all regions under the user account \times 5	Yes
Quota of system disks for a single instance	1	No
Quota of data disks for a single instance	16 (including cloud disks and Shared Block Storage)	No
Quota of instances to which a single Shared Block Storage can be attached	8	No
Quota of Shared Block Storage in all regions for a single account	10	Yes
Capacity of a single Basic Cloud Disk	5 GiB-2,000 GiB	No
Capacity of a single SSD Cloud Disk	20 GiB-32,768 GiB	No
Capacity of a single Ultra Cloud Disk	20 GiB-32,768 GiB	No
Capacity of a single local SSD disk	5 GiB-800 GiB	No
Capacity of local SSD disks for a single instance	1,024 GiB	No
Capacity of a single local NVMe SSD disk	1,456 GiB	No
Capacity of local NVMe SSD disks for a single instance	2,912 GiB	No

Item	Limit	Can I open a ticket to change the limit?
Capacity of a single local SATA HDD disk	5,500 GiB	No
Capacity of local SATA HDD disks for a single instance	154,000 GiB	No
Capacity of a single SSD Shared Block Storage	32,768 GiB	No
Capacity of SSD Shared Block Storage for a single instance	128 TiB	No
Capacity of a single Ultra Shared Block Storage	32,768 GiB	No
Capacity of Ultra Shared Block Storage for a single instance	128 TiB	No
Capacity of a single ESSD disk	32,768 GiB	No
Capacity of a single system disk	<ul style="list-style-type: none"> Windows: 40 GiB-500 GiB Linux (excluding CoreOS) and FreeBSD: 20 GiB-500 GiB CoreOS: 30 GiB-500 GiB 	No
Capacity of a single data disk	<ul style="list-style-type: none"> Basic Cloud Disk: 5 GiB-2,000 GiB SSD Cloud Disk/Ultra Cloud Disk/SSD Shared Block Storage/Ultra Shared Block Storage: 20 GiB-32,768 GiB Local disk: dependent on specific disks 	No
Attach a new local disk to an instance with local disks	This feature is not supported.	No

Item	Limit	Can I open a ticket to change the limit?
Change configuration of an instance with local disks	Only bandwidth changes are allowed.	No
System disk mount points	/dev/xvda	No
Data disk mount points	/dev/xvd[b-z]	No

**Note:**

Block storage capacity is measured in binary units. 1 KiB is 1,024 bytes. 1 MiB is 1,024 KiB. 1 GiB is 1,024 MiB. 1 TiB is 1,024 GiB.

Snapshots

Item	Limit	Can I open a ticket to change the limit?
The number of snapshots that can be created for each cloud disk and Shared Block Storage	64	No

Images

Item	Limit	Can I open a ticket to change the limit?
Quota of custom images in a region for each account	100	Yes
Maximum number of users with whom a single image can be shared	50	Yes
Usage of images on instance types	32-bit images are not supported on an instance with 4 GiB or more RAM.	No

Key pairs

Item	Limit	Can I open a ticket to change the limit?
Quota of key pairs in a region for each account	500	No

Item	Limit	Can I open a ticket to change the limit?
Instance types supporting key pairs	All instance types except non-I/O optimized instance types in Generation I	No
Images supporting key pairs	Linux images only	No

Internet bandwidth

Item	Limit	Can I open a ticket to change the limit?
Maximum inbound Internet bandwidth	200 Mbit/s	No
Maximum outbound Internet bandwidth	<ul style="list-style-type: none"> Subscription instance: up to 200 Mbit/s Pay-As-You-Go instance : up to 100 Mbit/s 	No
Change the assigned Internet address for a single instance	The instance has existed for less than six hours. You can change the Internet address of an instance three times.	No

Security groups

Item	Limit on common security groups	Limit on enterprise-level security groups
Number of security groups in a region for each account	100 You can view the quota on the Overview page of the ECS console. If you require a higher quota, open a ticket.	Same as common security groups.

Item	Limit on common security groups	Limit on enterprise-level security groups
Number of classic network instances that can be added to a classic network security group	1000*	Classic network is not supported.
Number of VPC-connected instances that can be added to a VPC security group	Depends on the number of private IP addresses supported by a security group .	N/A
Number of security groups that an ECS instance can be added to	5 If you require a higher quota, open a ticket to increase the quota to 10 or 16.	Same as common security groups.
Number of security groups that each ENI of an ECS instance can be added to		
Maximum number of rules that can be added for a security group (including inbound and outbound rules)	100***	Same as common security groups.
Maximum number of rules for an ENI in all the security groups that the ENI is added to (including inbound and outbound rules)	500	Same as common security groups.
Number of private IP addresses supported by a VPC security group	2000**	N/A
Number of ENIs supported by a VPC security group	Depends on the number of security group rules .	50000

Item	Limit on common security groups	Limit on enterprise-level security groups
Internet access port	For the outbound Internet traffic, the default STMP port is 25, which is disabled by default and cannot be enabled by setting security group rules. For information about how to enable port 25, see .	Same as common security groups.

***If more than 1000 classic network instances require intranet access, you can assign them to multiple security groups and authorize mutual access between these security groups by setting security group rules.**

****If more than 2000 private IP addresses require intranet access, you can assign ECS instances corresponding to these private IP addresses to multiple security groups and authorize mutual access between these security groups by setting security group rules.**

*****Up to 500 security group rules can be added for an ECS instance. If you increase the quota of security groups that an ECS instance can be added to, the maximum number of rules for each security group will decrease. The number of security groups multiplied by the maximum number of inbound and outbound rules for each security group must be less than or equal to 500.**

Deployment sets

Item	Limit	Can I open a ticket to change the limit?
Quota of deployment sets in a region for each account	2	No
Number of instances that can be included in a deployment set	Up to seven instances are allowed in each zone. The number of instances allowed in a region equals the number of zones \times 7.	No

Item	Limit	Can I open a ticket to change the limit?
Instance types that can be created in a deployment set	c5, g5, hfc5, hfg5, r5, se1ne, sn1ne, and sn2ne	No

Cloud assistants

Item	Limit	Can I open a ticket to change the limit?
Quota of cloud assistant scripts in a region for each account	100	Yes
Quota to invoke cloud assistant scripts in a region for each account	5000	Yes

ENIs

Item	Limit	Can I open a ticket to change the limit?
Quota of ENIs in a region for each account	100 You can click Privileges & Quota Management on the Overview page of the ECS console to view the quota.	Yes

Tags

Item	Limit	Can I open a ticket to change the limit?
Quota of tags that can be bound to a single instance	20	No

API

Item	Limit	Can I open a ticket to change the limit?
Quota of CreateInstance calls	200 times per minute	Yes

**Note:**

For the limits of VPC products, see [Limits](#).

6 Usage notes

This topic describes the general limits for the usage of ECS instances and limits that are specific to the installed OS.

General limits

- After you create an ECS instance, only you have administrator permissions on this instance.
- You are not allowed to resell or sublicense any bandwidth assigned to your ECS instances without authorization. Failure to comply will result in a suspension or termination of your right to use any portions or all Alibaba Cloud services.
- Do not use your ECS instances for malicious, fraudulent, or unlawful activities, such as click farming or fraudulent transactions, on e-commerce websites such as Taobao. Failure to comply will result in suspension or termination of your account.
- Do not uninstall relevant hardware drivers.
- Do not modify the MAC address of ENIs unless necessary.
- Do not enable SELinux.
- An ECS instance with more than 4 GiB RAM must use a 64-bit operating system (OS). The 64-bit OSs currently supported are as follows:
 - Aliyun Linux 64-bit
 - CoreOS 64-bit
 - CentOS 64-bit
 - Debian 64-bit
 - FreeBSD 64-bit
 - OpenSUSE 64-bit
 - SUSE Linux 64-bit
 - Ubuntu 64-bit
 - Red Hat 64-bit
 - Windows 64-bit

For the latest available OSs, go to the instance purchase page in the ECS console.

- To guarantee service continuity, we recommend that you enable automatic startup upon instance boot for relevant software. If service applications are connected

to databases, we recommend that you enable automatic re-connection for these service applications.

- We recommend that you do not upgrade the kernel and the operating system. If you need to upgrade the kernel, see [How to avoid Linux instance startup failure after kernel upgrade](#).

Windows limits

- Do not stop the built-in AliyunService or shutdownmon.exe process. Otherwise, ECS instances may not be stopped or restarted properly.
- Do not modify the hostname of the domain controller.
- We recommend that you do not create custom images by using a virtual machine that acts as a domain controller.
- Do not rename, delete, or disable the administrator account.
- We recommend that you do not use the virtual memory if Basic Cloud Disks are used. For Ultra Cloud Disks, SSD Cloud Disks, or ESSD Cloud Disks, you can use the virtual memory as needed.
- Exercise caution when you use the administrator account to resize cloud disks, operate spanned volumes or the registry, and update the system. Failure to comply can result in data loss. For more information, see [Resize a cloud disk](#).
- Windows 32-bit OSs support up to four CPU cores.
- Ensure that a minimum of 2 GiB RAM is available when you build a website or deploy a Web environment on a Windows instance. Instances with 1-core CPU and 1 GiB RAM cannot start MySQL.
- For more information see [Select an image](#).

Linux limits

- Do not modify the content of the default `/etc/issue` file on Linux instances. Otherwise, if you create a custom image of such an instance and then use the image to create a new instance, the new instance cannot start properly because the operating system edition cannot be recognized.
- Do not modify permissions of the directories in the root partition, especially `/etc`, `/sbin`, `/bin`, `/boot`, `/dev`, `/usr`, and `/lib`. Improper modifications of permissions may cause system errors.
- Do not rename, delete, or disable the root account.
- Do not compile the Linux kernel or perform any other operations on it.

- We recommend that you do not use the swap partition if Basic Cloud Disks are used . However, for Ultra Cloud Disks, SSD Cloud Disks, or ESSD Cloud Disks, you can use the swap partition as needed.
- Exercise caution when you use the root account to run fio, mkfs, or fsck commands or resize cloud disks. Failure to comply can result in data loss.
- For more information see [Select an image](#).