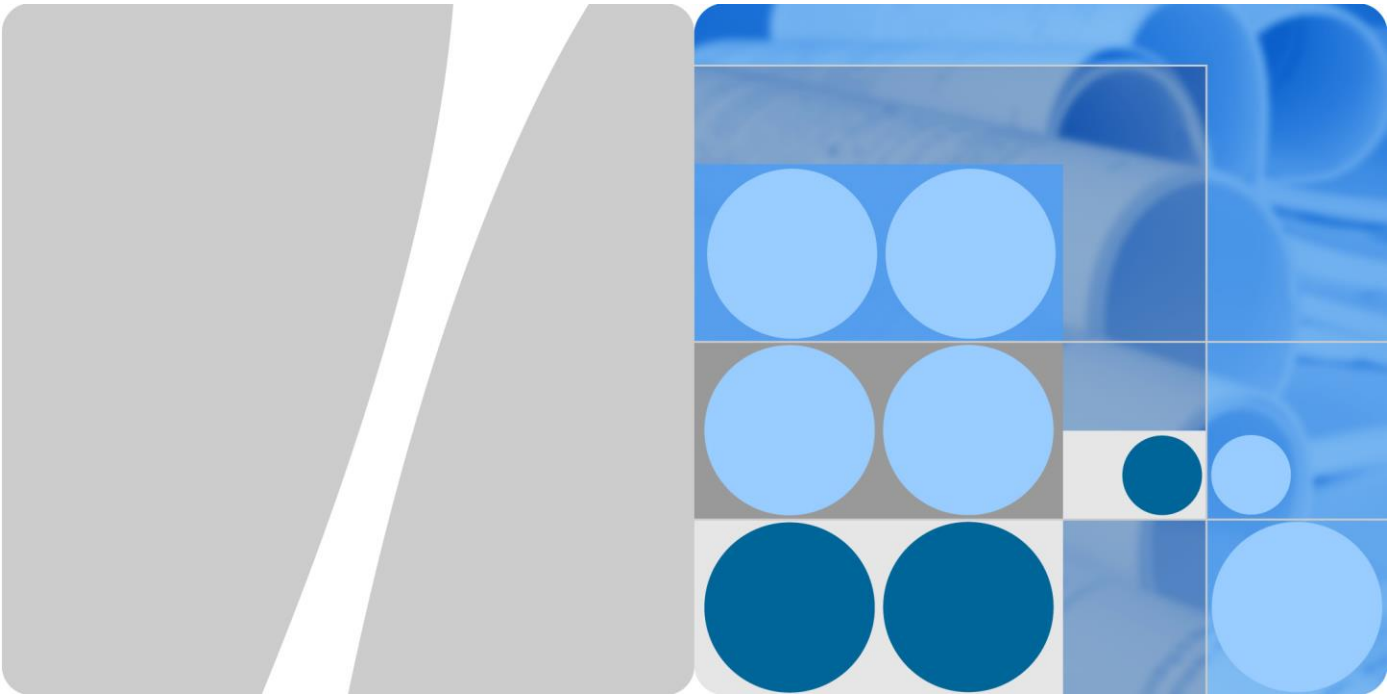

Elastic Cloud Server

Service Overview

Date **2019-10-28**

Contents

1 What Is ECS?

An Elastic Cloud Server (ECS) is a basic computing unit that consists of vCPUs, memory, OS, and Elastic Volume Service (EVS) disks. After creating an ECS, you can use it like using your local computer or physical server.

ECSs support self-service creation, modification, and operation. You can create an ECS by specifying its vCPUs, memory, OS, and login authentication. After the ECS is created, you can modify its specifications as required. This ensures an efficient, reliable, secure computing environment.

Why ECS?

- Rich specifications: A variety of ECS types are available for different scenario requirements. There are multiple customizable specifications for each type.
- Comprehensive images: Public, private, and shared images can be flexibly selected to request for ECSs.
- Differentiated EVS disks: Common I/O, high I/O, and ultra-high I/O EVS disks are available for all of your service requirements.
- Flexible billing modes: Yearly/Monthly and pay-per-use billing modes allow you to purchase and release resources at any time based on service fluctuation.
- Reliable data: Scalable, reliable high-throughput virtual block storage is based on distributed architecture.
- Security protection: The network is isolated and protected using security group rules from viruses and Trojan horses. Security services, such as Anti-DDoS, Web Application Firewall (WAF), and Vulnerability Scan Service (VSS) are included to further enhance ECS security.
- Flexible, easy-to-use: Elastic computing resources are automatically adjusted based on service requirements and policies to efficiently meet service requirements.
- Highly efficient O&M: Multi-choice management via the management console, remote access, and APIs with full rights.
- On-cloud monitoring: Cloud Eye samples monitored metrics in real time, correctly generates resource monitoring alarms, and sends notifications to related personnel immediately after the alarms are triggered.
- Load balancing: Elastic Load Balance (ELB) automatically distributes access traffic to multiple ECSs to balance their service load. It enables higher levels of fault tolerance in your applications and expands application service capabilities.

For more details, see 2 ECS Advantages and 3 ECS Application Scenarios.

System Architecture

ECS works with other products and services to provide computing, storage, network, and image installation functions.

- ECSs are deployed in multiple availability zones (AZs) connected with each other through an internal network. If an AZ becomes faulty, other AZs in the same region will not be affected.
- With the Virtual Private Cloud (VPC) service, you can build a dedicated network, set the subnet and security group, and allow the VPC to communicate with the external network through an EIP (bandwidth support required).
- With the Image Management Service (IMS), you can install images on ECSs, or create ECSs using private images and deploy services quickly.
- The Elastic Volume Service (EVS) provides storage and Volume Backup Service (VBS) provides data backup and recovery functions.
- Cloud Eye is a key measure to ensure ECS performance, reliability, and availability. Using Cloud Eye, you can determine ECS resource usage.
- Volume Backup Service (VBS) allows you to create data backups for EVS disks and use the backups to restore the EVS disks. This maximizes user data correctness and security.
- Cloud Server Backup Service (CSBS) backs up all EVS disks of an ECS, including the system disk and data disks, and uses the backup to restore the ECS.

Figure 1-1 System architecture

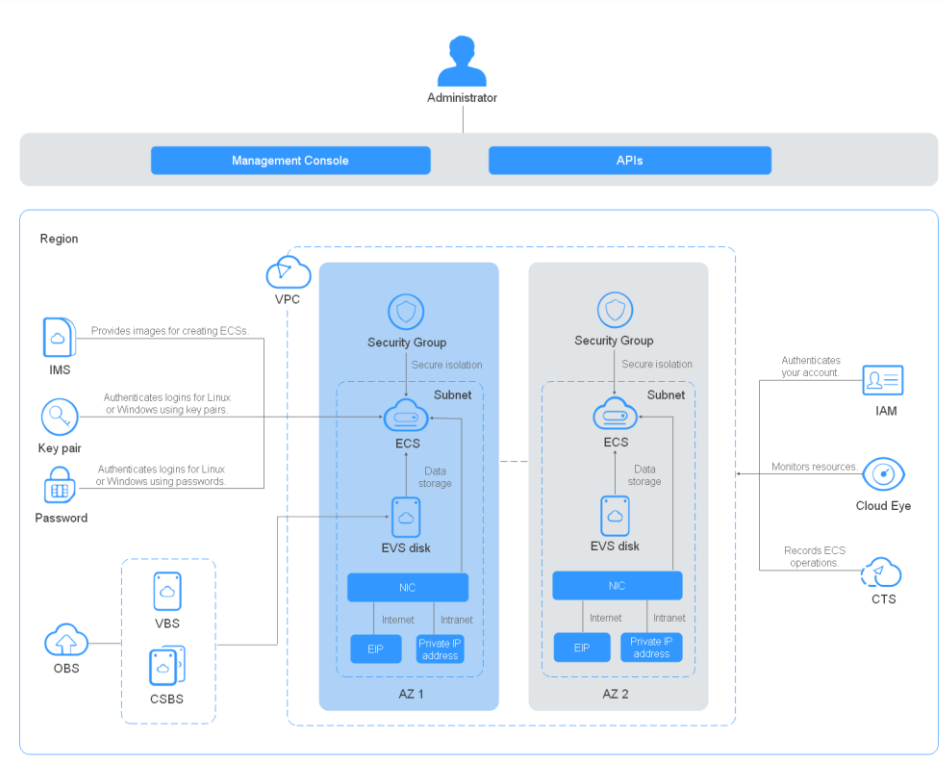

Access Methods

The public cloud provides a web-based service management platform. You can access ECSs through HTTPS-compliant application programming interfaces (APIs) or the management console. These two access modes differ as follows:

- Accessing ECSs through APIs

  Use this method if you are required to integrate the ECSs on the public cloud platform into a third-party system for secondary development. For detailed operations, see *Elastic Cloud Server API Reference*.

- Accessing ECSs through the management console

  Use this method if you are not required to integrate ECSs with a third-party system.

  After registering on the public cloud, log in to the management console and click **Elastic Cloud Server** under **Computing** on the homepage.

2 ECS Advantages

ECS supports automatic adjustment of computing resources based on service requirements and scaling policies. You can customize ECS configurations as needed, including vCPUs, memory, and bandwidth. Scalable, on-demand computing resources are provided by ECSs for secure, flexible, and efficient applications.

Stability and Reliability

- Differentiated EVS disks

  Common I/O, high I/O, and ultra-high I/O EVS disks are available for all of your service requirements.

  Common I/O EVS disks: feature secure, reliable, and scalable. They are ideal for applications requiring large capacity, normal read/write speed, and few transactions.

  High I/O EVS disks: feature high performance, scalability, and reliability. They are ideal for applications requiring high performance, high read/write speed, and real-time data storage.

  Ultra-high I/O EVS disks: feature low latency and high performance. They are ideal for intensive read/write applications requiring extremely high performance and read/write speed, and low latency.

- Reliable data

  Scalable, reliable high-throughput virtual block storage is based on distributed architecture. This ensures that data can be rapidly migrated and restored if any data replica is unavailable, preventing data loss caused by a single hardware fault.

- Backup and restoration of ECSs and EVS disks

  Automatic backup policies can be preset to back up in-service ECSs and EVS disks. Additionally, the data of ECSs and EVS disks at a specified time can be automatically backed up through the management console or API.

Security

- Various security services are provided for multi-dimensional protection.

  Security services, such as Anti-DDoS, WAF, and VSS are available.

- Security evaluation

  Cloud environment security evaluation helps you quickly identify security vulnerabilities and threats. Security configuration check and recommendations reduce or eliminate your loss from network viruses or attacks.

- Intelligent process management

Intelligent process management automatically prohibits the execution of unauthorized programs based on a customized whitelist, thereby ensuring ECS security.

- Vulnerability scan

  Various scanning services are provided, including general web vulnerability scanning, third-party application vulnerability scanning, port detection, and fingerprint identification.

Competitive Advantage

- Professional hardware devices

  ECSs are deployed on professional hardware devices that support in-depth virtualization optimization, relieving you of equipment-room concerns.

- Always available virtualization resources

  Scalable, dedicated resources can be obtained from the virtualized resource pool any time, ensuring reliable, secure, flexible, and efficient application environments. You can use your ECS like using your local computer.

Auto Scaling

- Automatic adjustment of computing resources

  Dynamic scaling: AS automatically increases or decreases the number of ECSs in an AS group based on monitored data.

  Periodic/Scheduled scaling: AS increases or decreases the number of ECSs in an AS group periodically or at a specified time based on service expectation and operation plan.

- Flexible adjustment of ECS configurations

  ECS specifications and bandwidth can be flexibly adjusted based on service requirements.

- Flexible billing modes

  Yearly/Monthly and pay-per-use billing modes allow you to purchase and release resources at any time based on service fluctuation.

3 ECS Application Scenarios

Internet

No special requirements on CPUs, memory, disk space, or bandwidth; strong security and reliability; application deployment based on one or only a few ECSs to minimize initial investment and maintenance costs, such as website R&D and testing, and small-scale databases

Use general-computing ECSs, which provide a balance of computing, memory, and network resources. This ECS type is appropriate for medium-load applications and meets the cloud service needs of both enterprises and individuals.

E-Commerce

Large amount of memory; capable of processing large volumes of data; fast network and rapid data processing such as precision marketing, E-Commerce, and mobile apps

Use memory-optimized ECSs, which have a large amount of memory and provide ultra-high I/O EVS disks and appropriate bandwidths.

Graphics Rendering

High-quality graphics and video; large amount of memory, capable of processing large volumes of data, and high I/O concurrency; fast network and rapid data processing; high GPU performance, such as graphics rendering and engineering drawing

Use GPU-accelerated ECSs, including G1 ECSs, which are based on NVIDIA Tesla M60 hardware virtualization and provide cost-effective graphics acceleration. These ECSs support DirectX and OpenGL, provide computing with up to 1 GB of GPU memory and 4,096 x 2,160 resolution.

Data Analysis

Capable of processing large volumes of data; high I/O performance and rapid data switching and processing, such as MapReduce and Hadoop

Use disk-intensive ECSs, which are designed for applications requiring sequential read/write on ultra-large datasets in local storage (such as distributed Hadoop computing) as well as large-scale parallel data processing and log processing. Disk-intensive ECSs are based on HDD and a default network bandwidth of 10GE, providing high PPS and low network latency. They also support up to 24 local disks, 48 vCPUs, and 384 GB of memory.

High-Performance Computing

High computing performance and throughput, such as scientific computing, genetic engineering, games and animation, biopharmaceuticals, and storage

Use high-performance computing ECSs to meet the computing, storage, and rendering needs of high-performance infrastructure services and applications that require a large number of parallel computing resources.

4 Notes on Using ECSs

Before using ECSs, read the following notes:

Notes on ECS Application Scenarios

- Do not use ECSs for any illegal or violation service, such as gambling, private service, or cross-border VPN.
- Do not use ECSs for fake transactions, such as click farming on e-commerce websites.
- Do not use ECSs to initiate network attacks, such as DDoS attacks, CC attacks, web attacks, brute force cracking, or spreading of viruses and Trojan horses.
- Do not use ECSs for traffic transit.
- Do not use ECSs to set up the crawler environment for data crawling.
- ECSs can be used for probe, such as scan or penetrate external systems only after being authorized by the external systems.
- Do not deploy any illegal websites or applications on ECSs.

Constraints

- Do not uninstall the driver on the ECS hardware.
- Do not change the MAC address of a NIC.
- ECSs do not support secondary virtualization.
- The authentication of certain software may bind a license to the physical server hardware. Once the ECS deployed on the physical server is migrated out, the bound license fails due to the change of the physical server.
- If an ECS migrates out of a faulty physical server, the ECS may be stopped or restarted. For service high availability, deploy services in cluster or active/standby mode, or configure automatic ECS startup upon a physical server failure or startup.
- Back up data for the ECSs where core services are deployed.
- Monitor application metrics on ECSs.
- Do not change the default DNS server. If you are required to configure public DNS, configure public and intranet DNS on ECSs.

Notes on Using Windows

- Do not stop system processes. Otherwise, blue screen of death (BSOD) may occur on the ECS, or the ECS may restart.

- Ensure that there is at least 2 GB of idle memory. Otherwise, BSOD, frame freezing, or service running failure may occur.
- Do not modify the registry. Otherwise, starting the system may fail. If the modification is mandatory, back up the registry before modifying it.
- Do not modify ECS clock settings. Otherwise, DHCP lease may fail, leading to the loss of IP addresses.
- Do not delete the CloudResetPwdAgent or CloudResetPwdUpdateAgent process. Otherwise, one-click password reset will not be available.
- Do not disable virtual memory. Otherwise, system performance may deteriorate, or system exceptions may occur.
- Do not delete the VMTool program. Otherwise, the ECS may fail to run properly.

Notes on Using Linux

- Do not modify the **/etc/issue** file. Otherwise, the system edition will not be identified.
- Do not delete system directories or files. Otherwise, the system may fail to start or run.
- Do not change the permissions or names of system directories. Otherwise, the system may fail to start or run.
- Upgrade Linux kernel only required. For details, see How Can I Upgrade the Kernel of a Linux ECS?
- Do not delete the CloudResetPwdAgent or CloudResetPwdUpdateAgent process. Otherwise, one-click password reset will not be available.
- Do not change the default DNS server **/etc/resolv.conf**. Otherwise, internal services, such as software sources and NTP may be unavailable.
- Do not modify default intranet configurations, such as IP addresses, subnet mask, and gateway address of an ECS. Otherwise, network exceptions may occur.

5 Spot Pricing

Concept

HUAWEI CLOUD sells idle computing resources at a discount. The price changes in real time depending on market demands. This is the spot price billing mode.

An ECS billed in spot price billing mode is a spot ECS.

In spot price billing mode, you can purchase and use ECSs at a discount price. The performance of spot ECSs is the same as that of the ECSs with the same specifications in other billing modes. However, when inventory resources are insufficient, or the market price increases and exceeds your expected price, the system will automatically release your ECS resources and reclaim the ECSs. Compared with pay-per-use and yearly/monthly ECSs, spot ECSs offer the same level of performance while at lower costs.

Working Rules

The market price for the ECSs of a certain flavor fluctuates due to supply-and-demand changes. You can purchase and use spot ECSs at a low market price to reduce computing costs.

When purchasing a spot ECS, you are required to set the highest price for a specified flavor you are willing to pay for. This process is also known as price quoting. A higher quoted price makes it more likely for you to purchase such an ECS.

- If the quotation is greater than or equal to the market price and the inventory resources are sufficient, the spot ECS is purchased. The ECS is billed at the market price.
- If the quotation is less than the market price, the spot ECS cannot be purchased.

After purchasing a spot ECS, you can use it like using the ECSs in other billing modes. However, the system will periodically compare your quotation with the market price and check the inventory resources.

- If the quotation is greater than or equal to the market price and the inventory resources are sufficient, you can continue using the ECS.
- If the quotation is less than the market price or the inventory resources are insufficient, the system notifies you of releasing the ECS resources (notification is enabled) and automatically deletes the ECS in about 5 minutes.

Figure 5-1 Lifecycle of a spot ECS

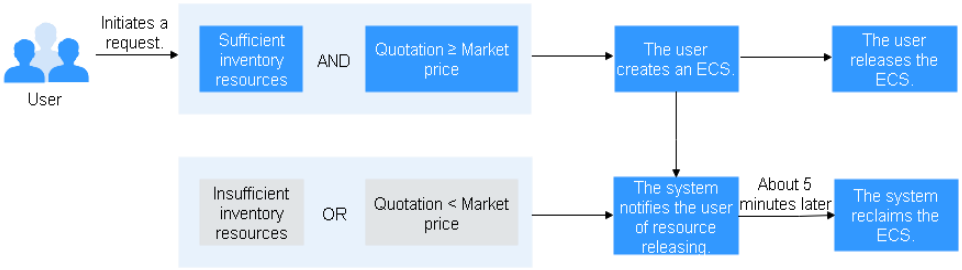

Application Scenarios

- What are supported?

  Spot ECSs are suitable for image rendering, stateless web service, gene sequencing, offline analysis, function calculation, batch calculation, sample analysis, CI/CD, and test.

  NOTE

  When the market price is higher than your quotation or the inventory resources are insufficient, the spot ECSs will be reclaimed. Therefore, back up data when using such ECSs.

- What are not supported?

  To prevent ECS reclamation from interrupting services, do not use spot ECSs to run the services requiring long-time operations or high stability.

Notes

- Only KVM ECSs support spot price payments. For details about the ECS flavors, see the information displayed on the management console.
- The market prices of the ECSs of the same flavor may vary depending on AZs.
- Spot ECSs do not support OS changing.
- Spot ECSs do not support automatic recovery.
- Spot ECSs do not support specifications modification.
- Spot ECSs cannot be created using a Marketplace image.
- Spot ECSs cannot be switched to yearly/monthly ECSs.
- When a spot ECS is being reclaimed,
  - It cannot be used to create system disk images and full-ECS images. However, data disks of the ECS can be used to create data disk images.
  - It cannot be deleted.

Purchasing a Spot ECS

You can purchase a spot ECS on the management console or by calling APIs.

- For instructions about how to purchase a spot ECS on the management console, see Purchasing a Spot ECS.
- For instructions about how to purchase a spot ECS by calling APIs, see Creating ECSs in *Elastic Cloud Server API Reference*.

Reclaiming an ECS

HUAWEI CLOUD may reclaim and terminate your spot ECS at any time. A spot ECS that is being reclaimed cannot be used to create images.

The reclamation may be due to:

- Higher market price than your quotation
- Insufficient inventory resources

NOTE

- If a spot ECS is reclaimed within the first hour after it is provisioned, the spot ECS is not billed.
- In the first settlement period (in hours) of a spot ECS, the spot ECS is billed, regardless of whether it runs.
- The time required for reclaiming a spot ECS is 5 minutes. During the reclaiming, if the sharp clock is exceeded, the spot ECS is billed at the market price for the time after the sharp clock.
- During the running of a spot ECS, its price is updated once an hour. After a spot ECS is restarted, or it is stopped and then started, it is billed at the market price when the ECS starts.

Back up data on spot ECSs. Before the system reclaims your spot ECS, it will notify you of resource releasing if notification is enabled. Use either of the following methods to enable notification:

- Method 1: Use the spot metadata. For details, see Obtaining ECS Metadata.
- Method 2: Use Cloud Trace Service (CTS) and Simple Message Notification (SMN) provided on the public cloud. For details, see Purchasing a Spot ECS.

6 Reserved Instance Overview

Concept

A reserved instance (RI) is not an actual instance, but a billing discount that can be applied to the use of pay-per-use ECSs in your account. When the attributes of your pay-per-use ECSs match those of an RI, the RI billing benefit automatically applies to your ECSs. The combination of RIs and pay-per-use payments fully utilizes the flexibility of pay-per-use resources at lower costs.

NOTE

RIs are in OBT phase.

Table 6-1 Comparison between RI, pay-per-use ECS, yearly/monthly ECS, and spot ECS

Item	What It is	How to Use
RI	A billing discount applied to pay-per-use ECSs.	When the attributes of your pay-per-use ECSs match those of an RI, the RI billing benefit automatically applies to your ECSs.
Pay-per-use ECS	ECS billed based on usage frequency and duration. Such an ECS can be created or deleted at any time.	A pay-per-use ECS is a basic computing unit that consists of vCPUs, memory, OS, and EVS disks. After purchasing such an ECS, you can use it on the cloud.
Yearly/Monthly ECS	ECS billed based on the service duration. This mode is ideal when the duration of ECS usage is predictable.	A yearly/monthly ECS is a basic computing unit that consists of vCPUs, memory, OS, and EVS disks. After purchasing such an ECS, you can use it on the cloud.
Spot ECS	ECS billed in spot price billing mode.	A spot ECS is a basic computing unit that consists of vCPUs, memory, OS, and EVS disks. After purchasing such an ECS, you can use it on the cloud.

- For instructions about how to purchase an RI, see Purchasing an RI.

- For instructions about how to modify an RI, see Modifying RI Attributes.

What Is Attribute Mapping Between an RI and a Pay-per-Use ECS?

A regional RI is purchased within a region and without an AZ specified. A zonal RI is purchased within an AZ.

- Attribute mapping of a regional RI: indicates whether the region, OS type, ECS series, and vCPU/memory ratio of a pay-per-use ECS are the same as those specified in a regional RI.
- Attribute mapping of a zonal RI: indicates whether the AZ, OS type, flavor of a pay-per-use ECS are the same as those specified in a zonal RI.

Application Scenarios

If your ECSs are to be used in a short term, you are advised to pay pay-per-use rates. If you plan to use ECSs for one or three years, you are advised to use RIs. RIs offer discounts for pay-per-use ECSs with matched attributes.

For example, after you purchase two s3.2xlarge Linux RIs with a one-year term in AZ 1, the billing benefit of the RIs is immediately applied to up to two pay-per-use s3.2xlarge Linux ECSs running in AZ 1.

Working Rules

For example, you have a running pay-per-use ECS in your account. After you purchase an RI that matches the attributes of this ECS, the billing benefit of the RI is automatically applied to your ECS when the RI takes effect. A purchased RI takes effect at the next hour.

Table 6-2 lists RI attributes. Based on these attributes, you can select your desired RIs.

Table 6-2 RI attributes

Parameter	Description
Region or AZ	• Regional RI: indicates an RI purchased in a region, without an AZ specified. Capacity reservations are not supported for regional RIs. • Zonal RI: indicates an RI purchased with an AZ specified. Capacity reservations are supported for zonal RIs.
Flavor	• When purchasing a regional RI, ensure that the ECS series and vCPU/memory ratio specified in the RI are the same as those of the target pay-per-use ECS. • When purchasing a zonal ECS, ensure that the flavor specified in the RI is the same as that of the target pay-per-use ECS.
OS	Specifies the OS of the ECS to be bought, which must match the OS specified in your RI. For example, if you want to use a Linux RI, select a Linux public or private image when purchasing an ECS.
Term	Specifies the service duration of an RI. A year is defined as 31,536,000 seconds (365 days).
Offering Class	Standard: Certain attributes, such as the ECS size can be modified

Parameter	Description
	during the term. However, the ECS type cannot be changed.
Payment Option	All upfront

Zonal RIs

A zonal RI, which is purchased within a specified AZ, offers a billing discount for the ECSs with the same OS and flavor as the RI in that AZ.

For example, after you purchase two c3.xlarge.2 Linux RIs with a one-year term in AZ 1, the billing benefit of the RIs is immediately applied to up to two pay-per-use c3.xlarge.2 Linux ECSs running in AZ 1.

Regional RIs

A regional RI, which is purchased within a specified region, has the following attributes:

- AZ flexibility: The RI discount applies to pay-per-use ECS usage in any AZ in a region.
- ECS size flexibility: The RI discount applies to pay-per-use ECS usage when the ECS series and vCPU/memory ratio of the target ECS are the same as those specified in the regional RI. ECS size flexibility is determined based on the normalization factor of the ECS size. ECS size flexibility does not apply to zonal RIs.

ECS size flexibility is applied from the smallest to the largest ECS size within the ECS series based on the normalization factor. Table 6-3 describes ECS size within an ECS type and corresponding normalization factor per hour.

NOTE

An ECS automatically benefits from the billing discount offered by a regional RI only when the ECS series and vCPU/memory ratio are the same as those specified in the RI.

For example, a regional c3.large.4 RI cannot be used on a c3.large.2 ECS because their vCPU/memory ratios are different.

Table 6-3 Normalization factors

ECS Size	Normalization Factor
small	1
medium	1
large	2
xlarge	4
2xlarge	8
4xlarge	16
6xlarge	24
7xlarge	28
8xlarge	32

ECS Size	Normalization Factor
9xlarge	36
12xlarge	48
14xlarge	56
15xlarge	60
16xlarge	64
26xlarge	104
52xlarge	208
nxlarge	n x 4

For example, an s3.large.2 ECS has a normalization factor of 2. You purchase an s3.large.2 Linux RI for the AP-Hong Kong region of HUAWEI CLOUD with a one-year term.

- If you have two running s3.medium.2 pay-per-use Linux ECSs in this region, the billing benefit is fully applied to both ECSs.

  Figure 6-1 Example RI 1

  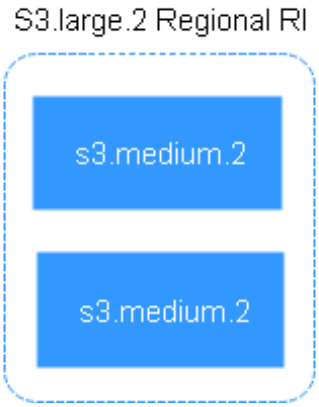

- If you have one running s3.xlarge.2 pay-per-use Linux ECS with a normalization factor of 4 in this region, the billing benefit is applied to 50% of the usage of the ECS.

Figure 6-2 Example RI 2

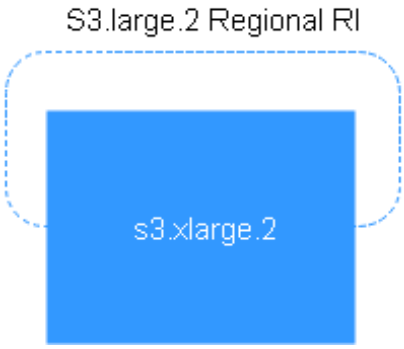

Table 6-4 Comparison between regional and zonal RIs

RI Type	AZ Flexibility	ECS Size Flexibility	Capacity Reservation
Regional RI	Supported A regional RI applies to any AZ in the region.	Supported A regional RI applies when the ECS series and vCPU/memory ratio of the target ECS are the same as those specified in the RI.	Not supported When available resources cannot meet service requirements of a pay-per-use ECS, creating the ECS failed.
Zonal RI	Not supported A zonal RI applies only in a specified AZ.	Not supported A zonal RI applies only when the flavor of the target ECS is the same as that specified in the RI.	Supported Desired resources can be reserved for creating a pay-per-use ECS.

Use Cases

You are running the following pay-per-use ECSs in account A:

- Five s3.large.2 Windows ECSs in AZ 1
- Three m3.xlarge.2 Windows ECSs in AZ 2
- One c3.xlarge.2 Windows ECS in AZ 3

You purchase the following RIs in region A:

- Five s3.large.2 Windows RIs with a one-year term in AZ 1
- Six m3.large.2 Windows RIs with a one-year term in region A
- One c3.large.2 Windows RI with a one-year term in region A

The RI benefits are applied as follows:

- The discount of the five s3.large.2 zonal RIs is used by the five s3.large.2 ECSs because the attributes (AZ, OS, and ECS type) between them match.
- The m3.large.2 regional RIs offer AZ flexibility and ECS size flexibility.

An m3.large.2 RI is equivalent to two normalization factors. The six m3.large.2 regional RIs are equal to 12 normalization factors (6 x 2). In account A, there are three running m3.xlarge.2 ECSs, which are equivalent to 12 normalization factors (3 x 4). In this case, the six m3.large.2 regional RIs are equivalent to three m3.xlarge.2 ECSs.

- The c3.large.2 regional RI offers AZ flexibility and ECS size flexibility and can be applied to c3.xlarge.2 ECSs.

  A c3.large.2 RI is equivalent to two normalization factors (1 x 2). A c3.xlarge.2 ECS requires an RI with four normalization factors (1 x 4). Therefore, the c3.large.2 RI billing discount applies to 50% of c3.xlarge.2 usage. The remaining c3.xlarge.2 usage is billed at the pay-per-use rate.

7 ECSs

7.1 Overview

An ECS is a basic computing unit that consists of vCPUs, memory, OS, and EVS disks.

After creating an ECS, you can use it like using your local computer or physical server, ensuring secure, reliable, and efficient computing. ECSs support self-service creation, modification, and operation. You can create an ECS by specifying its vCPUs, memory, OS, and login authentication. After the ECS is created, you can modify its specifications as required.

The cloud platform provides multiple ECS types for different computing and storage capabilities. One ECS type provides various flavors with different vCPU and memory configurations for you to select.

- For details about ECS types, see 7.3 ECS Types.
- For details about all ECS statuses in a lifecycle, see 7.2 ECS Lifecycle.

7.2 ECS Lifecycle

A lifecycle indicates the ECS statuses recorded from the time when the ECS is created through the time when the ECS is deleted or released.

Table 7-1 ECS statuses

Status	Status Attribute	Description
Creating	Intermediate	The ECS has been created but is not running.
Starting	Intermediate	The ECS is between the **Stopped** and **Running** states.
Running	Stable	The ECS is running properly. An ECS in this state can provide services.
Stopping	Intermediate	The ECS is between the **Running** and **Stopped** states.
Stopped	Stable	The ECS has been properly stopped.

Status	Status Attribute	Description
		An ECS in this state cannot provide services.
Restarting	Intermediate	The ECS is being restarted.
Resizing	Intermediate	The ECS has received a resizing request and has started to resize.
Verifying resizing	Intermediate	The ECS is verifying the modified configuration.
Deleting	Intermediate	The ECS is being deleted. If the ECS remains in this state for a long time, exceptions may have occurred. In such an event, contact the administrator.
Deleted	Intermediate	The ECS has been deleted. An ECS in this state cannot provide services and will be promptly cleared from the system.
Faulty	Stable	An exception has occurred on the ECS. An ECS in this state cannot provide services. Contact the administrator.
Reinstalling OS	Intermediate	The ECS has received a request to reinstall the OS and has begun the reinstallation.
Reinstalling OS failed	Stable	The ECS received a request to reinstall the OS, but due to exceptions, the reinstallation failed. An ECS in this state cannot provide services. Contact the administrator.
Changing OS	Intermediate	The ECS received a request to change the OS and has begun implementing the changes.
OS Change failed	Stable	The ECS has received a request to change the OS, but due to exceptions, the changes failed to be implemented. An ECS in this state cannot provide services. Contact the administrator.
Forcibly restarting	Intermediate	The ECS is being forcibly restarted.
Rolling back resizing	Intermediate	The ECS is rolling back resizing.
Frozen	Stable	The ECS has been stopped by the administrator because the order has expired or is overdue. An ECS in this state cannot provide services. The system retains it for a period of time. If it is not renewed after the time expires, the system will automatically delete the ECS.

7.3 ECS Types

The public cloud provides the following ECS types for different application scenarios:

- General computing
- General computing-plus
- Memory-optimized
- Large-memory
- Disk-intensive
- Ultra-high I/O
- High-performance computing

ECS Flavor Naming Rules

ECS flavors are named using the format "AB.C.D".

The format is defined as follows:

- **A** specifies the ECS type. For example, **s** indicates a general computing ECS, **c** a computing ECS, and **m** a memory-optimized ECS.
- **B** specifies the type ID. For example, the **1** in **s1** indicates a general computing first-generation ECS, and the **2** in **s2** indicates a general computing second-generation ECS.
- **C** specifies the flavor size, such as medium, large, or xlarge.
- **D** specifies the ratio of vCPUs to memory expressed in a digit. For example, value **4** indicates that the ratio of vCPUs to memory is 4.

Network Bandwidth

The intranet bandwidth and PPS of an ECS are determined based on ECS specifications.

- Assured intranet bandwidth: indicates the assured ECS bandwidth.
- Maximum intranet bandwidth: indicates the maximum ECS bandwidth.
- Maximum intranet PPS: indicates the maximum ECS capabilities in transmitting and receiving packets.

7.4 General Computing ECSs

Overview

General computing ECSs provide a baseline level of CPU performance with the ability to burst above the baseline and a balance of computing, memory, and network resources. These ECSs are suitable for many applications, such as web servers, enterprise R&D, and small-scale databases.

Specifications

Table 7-2 S2 ECS specifications

Flavor	vCPUs	Memory (GB)	Maximum/Assured Bandwidth (Gbit/s)	Maximum PPS (10,000)	NIC Multi-Queue	Virtualization Type
s2.small.1	1	1	0.5/0.1	5	1	KVM
s2.medium.2	1	2	0.5/0.1	5	1	KVM
s2.large.2	2	4	0.8/0.2	10	1	KVM
s2.xlarge.2	4	8	1.5/0.4	15	1	KVM
s2.2xlarge.2	8	16	3/0.8	20	2	KVM
s2.4xlarge.2	16	32	4/1.5	30	4	KVM
s2.medium.4	1	4	0.5/0.1	5	1	KVM
s2.large.4	2	8	0.8/0.2	10	1	KVM
s2.xlarge.4	4	16	1.5/0.4	15	1	KVM
s2.2xlarge.4	8	32	3/0.8	20	2	KVM
s2.4xlarge.4	16	64	4/1.5	30	4	KVM

Table 7-3 S3 ECS specifications

Flavor	vCPUs	Memory (GB)	Maximum/Assured Bandwidth (Gbit/s)	Maximum PPS (10,000)	NIC Multi-Queue	Virtualization Type
s3.small.1	1	1	0.5/0.1	5	1	KVM
s3.medium.2	1	2	0.5/0.1	5	1	KVM
s3.large.2	2	4	0.8/0.2	10	1	KVM
s3.xlarge.2	4	8	1.5/0.4	15	1	KVM

Flavor	vCPUs	Memory (GB)	Maximum/Assured Bandwidth (Gbit/s)	Maximum PPS (10,000)	NIC Multi-Queue	Virtualization Type
s3.2xlarge.2	8	16	3/0.8	20	2	KVM
s3.4xlarge.2	16	32	4/1.5	30	4	KVM
s3.medium.4	1	4	0.5/0.1	5	1	KVM
s3.large.4	2	8	0.8/0.2	10	1	KVM
s3.xlarge.4	4	16	1.5/0.4	15	1	KVM
s3.2xlarge.4	8	32	3/0.8	20	2	KVM
s3.4xlarge.4	16	64	4/1.5	30	4	KVM

Scenarios

- Applications

  General computing ECSs are suitable for applications that have no special requirements on CPU performance, memory, disk capacity, or bandwidth, but have high requirements on security and reliability. They feature low initial investment and maintenance costs.

- Application scenarios

  Enterprise website deployment, enterprise office environment setup, enterprise R&D and testing activities, Web servers, R&D and testing environments for developers, and small-scale databases

7.5 General Computing-plus ECSs

Overview

Compared with general computing ECSs, the general computing-plus ECSs provide the combinations of vCPUs and memory with larger specifications, offering more options for you to select. In addition, the ECSs use latest-generation network acceleration engines and DPDK rapid packet processing mechanism to provide higher network performance, meeting requirements in different scenarios.

- C3 ECSs are newly released. They use latest-generation Intel Xeon Skylake CPUs and feature high and stable computing performance. Working in high-performance networks, the C3 ECSs provide higher performance and stability, meeting enterprise-class application requirements.

Specifications

Table 7-4 C3 ECS specifications

Flavor	vCPUs	Memory (GB)	Maximum/Assured Bandwidth (Gbit/s)	Maximum PPS (10,000)	NIC Multi-Queue	Virtualization Type
c3.large.2	2	4	1.5/0.6	30	2	KVM
c3.xlarge.2	4	8	3/1	50	2	KVM
c3.2xlarge.2	8	16	5/2	90	4	KVM
c3.3xlarge.2	12	24	7/3	110	4	KVM
c3.4xlarge.2	16	32	10/4	130	4	KVM
c3.6xlarge.2	24	48	12/6	200	8	KVM
c3.8xlarge.2	32	64	15/8	260	8	KVM
c3.15xlarge.2	60	128	17/16	500	16	KVM
c3.large.4	2	8	1.5/0.6	30	2	KVM
c3.xlarge.4	4	16	3/1	50	2	KVM
c3.2xlarge.4	8	32	5/2	90	4	KVM
c3.3xlarge.4	12	48	7/3	110	4	KVM
c3.4xlarge.4	16	64	10/4	130	4	KVM
c3.6xlarge.4	24	96	12/6	200	8	KVM
c3.8xlarge.4	32	128	15/8	260	8	KVM
c3.15xlarge.4	60	256	17/16	500	16	KVM

Scenarios

- C3 ECSs

  Small- and medium-scale databases, cache servers, and search clusters with high requirements on stability; enterprise-class applications of diverse types and in various scales

7.6 General Computing-Basic ECSs

Overview

General computing-basic ECSs provide a balance of computing, memory, and network resources. Each ECS provides a baseline CPU performance, accumulating CPU credits when workloads operate below the baseline threshold. They are suitable for applications requiring burstable performance while keeping costs low.

For details about CPU usage calculations, see 12 CPU Credits.

Specifications

Table 7-5 T6 ECS specifications

Flavor	vCPUs	Memory (GB)	Initial Credits	Maximum Credits	CPU Credits/Hour	CPU Baseline (%)	Average CPU Baseline (%)	Maximum NICs	Virtualization Type
t6.small.1	1	1	30	144	6	10	10	1	KVM
t6.large.1	2	2	60	576	24	40	20	1	KVM
t6.xlarge.1	4	4	120	1152	48	80	20	2	KVM
t6.2xlarge.1	8	8	120	1728	72	120	15	2	KVM
t6.4xlarge.1	16	16	160	3456	144	240	15	2	KVM
t6.medium.2	1	2	30	144	6	10	10	1	KVM
t6.large.2	2	4	60	576	24	40	20	1	KVM
t6.xlarge.2	4	8	120	1152	48	80	20	2	KVM
t6.2xlarge.2	8	16	120	1728	72	120	15	2	KVM
t6.4xlarge.2	16	32	160	3456	144	240	15	2	KVM
t6.large.4	2	8	60	576	24	40	20	1	KVM
t6.xlarge.4	4	16	120	1152	48	80	20	2	KVM

Flavor	vCPUs	Memory (GB)	Initial Credits	Maximum Credits	CPU Credits/Hour	CPU Baseline (%)	Average CPU Baseline (%)	Maximum NICs	Virtualization Type
t6.2xlarge.4	8	32	120	1728	72	120	15	2	KVM

7.7 Memory-optimized ECSs

Overview

Memory-optimized ECSs have a large memory size and provide high memory performance. They are designed for memory-intensive applications that process a large amount of data, such as precision advertising, e-commerce big data analysis, and IoV big data analysis.

- M3 ECSs are developed based on the KVM virtualization platform and designed for processing large-scale data sets in the memory. They use latest-generation Intel Xeon Skylake CPUs, network acceleration engines, and DPDK rapid packet processing mechanism to provide higher network performance, providing a maximum memory size of 512 GB based on DDR4 for high-memory computing applications.
- M2 ECSs use Intel Xeon E5-2690 v4 CPUs and are designed for memory-optimized applications.

Specifications

Table 7-6 M3 ECS specifications

Flavor	vCPUs	Memory (GB)	Maximum/Assured Bandwidth (Gbit/s)	Maximum PPS (10,000)	NIC Multi-Queue	Virtualization Type
m3.large.8	2	16	1.5/0.6	30	2	KVM
m3.xlarge.8	4	32	3/1.1	50	2	KVM
m3.2xlarge.8	8	64	5/2	90	4	KVM
m3.3xlarge.8	12	96	8/3.5	110	4	KVM
m3.4xlarge.8	16	128	10/4.5	130	4	KVM
m3.6xlarge.	24	192	12/6.5	200	8	KVM

Flavor	vCPUs	Memory (GB)	Maximum/Assured Bandwidth (Gbit/s)	Maximum PPS (10,000)	NIC Multi-Queue	Virtualization Type
8						
m3.8xlarge.8	32	256	15/9	260	8	KVM
m3.15xlarge.8	60	512	17/17	500	16	KVM

Table 7-7 Memory-optimized ECS specifications

ECS Type	vCPUs	Memory (GB)	Flavor	Virtualization Type
Memory-optimized	2	16	m2.large.8	KVM
	4	32	m2.xlarge.8	KVM
	8	64	m2.2xlarge.8	KVM
	16	128	m2.4xlarge.8	KVM

Scenarios

- Applications

  Memory-optimized ECSs are suitable for applications that process large volumes of data and require a large amount of memory, rapid data switching and processing, and low-latency storage resources.

- Application scenarios

  Big data analysis, precision advertising, e-commerce big data analysis, IoV big data analysis, relational databases, NoSQL databases, and memory data analysis

Notes on Using M2 ECSs

- M2 ECSs support all rollout OSs.
- M2 ECSs do not have InfiniBand or SSD NICs configured.
- M2 ECSs support modifying specifications if the source and target ECSs are of the same type.
- To improve network performance, you can set the NIC MTU to **8888**.

Notes on Using M3 ECSs

- M3 ECSs support all rollout OSs.
- M3 ECSs do not have InfiniBand or SSD NICs configured.

- M3 ECSs support modifying specifications if the source and target ECSs are of the same type.

7.8 Large-Memory ECSs

Overview

Large-memory ECSs provide an even larger amount of memory than memory-optimized ECSs. They are used for applications that require a large amount of memory, rapid data switching, low latency, and process large volumes of data. Large-memory ECSs provide large memory and high computing, storage, and network performance.

- Applications

  Large-memory ECSs are suitable for applications that require a large amount of memory, rapid data switching, and low latency, and process large volumes of data.
- Application scenarios

  E3 ECSs: OLAP and OLTP applications with hyper-threading enabled

Specifications

Table 7-8 E3 ECS specifications

Flavor	vCPUs	Memory (GB)	Virtualization Type
e3.7xlarge.12	28	348	KVM

Notes

- Large-memory ECSs do not support NIC hot swapping.
- E3 ECSs support the following OS that has been verified:

  SUSE Enterprise Linux Server 12 SP2 64bit
- E1 and E2 ECSs can use the following types of EVS disks as system disk and data disk:
  - High I/O (performance-optimized I)
  - Ultra-high I/O (latency-optimized)
- The primary and extension NICs of a large-memory ECS have specified application scenarios. For details, see Table 7-9.

Table 7-9 Application scenarios of the NICs of a large-memory ECS

NIC Type	Application Scenario	Remarks
Primary NIC	Applies to vertical layer 3 communication.	N/A
Extension NIC	Applies to horizontal layer 2 communication.	To improve network performance, you can set the MTU of the extension NIC to **8888**.

7.9 Disk-intensive ECSs

Overview

D3 ECSs use latest-generation Intel Xeon Skylake CPUs to offer robust, stable computing performance. Equipped with Huawei 25GE high-speed intelligent NICs and local SAS disks, D3 ECSs ultra-high network bandwidth, PPS, and local storage with excellent price/performance ratio.

D2 ECSs are developed based on KVM virtualization. They use local storage and provide high storage performance and intranet bandwidth for distributed Hadoop computing, large data warehouse, distributed file system, data processing, and log processing.

D1 ECSs are developed based on the Xen virtualization platform, use local storage, and provide high network performance. They are designed for applications requiring sequential read/write on ultra-large datasets in local storage (such as distributed Hadoop computing) as well as large-scale parallel data processing and log processing.

Specifications

Table 7-10 D2 ECS specifications

Flavor	vCPUs	Memory (GB)	Maximum/Assured Bandwidth (Gbit/s)	Maximum PPS (10,000)	NIC Multi-Queue	Local Disks (GB)	Virtualization Type
d2.xlarge.8	4	32	3/1	15	2	2 x 1800	KVM
d2.2xlarge.8	8	64	5/2	30	2	4 x 1800	KVM
d2.4xlarge.8	16	128	8/4	40	4	8 x 1800	KVM
d2.6xlarge.8	24	192	10/6	50	6	12 x 1800	KVM
d2.8xlarge.8	32	256	13/8	60	8	16 x 1800	KVM
d2.12xlarge.8	48	384	13/13	90	8	24 x 1800	KVM

D3 ECS Application Scenarios

- Applications

D3 ECSs are suitable for applications that require large volumes of data to process, high I/O performance, and rapid data switching and processing.

- Application scenarios

  MPP data warehouse, distributed MapReduce and Hadoop computing, distributed file systems, network file systems, and log/data processing

D1 and D2 ECS Application Scenarios

- Applications

  D1 and D2 ECS ECSs are suitable for applications that require large volumes of data to process, high I/O performance, and rapid data switching and processing.

- Application scenarios

  MapReduce, Hadoop, and data-intensive computing, big data computing, network file systems, and data processing

Features of D2 ECSs

- D2 ECSs use local disks to provide high sequential read/write performance and low latency, improving file read/write performance.
- D2 ECSs provide powerful and stable computing capabilities, ensuring efficient data processing.
- D2 ECSs with a CPU/memory ratio of 1:8 process large volumes of data.
- D2 ECSs provide high intranet performance, including high intranet bandwidth and packet per second (pps), meeting requirements for data exchange between ECSs during peak hours.

Table 7-11 Specifications of a single SAS HDD disk attached to a D2 ECS

Metric	Performance
Disk capacity	1800 GB
Maximum throughput	230 MB/s
Access latency	Millisecond-level

Notes on Using D2 ECSs

- D2 ECSs support the following OSs:
  - CentOS 6.7/6.8/7.2/7.3/7.4 64bit
  - SUSE Enterprise Linux Server 11 SP3/SP4 64bit
  - SUSE Enterprise Linux Server 12 SP1/SP2 64bit
  - Red Hat Enterprise Linux 6.8/7.3 64bit
  - Windows Server 2008 R2 Enterprise 64bit
  - Windows Server 2012 R2 Standard 64bit
  - Windows Server 2016 Standard 64bit
  - Debian 8.7/9/9.0.0 64bit
  - EulerOS 2.2 64bit

  - Fedora 25/26 64bit
  - OpenSUSE 42.2/42.3 64bit
- When the physical host where a D2 ECS is deployed becomes faulty, the ECS cannot be migrated.
- To improve network performance, you can set the NIC MTU to **8888**.
- D2 ECSs do not support modifying specifications.
- D2 ECSs do not support local disk snapshot or backup.
- D2 ECSs do not support OS reinstallation or change.
- D2 ECSs do not support automatic recovery.
- D2 ECSs can use both local disks and EVS disks to store data. In addition, they can have EVS disks attached to provide a larger storage size. Use restrictions on the two types of storage media are as follows:
  - Only an EVS disk, not a local disk, can be used as the system disk of a D2 ECS.
  - Both an EVS disk and a local disk can be used as the data disk of a D2 ECS.
  - A D2 ECS can be attached with a maximum of 60 disks (including VBD, SCSI, and local disks). Among the 60 disks, the maximum number of SCSI disks is 30, and the maximum number of VBD disks is 24 (including the system disk). For details about restrictions, see Can Multiple Disks Be Attached to an ECS?
  - You are advised to use World Wide Names (WWNs), but not drive letters, in applications to perform operations on local disks to prevent drive letter drift (low probability) on Linux. Take local disk attachment as an example:

    If the local disk WWN is wwn-0x50014ee2b14249f6, run the **mount /dev/disk/by-id/wwn-0x50014ee2b14249f6** command.

NOTE

How can I view the local disk WWN?

1. Log in to the ECS.
2. Run the following command to view the WWN:

   ll /dev/disk/by-id

- The basic resources, including vCPUs, memory, and image of a stopped D2 ECS are still billed. To stop billing such an ECS, delete it.
- The local disk data of a D2 ECS may be lost due to some reasons, such as host machine breakdown or local disk damage. If your application does not use the data reliability architecture, you are strongly advised to use EVS disks to build your ECS.
- When a D2 ECS is deleted, its local disk data is automatically deleted. Back up the data before deleting such an ECS. Deleting local disk data is time-consuming. Therefore, a D2 ECS requires a longer period of time than other ECSs for releasing resources.
- Do not store service data for a long time in local disks. Instead, use EVS disks to store the data. In addition, back up data in a timely manner and use a high availability architecture.
- You are not allowed to buy additional local disks. The quantity and capacity of your local disks are determined according to the specifications of your ECS. For D2 ECSs, if additional local disks are required, buy them when creating the ECSs.

7.10 Ultra-high I/O ECSs

Overview

Ultra-high I/O ECSs use high-performance local NVMe SSDs to provide high storage IOPS and low read/write latency. The ratio of vCPU to memory is 1:8. You can create such ECSs with high-performance NVMe SSDs attached on the management console.

Ultra-high I/O ECSs can be used for high-performance relational databases, NoSQL databases (such as Cassandra and MongoDB), and ElasticSearch.

Specifications

Table 7-12 I3 ECS specifications

Flavor	vCPUs	Memory (GB)	Maximum/Assured Bandwidth (Gbit/s)	Maximum PPS (10,000)	NIC Multi-Queue	Local Disks	Maximum NICs	Virtualization Type
i3.2xlarge.8	8	64	8/3.5	100	4	1 x 1600 GB NVMe	4	KVM
i3.4xlarge.8	16	128	15/7	160	4	2 x 1600 GB NVMe	8	KVM
i3.8xlarge.8	32	256	20/14	280	8	4 x 1600 GB NVMe	8	KVM
i3.12xlarge.8	48	384	25/20	420	8	6 x 1600 GB NVMe	8	KVM
i3.15xlarge.8	60	512	25/25	500	16	7 x 1600 GB NVMe	8	KVM

Features

Table 7-13 lists the IOPS performance of I3 ECSs.

Table 7-13 IOPS performance

Flavor	Maximum IOPS for Random 4 KB Read
i3.2xlarge.8	750,000
i3.4xlarge.8	1,500,000
i3.8xlarge.8	3,000,000
i3.12xlarge.8	4,500,000
i3.15xlarge.8	5,250,000

Table 7-14 Specifications of a single NVMe disk attached to an I3 ECS

Metric	Performance
Disk capacity	1.6 TB
IOPS for random 4 KB read	750,000
IOPS for random 4 KB write	200,000
Read throughput	2.9 GB/s
Write throughput	1.9 GB/s
Access latency	Microsecond-level

Notes

- I3 ECSs do not support specifications modification.
- After an I3 ECS is deleted, the data on the local NVMe SSD is automatically deleted. Back up the data before deleting such an ECS. Deleting local disk data is time-consuming. Therefore, an I3 ECS requires a longer period of time than other ECSs for releasing resources.
- The data reliability of local disks depends on the reliability of physical servers and hard disks, which are SPOF-prone. Therefore, you are advised to perform data redundancy at the application layer to ensure data availability. Use EVS disks to store long-term service data.
- The device name of a local disk attached to an I3 ECS is **/dev/nvme0n1** or **/dev/nvme0n2**.
- I3 ECSs support the following OSs:
  - EulerOS 2.2
  - CentOS 7.2
  - CentOS 7.3

  - Ubuntu Server 16.04
  - SUSE Linux Enterprise Server 12 SP2
  - Fedora 25 64bit
  - OpenSUSE 42.2 64bit
- An I3 ECS can be attached with a maximum of 60 disks (including VBD, SCSI, and local disks). Among the 60 disks, the maximum number of SCSI disks is 30, and the maximum number of VBD disks is 24 (including the system disk).
- The basic resources, including vCPUs, memory, and image of an I3 ECS are still billed. To stop billing such an ECS, delete it.

7.11 High-Performance Computing ECSs

Overview

H3 ECSs use high-performance Intel Xeon SkyLake CPUs. Each vCPU corresponds to the hyper-thread of an Intel Xeon processor core, providing stable computing capabilities. H3 ECSs are suitable for high-performance computing services. In addition, the ECSs use latest-generation network acceleration engines and DPDK rapid packet processing mechanism to provide high network performance.

The processor and memory ratio of an HC2 ECS is 1:2 or 1:4. Each vCPU corresponds to the hyperthreading of an Intel Xeon processor core. HC2 ECSs can be used for high-performance computing services. They provide a large number of parallel computing resources and high-performance infrastructure services to meet the requirements of high-performance computing and massive storage and ensure rendering efficiency. HC2 ECSs are frequently used in the following scenarios:

- Computing and storage systems for genetic engineering, games, animations, biopharmaceuticals, and scientific computing
- Public rendering platforms for renderfarms and animation and film bases; other rendering platforms for movies and videos
- High-performance frontend clusters, web servers, high-performance science and engineering applications, advertisements, video coding, and distributed analysis

Specifications

Table 7-15 H3 ECS specifications

Flavor	vCPUs	Memory (GB)	Maximum/Assured Bandwidth (Gbit/s)	Maximum PPS (10,000)	NIC Multi-Queue	Virtualization Type
h3.large.2	2	4	2/1	30	2	KVM
h3.xlarge.2	4	8	4/2	60	2	KVM
h3.2xlarge.2	8	16	6/3.5	120	4	KVM

Flavor	vCPUs	Memory (GB)	Maximum/Assured Bandwidth (Gbit/s)	Maximum PPS (10,000)	NIC Multi-Queue	Virtualization Type
h3.3xlarge.2	12	24	6/5.5	160	4	KVM
h3.4xlarge.2	16	32	12/7.5	200	8	KVM
h3.large.4	2	8	2/1	30	2	KVM
h3.xlarge.4	4	16	4/2	60	2	KVM
h3.2xlarge.4	8	32	6/3.5	120	4	KVM
h3.3xlarge.4	12	48	6/5.5	160	4	KVM
h3.4xlarge.4	16	64	12/7.5	200	8	KVM

Table 7-16 HC2 ECS specifications

Flavor	vCPUs	Memory (GB)	Maximum/Assured Bandwidth (Gbit/s)	Maximum PPS (10,000)	NIC Multi-Queue	Virtualization Type
hc2.large.2	2	4	1.5/0.5	10	1	KVM
hc2.xlarge.2	4	8	3/1	15	1	KVM
hc2.2xlarge.2	8	16	5/2	30	2	KVM
hc2.4xlarge.2	16	32	8/4	40	4	KVM
hc2.large.4	2	8	1.5/0.5	10	1	KVM
hc2.xlarge.4	4	16	3/1	15	1	KVM
hc2.2xlarge.4	8	32	5/2	30	2	KVM
hc2.4xlarge.4	16	64	8/4	40	4	KVM

Scenarios

- Computing and storage systems for genetic engineering, games, animations, and biopharmaceuticals
- Public rendering platforms for renderfarms and animation and film bases; other rendering platforms for movies and videos
- High-performance frontend clusters, web servers, high-performance science and engineering applications, advertisements, video coding, and distributed analysis
- Batch-processed workload, high-performance computing (HPC), and SAP applications
- Computing-intensive services, such as large-scale multiplayer online (MMO) gaming

7.12 GPU-accelerated ECSs

GPU-accelerated ECSs provide outstanding floating-point computing capabilities. They are suitable for scenarios that require real-time, highly concurrent massive computing.

GPU-accelerated ECS Types

- G series
- P series
  - Computing-accelerated P2v
  - Inference-accelerated PI1

Table 7-17 GPU-accelerated ECSs

Classification	ECS Type	GPU	Application Scenario
Computing-accelerated	P2v	NVIDIA V100 NVLink (GPU passthrough)	Machine learning, deep learning, inference training, scientific computing, seismic analysis, computing finance, rendering, multimedia encoding and decoding
Inference-accelerated	PI1	NVIDIA P4 (GPU passthrough)	

Images Supported by GPU-accelerated ECSs

Table 7-18 Image list

Classification	ECS Type	Supported Image
Computing-accelerated	P2v	Windows Server 2016 Standard 64bit Windows Server 2012 R2 Standard 64bit Ubuntu 16.04 64bit CentOS 7.4 64bit EulerOS 2.2 64bit
Inference-accelerated	PI1	CentOS 7.3 64bit

Classification	ECS Type	Supported Image
		Ubuntu 16.04 64bit Ubuntu 14.04 64bit

Computing-accelerated P2v

Overview

P2v ECSs use NVIDIA Tesla V100 GPUs (16 GB GPU memory) and provide flexibility, high-performance computing, and cost-effectiveness. These ECSs use GPU NVLink for direct communication between GPUs, improving data transmission efficiency. P2v ECSs provide outstanding general computing capabilities and have strengths in AI-based deep learning, scientific computing, Computational Fluid Dynamics (CFD), computing finance, seismic analysis, molecular modeling, and genomics.

Specifications

Table 7-19 P2v ECS specifications

Flavor	vCPUs	Memory (GB)	Maximum/Assured Bandwidth (Gbit/s)	Maximum PPS (10,000)	NIC Multi-Queue	GPUs	GPU Connection	GPU Memory (GB)	Virtualization Type
p2v.2xlarge.8	8	64	10/4	50	4	1 x V100	N/A	1 x 16 GB	KVM
p2v.4xlarge.8	16	128	15/8	100	8	2 x V100	NVLink	2 x 16 GB	KVM
p2v.8xlarge.8	32	256	25/15	200	16	4 x V100	NVLink	4 x 16 GB	KVM
p2v.16xlarge.8	64	512	30/30	400	32	8 x V100	NVLink	8 x 16 GB	KVM

P2v ECS Features

- Up to eight NVIDIA Tesla V100 GPUs on an ECS
- NVIDIA CUDA parallel computing and common deep learning frameworks, such as TensorFlow, Caffe, PyTorch, and MXNet
- 15.7 TFLOPS of single-precision computing and 7.8 TFLOPS of double-precision computing
- NVIDIA Tensor cores with 125 TFLOPS of single- and double-precision computing for deep learning
- Up to 30 GB/s of network bandwidth on a single ECS
- 16 GB of HBM2 GPU memory with a bandwidth of 900 GB/s
- Comprehensive basic capabilities

Networks are user-defined, subnets can be divided, and network access policies can be configured as needed. Mass storage is used, elastic capacity expansion as well as backup and restoration are supported to make data more secure. Auto Scaling allows you to add or reduce the number of ECSs quickly.

- Flexibility

  Similar to other types of ECSs, P2v ECSs can be provisioned in a few minutes.

- Excellent supercomputing ecosystem

  The supercomputing ecosystem allows you to build up a flexible, high-performance, cost-effective computing platform. A large number of HPC applications and deep-learning frameworks can run on P2v ECSs.

Supported Common Software

P2v ECSs are used in computing acceleration scenarios, such as deep learning training, inference, scientific computing, molecular modeling, and seismic analysis. If the software requires GPU CUDA parallel computing, use P2v ECSs. P2v ECSs support the following commonly used software:

- Common deep learning frameworks, such as TensorFlow, Caffe, PyTorch, and MXNet
- CUDA GPU rendering supported by RedShift for Autodesk 3dsMax and V-Ray for 3ds Max
- Agisoft PhotoScan
- MapD

Notes

- If a P2v ECS is created using a private image, make sure that the NVIDIA driver has been installed during the private image creation. If not, install the driver after the P2v ECS is created for computing acceleration. For details, see Installing the NVIDIA GPU Driver and CUDA Toolkit.
- P2v ECSs do not support specifications modification.
- P2v ECSs support the following OSs:
  - Ubuntu Server 16.04 64bit
  - CentOS 7.4 64bit
  - EulerOS 2.2 64bit
  - Windows Server 2012 R2 Standard 64bit
  - Windows Server 2016 Standard 64bit

Inference-accelerated PI1

Overview

PI1 ECSs use NVIDIA Tesla P4 GPUs dedicated for real-time AI inference. Working with P4 INT8 calculators, PI1 ECSs have shortened the inference latency by 15 times. Working with hardware decoding engines, PI1 ECSs concurrently support real-time 35-channel HD video transcoding and inference.

Specifications

Table 7-20 PI1 ECS specifications

Flavor	vCPUs	Memory (GB)	Max./Assured Bandwidth (Gbit/s)	Maximum PPS (10,000)	NIC Multi-Queue	GPUs	GPU Memory (GB)	Local Disks	Virtualization Type
pi1.2xlarge.4	8	32	5/1.6	40	2	1 x P4	1 x 8 GB	N/A	KVM
pi1.4xlarge.4	16	64	8/3.2	70	4	2 x P4	2 x 8 GB	N/A	KVM
pi1.8xlarge.4	32	128	10/6.5	140	8	4 x P4	4 x 8 GB	N/A	KVM

PI1 ECS Features

- Up to four NVIDIA Tesla P4 GPUs on an ECS
- GPU hardware passthrough
- Up to 5.5 TFLOPS of single-precision computing on a single GPU
- Up to 22 TOPS of INT8 computing on a single GPU
- 8 GB of ECC GPU memory with a bandwidth of 192 GB/s on a single GPU
- Hardware video encoding and decoding engines embedded in GPUs for concurrent real-time 35-channel HD video transcoding and inference

Supported Common Software

PI1 ECSs are used in GPU-based inference computing scenarios, such as image recognition, voice recognition, and natural language processing.

PI1 ECSs support the following commonly used software:

- Deep learning frameworks, such as TensorFlow, Caffe, PyTorch, and MXNet

Notes

- The basic resources, including vCPUs, memory, and image of a pay-per-use PI1 ECS of flavor pi1.2xlarge.4, pi1.4xlarge.4, or pi1.8xlarge.4 are not billed after the ECS is stopped, but the system disk of the ECS is still being billed according to the disk capacity. If other service products, such as EVS disks, EIPs, and bandwidth are bound to the ECS, these products are billed using their own billing mode (yearly/monthly or pay-per-use).

NOTICE

The resources of a pay-per-use PI1 ECS of flavor pi1.2xlarge.4, pi1.4xlarge.4, or pi1.8xlarge.4 are released after the ECS is stopped. If the underlying resources are insufficient when the ECS is started, starting the ECS may fail. If you want to use such an ECS for a long period of time, you are advised to keep the ECS running or select the yearly/monthly payment.

- If a PI1 ECS is created using a private image, make sure that the GPU driver has been installed during the private image creation. If not, install the driver after the PI1 ECS is

created for inference acceleration. For details, see Installing the NVIDIA GPU Driver and CUDA Toolkit.

- PI1 ECSs support the following OSs:
  - Ubuntu Server 14.04 64bit
  - Ubuntu Server 16.04 64bit
  - CentOS 7.3 64bit
- PI1 ECSs do not support specifications modification.
- PI1 ECSs support automatic recovery when the hosts accommodating such ECSs become faulty.

8 Images

Introduction

Image Management Service (IMS) allows you to create ECSs using images. An image is an ECS template that contains an OS and may also contain proprietary software and application software, such as database software.

Images can be public, private, or shared. Public images are provided by the system by default, private images are manually created, and shared images are private images that are shared by another user. You can use any type of image to create an ECS. You can also create a private image using an existing ECS. This provides you with a simple way to create ECSs that comply with your service requirements. For example, if you use web services, your image can contain a web server, static configurations, and dynamic page code. After you use this image to create an ECS, the web server will run.

Image Types

Table 8-1 Image types

Image Type	Description
Public image	A standard, widely used image. It contains an OS and preinstalled public applications and is available to all users. For details about the OSs supported by public images, see OSs Supported by ECSs.
Private image	An image that is available only to the user who created it based on an ECS or EVS backup (system disk backup). Such a private image contains an OS, preinstalled public applications, and the user's private applications.
Shared image	A private image shared by another user.

9 EVS Disks

Disk Types

ECSs support the following types of EVS disks for storing data:

- Common I/O: EVS disks of this type deliver a maximum of 2,200 IOPS. This disk type is suitable for application scenarios that require large capacity, a medium read/write rate, and fewer transactions, such as enterprise office applications and small-scale testing.
- High I/O: EVS disks of this type deliver a maximum of 5,000 IOPS and a minimum of 1 ms read/write latency. This disk type is designed to meet the needs of mainstream high-performance, high-reliability application scenarios, such as enterprise applications, large-scale development and testing, and web server logs.
- Ultra-high I/O: EVS disks of this type deliver a maximum of 33,000 IOPS and a minimum of 1 ms read/write latency. This disk type is excellent for ultra-high I/O, ultra-high bandwidth, and read/write-intensive application scenarios, such as distributed file systems in HPC scenarios or NoSQL/RDS in I/O-intensive scenarios.

EVS disks with different I/O capacities provide different features at different prices. Choose EVS disks based on your requirements. For more information about EVS disk specifications and performance, see *Elastic Volume Service User Guide*.

10 Network

VPC

Virtual Private Cloud (VPC) allows you to create customized virtual networks in your logically isolated AZ. Such networks are dedicated zones that are logically isolated for your ECSs. You can define security groups, virtual private networks (VPNs), IP address segments, and bandwidth for a VPC. This facilitates internal network configuration and management as well as secure and convenient network modification. You can also customize the ECS access rules within a security group and between security groups to strengthen ECS security protection.

For more information about VPC, see *Virtual Private Cloud User Guide*.

Subnet

A subnet is a range of IP addresses in your VPC and provides IP address management and DNS resolution functions for ECSs in it. The IP addresses of all ECSs in a subnet belong to the subnet.

Figure 10-1 Subnet

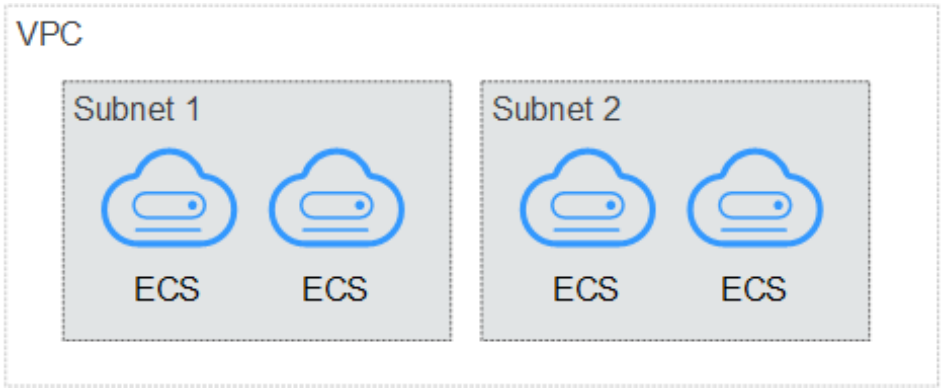

By default, ECSs in all subnets of the same VPC can communicate with each another, while ECSs in different VPCs cannot.

Security Group

A security group is a collection of access control rules for ECSs that have the same security protection requirements and are mutually trusted in a VPC. After a security group is created, you can create different access rules for the security group to protect the ECSs that are added to this security group.

Your account automatically comes with a default security group. The default security group allows all outbound traffic, denies all inbound traffic, and allows all traffic between ECSs in the group. Your ECSs in the security group can communicate with each other without the need to add rules.

Figure 10-2 Default security group

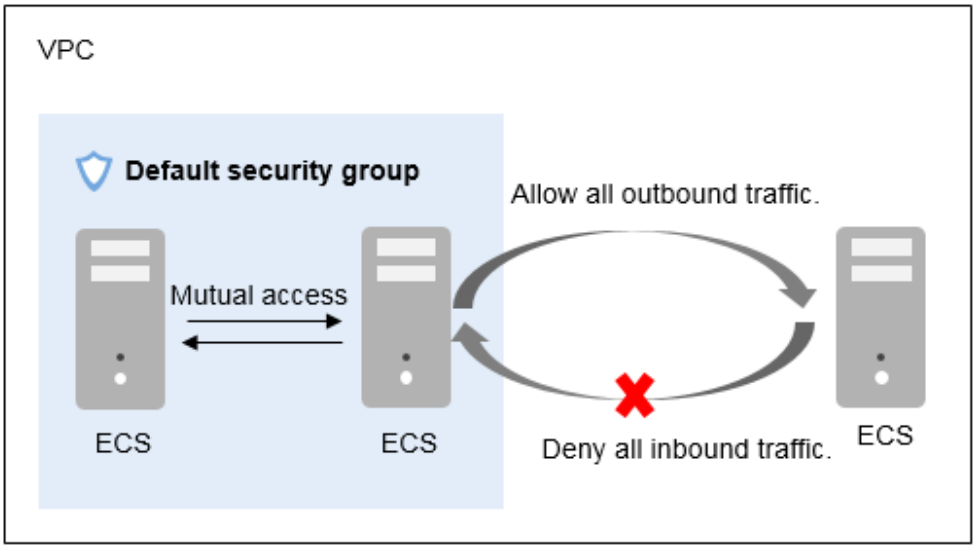

Table 10-1 describes default security group rules.

Table 10-1 Default security group rules

Direction	Protocol	Port/Range	Source/Destination	Description
Outbound	All	All	Destination: 0.0.0.0/0	Allow all outbound traffic.
Inbound	All	All	Source: ID of the current security group (for example, sg-*xxxxx*)	Allow communication among ECSs within the security group and deny all inbound traffic (incoming data packets).
Inbound	TCP	22	Source: 0.0.0.0/0	Allow all IP addresses to access Linux ECSs over SSH.
Inbound	TCP	3389	Source: 0.0.0.0/0	Allow all IP addresses to access Windows ECSs over RDP.

EIP

An EIP is a public IP address that can be directly accessed over the Internet. An EIP consists of the public IP address and public network egress bandwidth. EIPs can be bound to or unbound from ECSs, virtual IP addresses, NAT gateways, and load balancers. Various billing modes are provided to meet diversified service requirements.

Each EIP can be used by only one cloud resource at a time.

Figure 10-3 Accessing the Internet using an EIP

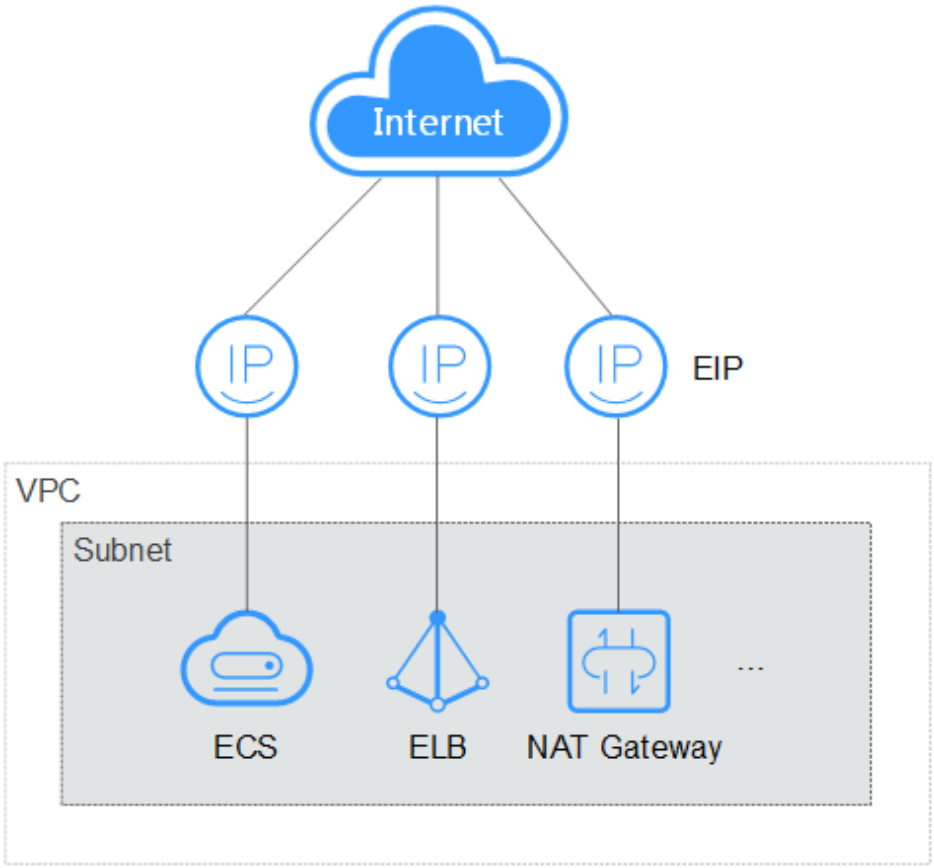

11 Security

11.1 User Encryption

User encryption allows you to use the encryption feature provided on the public cloud platform to encrypt ECS resources, improving data security. User encryption includes image encryption and EVS disk encryption.

Image Encryption

Key encryption supports encrypting private images. When creating an ECS, if you select an encrypted image, the system disk of the created ECS automatically has encryption enabled, implementing system disk encryption and improving data security.

Use either of the following methods to create an encrypted image:

- Create an encrypted image using an existing encrypted ECS.
- Create an encrypted image using an external image file.

For more information about image encryption, see Encrypting Images.

EVS Disk Encryption

EVS disk encryption supports system disk encryption and data disk encryption.

- When creating an ECS, you can encrypt added data disks.
- When creating an ECS, if you select an encrypted image, the system disk of the created ECS automatically has encryption enabled, and the encryption mode complies with the image encryption mode.

For more information about EVS disk encryption, see *EVS Disk Encryption*.

Impact on AS

If you use an encrypted ECS to create an Auto Scaling (AS) configuration, the encryption mode of the created AS configuration complies with the ECS encryption mode.

About Keys

The key required for encryption relies on Data Encryption Workshop (DEW). DEW uses a data encryption key (DEK) to encrypt data and a customer master key (CMK) to encrypt the DEK.

Figure 11-1 Data encryption process

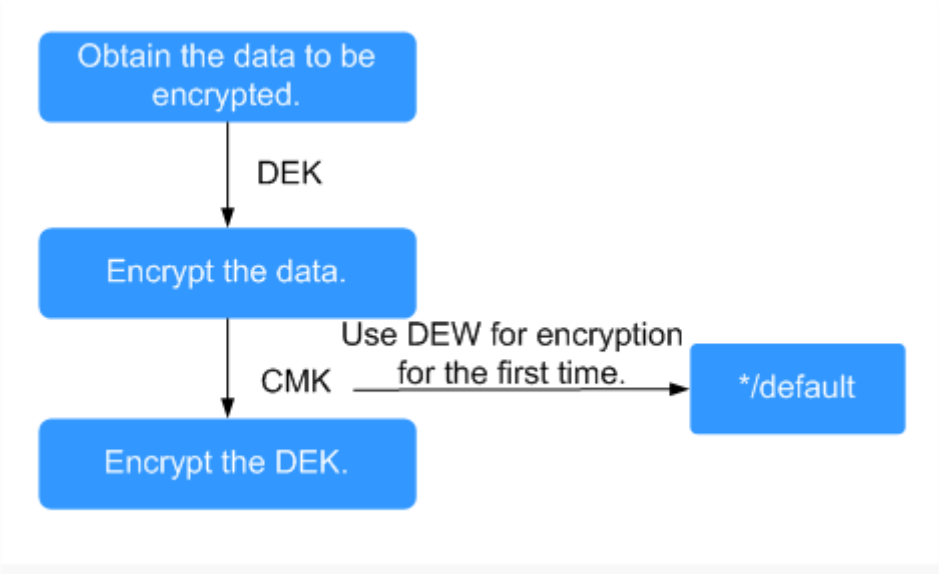

Table 11-1 describes the keys involved in the data encryption process.

Table 11-1 Keys

Name	Description	Function
DEK	An encryption key that is used for encrypting data.	Encrypts specific data.
CMK	An encryption key created using DEW for encrypting DEKs. A CMK can encrypt multiple DEKs.	Supports CMK disabling and scheduled deletion.
Default CMK	A master key automatically generated by the system when you use DEW for encryption for the first time. The name extension of a default CMK is **/default**, for example, **evs/default**.	• Supports viewing details of the default CMK on the KMS console. • Does not support CMK disabling or scheduled deletion.

> **NOTE**
>
> After disabling a CMK or scheduling the deletion of a CMK takes effect, the EVS disk encrypted using this CMK can still be used until the disk is detached from and then attached to an ECS again. During this process, the disk fails to be attached to the ECS because the CMK cannot be obtained. Therefore, the EVS disk becomes unavailable.

For details about DEW, see *Data Encryption Workshop User Guide*.

11.2 Cloud-Init

Cloud-Init is an open-source cloud initialization program, which initializes specified customized configurations, such as the hostname, key pair, and user data, of a newly created ECS.

Using Cloud-Init to initialize your ECSs will affect your ECS, IMS, and AS services.

Impact on IMS

To ensure that ECSs created using private images support customized configurations, you must install Cloud-Init or Cloudbase-Init before creating private images.

- For Windows OSs, download and install Cloudbase-Init.
- For Linux OSs, download and install Cloud-Init.

After Cloud-Init or Cloudbase-Init is installed in an image, Cloud-Init or Cloudbase-Init automatically configures initial ECS attributes when the ECS is created.

For instructions about the installation, see Installing Cloud-Init.

Impact on ECS

- When creating an ECS, if the selected image supports Cloud-Init, you can use user data injection to inject customized configuration, such as ECS login password, for initializing.
- After Cloud-Init is supported, you can view and use metadata to configure and manage running ECSs.

Impact on AS

- When creating an AS configuration, you can use user data injection to specify ECS configurations for initialization. If the AS configuration has taken effect in an AS group, the ECSs newly created in the AS group will automatically initialize their configurations.
- For an existing AS configuration, if its private image does not have Cloud-Init or Cloudbase-Init installed, the login mode of the ECSs created in the AS group where the AS configuration takes effect will be affected.

  To resolve this issue, see "How Does Cloud-Init Influence the AS Service?" in *Auto Scaling User Guide*.

Notes

- When using Cloud-Init, enable DHCP in the VPC to which the ECS belongs.
- When using Cloud-Init, ensure that security group rules in the outbound direction meet the following requirements:
  - **Protocol**: **TCP**

- **Port Range**: **80**
- **Remote End**: **169.254.0.0/16**

NOTE

If you use the default security group rules in the outbound direction, the preceding requirements are met, and the metadata can be accessed. Default security group rules in the outbound direction are as follows:

- **Protocol**: **ANY**
- **Port Range**: **ANY**
- **Remote End**: **0.0.0.0/0**

12 CPU Credits

Concept

CPU credits measure computing, storage, and network resource usage of an ECS. ECSs use CPU credits to ensure baseline performance, preventing issues caused by CPU overcommitment.

CPU-credit-based ECSs are suitable for the applications requiring baseline level of CPU performance generally and burstable performance in case of traffic bursts.

General computing-basic ECSs run based on CPU credits. For more details, see 7.6 General Computing-Basic ECSs.

Working Rules

After a CPU-credit-based ECS is created, the cloud platform automatically allocates initial CPU credits to the ECS for its burstable performance.

After the ECS runs, its credits are accrued or spent. When the actual computing performance of the ECS is higher than the baseline CPU performance, the CPU credits are spent to meet the performance requirements. When the actual computing performance is lower than the baseline CPU performance, the CPU credits are accrued until the CPU credit balance limit is reached.

NOTE

- CPU credits can be accrued. However, after the credits reach the CPU credit balance limit, any new credits that are earned will be discarded.
- Initial credits are not counted in the CPU credit balance limit.
- When an ECS starts to spend CPU credits, it preferentially uses the initial CPU credits.
- One CPU credit is equal to one vCPU running at 100% usage for one minute.
- When the actual computing performance is higher than the baseline performance, the accrued credits are spent until they are used up. Then, the actual computing performance cannot exceed the baseline performance.

Related Terms

Table 12-1 Terms related to CPU credits (taking a T6 ECS as an example)

Term	Description	Example

Term	Description	Example
Initial CPU credits	After a T6 ECS is created, the cloud platform automatically allocates CPU credits to this ECS. These credits are initial CPU credits. Initial CPU credits are allocated only after an ECS is created.	After a t6.large.1 is created, it has 60 initial CPU credits.
CPU credit balance limit	When the actual computing performance is lower than the baseline CPU performance, the CPU credits are accrued. The accrued credits will not expire on a running ECS. When the credits reach the maximum value allowed, which is specified by the CPU credit balance limit, any new credits that are earned will be discarded. The CPU credit balance limit varies depending on ECS flavors.	The CPU credit balance limit for a t6.large.1 ECS is 576. When its accrued CPU credits reach 576, no more credits will be accrued. When its accrued CPU credits are smaller than 576, the CPU credits can be accrued again.
CPU credit earn rate (credits/hour)	The number of CPU credits earned by an ECS per hour, which corresponds to CPU baseline. One CPU credit is equal to one vCPU running at 100% usage for one minute.	The CPU credit earn rate of a t6.large.1 ECS is 24, indicating that a t6.large.1 ECS can earn 24 CPU credits per hour.
CPU baseline (%)	When the number of CPU credits that an ECS spends per minute is the same as the number of CPU credits that the ECS earns per minute, the ECS runs at the CPU baseline.	The CPU baseline of a t6.large.1 ECS is 40%. When the actual computing performance of a t6.large.1 ECS reaches 40%, the number of credits spent by the ECS per minute is the same as the number of credits earned by the ECS per minute.
Average CPU baseline (%)	When an ECS runs at CPU baseline, the computing performance of each vCPU is the average CPU baseline, which is calculated using the following formula: **Average CPU baseline = CPU baseline/Number of vCPUs**	The CPU baseline of a t6.large.1 ECS is 40%, and the ECS has two vCPUs. Then, the average CPU baseline is 20%.
Spent CPU credits	When the actual computing performance of an ECS is higher than the baseline CPU performance, the CPU credits are spent to meet the performance requirements. One CPU credit is spent for one vCPU running at 100% usage for one minute. The formula for calculating the CPU	When a t6.large.1 ECS runs at the computing performance of 20% for one minute, the ECS spends 0.2 CPU credits.

Term	Description	Example
	credits spent per minute is as follows: **Number of CPU credits spent per minute = 1 CPU credit x Actual computing performance**	
Accrued CPU credits	• When the actual computing performance of an ECS is less than the baseline CPU performance, the number of CPU credits spent per minute is smaller than the number of CPU credits earned per minute. Therefore, the remaining CPU credits are accrued until the CPU credit balance limit is reached. • When the actual computing performance is higher than the baseline CPU performance, the number of CPU credits spent per minute is greater than the number of CPU credits earned per minute. In such a case, the ECS spends accrued CPU credits (initial CPU credits preferentially used) to comply with burstable CPU performance. The formula for calculating the number of CPU credits accrued per minute is as follows: **Number of CPU credits accrued per minute = 1 CPU credit x (CPU baseline – Actual computing performance)**	The CPU baseline of a t6.large.1 ECS is 40%. When the actual computing performance of the ECS is 10%, the ECS accrues 0.3 CPU credits per minute.

Impact of CPU Credits After an ECS Is Stopped

The change of CPU credits varies depending on the ECS billing mode and network type.

Table 12-2

Billing Mode	CPU Credit Change After an ECS Is Stopped
Yearly/Monthly	The existing CPU credits are retained and accrued until the CPU credit balance limit is reached.
Pay-per-use	The existing CPU credits are retained but not accrued.
Spot price	The existing CPU credits are retained but not accrued.

13 Region and AZ

Concept

A region and availability zone (AZ) identify the location of a data center. You can create resources in a specific region and AZ.

- Regions are divided based on geographical location and network latency. Public services, such as Elastic Cloud Server (ECS), Elastic Volume Service (EVS), Object Storage Service (OBS), Virtual Private Cloud (VPC), Elastic IP (EIP), and Image Management Service (IMS), are shared within the same region. Regions are classified into universal regions and dedicated regions. A universal region provides universal cloud services for common tenants. A dedicated region provides specific services for specific tenants.
- An AZ contains one or more physical data centers. Each AZ has independent cooling, fire extinguishing, moisture-proof, and electricity facilities. Within an AZ, computing, network, storage, and other resources are logically divided into multiple clusters. AZs within a region are interconnected using high-speed optical fibers to support cross-AZ high-availability systems.

Figure 13-1 shows the relationship between regions and AZs.

Figure 13-1 Regions and AZs

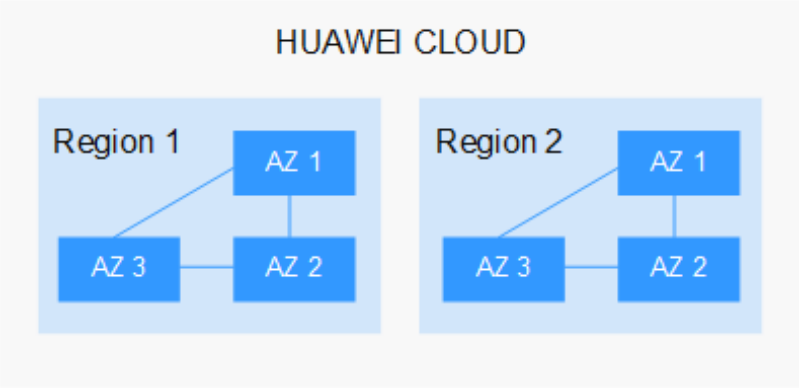

HUAWEI CLOUD provides services in many regions around the world. Select a region and AZ based on requirements. For more information, see HUAWEI CLOUD Global Regions.

Selecting a Region

When selecting a region, consider the following factors:

- Location

  It is recommended that you select the closest region for low network latency and quick access. Regions within the Chinese mainland provide the same infrastructure, BGP network quality, as well as resource operations and configurations. Therefore, if your target users are on the Chinese mainland, you do not need to consider the network latency differences when selecting a region.

  The countries and regions outside the Chinese mainland, such as Hong Kong, provide services for users outside the Chinese mainland. If your target users are on the Chinese mainland, these regions are not recommended due to high access latency.

  - If your target users are in Asia Pacific (excluding the Chinese mainland), select the **AP-Hong Kong**, **AP-Bangkok**, or **AP-Singapore** region.
  - If your target users are in Africa, select the **AF-Johannesburg** region.
  - If your target users are in Europe, select the **EU-Paris** region.
  - If your target users are in Latin America, select the **LA-Santiago** region.

  NOTE

  The **LA-Santiago** region is located in Chile.

- Resource price

  Resource prices may vary in different regions. For details, see Product Pricing Details.

Selecting an AZ

When deploying resources, consider your applications' requirements on disaster recovery (DR) and network latency.

- For high DR capability, deploy resources in different AZs within the same region.
- For low network latency, deploy resources in the same AZ.

Regions and Endpoints

Before you use an API to call resources, specify its region and endpoint. For more details, see Regions and Endpoints.

14 ECS and Other Services

Figure 14-1 shows the relationships between ECS and other services.

Figure 14-1 Relationships between ECS and other services

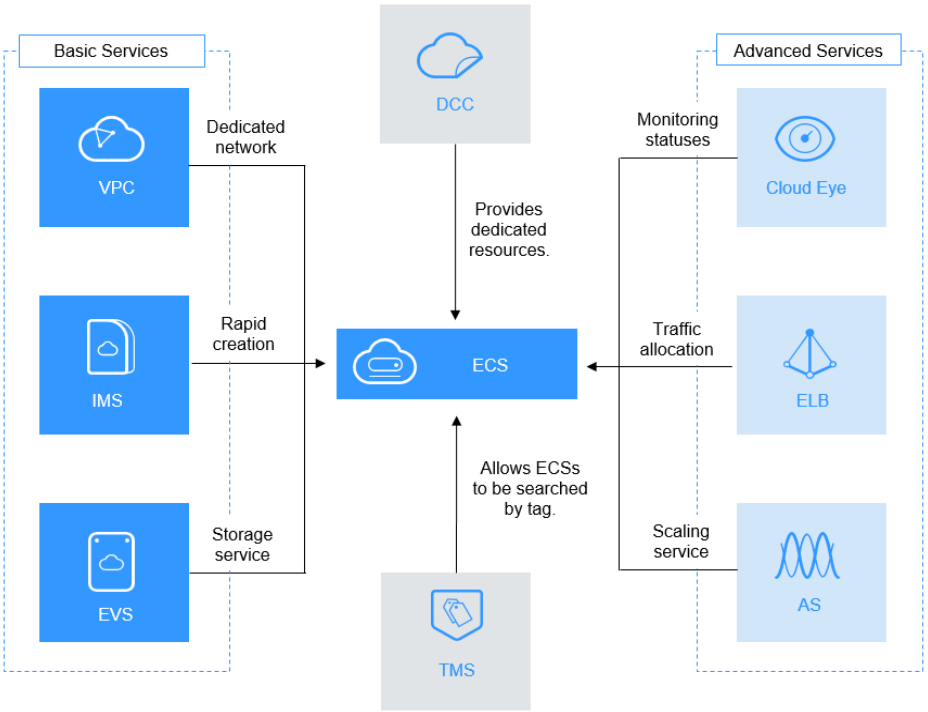

ECS-related Services

Table 14-1 ECS-related services

Service	Function	Related Operation
Auto Scaling (AS)	Automatically adjusts ECS service resources based on the configured AS policies. This improves resource usage and	• Using an Existing ECS to Create an AS Configuration • Using a New Specifications Template to Create an AS

Service	Function	Related Operation
	reduces resource costs.	Configuration
Load Balancer	Automatically distributes traffic to multiple ECSs. This enhances system service and fault tolerance capabilities.	• Backend Server
Elastic Volume Service (EVS)	Enables you to attach EVS disks to an ECS and expand their capacity.	• Attaching a Non-shared EVS Disk • Attaching a Shared EVS Disk
Virtual Private Cloud (VPC)	Enables you to configure internal networks and change network configurations by customizing security groups, VPNs, IP address segments, and bandwidth. This simplifies network management. You can also customize the ECS access rules within a security group and between security groups to strengthen ECS security protection.	• Assigning an EIP and Binding It to an ECS • Adding a Security Group Rule
Image Management Service (IMS)	Enables you to create ECSs using images. This improves the efficiency of ECS creation. You can also use an existing ECS to create a private image and export the data of the ECS system disk or data disks.	• Creating a Data Disk Image Using an ECS Data Disk • Creating a Full-ECS Image Using an ECS
Dedicated Computing Cluster (DCC)	To physically isolate your ECS, apply for a DCC before creating the ECS. After you obtain the DCC and set a region for it, your ECS is automatically allocated to the DCC.	• Enabling a DeC • Applying for DCC Resources
Cloud Eye	Allows you to check the status of monitored service objects after you have obtained an ECS. This can be done without requiring additional plug-ins be installed.	• ECS Metrics • ECS Metrics Under OS Monitoring
Data Encryption Workshop (DEW)	The encryption feature relies on DEW. You can use an encrypted image or EVS disks when creating an ECS. In such a case, you are required to use the key provided by DEW to improve data security.	• EVS Disk Encryption • Encrypting Images • Creating a Key Pair

Service	Function	Related Operation
Cloud Trace Service (CTS)	Records ECS-related operations for later query, audit, and backtrack.	• Key Operations on ECS
Cloud Server Backup Service (CSBS)	Backs up ECSs for restoration. CSBS backs up all EVS disks of an ECS, including the system disk and data disks, and uses the backup to restore the ECS.	• Creating a CSBS Backup
Volume Backup Service (VBS)	Allows you to create data backups for EVS disks and use the backups to restore the EVS disks. This maximizes data security and correctness and ensures service security.	• Creating a VBS Backup

15 Change History

Released On	Description
2019-10-28	This issue is the tenth official release. • Added 7.12 GPU-accelerated ECSs.
2019-10-24	This issue is the ninth official release. • Added 7.6 General Computing-Basic ECSs. • Added 12 CPU Credits.
2019-10-16	This issue is the eighth official release. • Added 7.10 Ultra-high I/O ECSs. • Moved "Spot ECSs" and "Reserved Instances" to User Guide.
2019-09-06	This issue is the seventh official release. • Added spot ECSs. • Modified 13 Region and AZ.
2019-07-12	This issue is the sixth official release. • Added 2 ECS Advantages. • Added 3 ECS Application Scenarios. • Added "Why ECS" in 1 What Is ECS? • Modified the ECS architecture in 1 What Is ECS? • Optimized the document structure. • Deleted API statuses in 7.2 ECS Lifecycle.
2019-06-24	This issue is the fifth official release. • Added 4 Notes on Using ECSs. • Modified 13 Region and AZ.
2019-05-07	This issue is the fourth official release. Modified the following content: • Added the newly released c3.6xlarge.2, c3.8xlarge.2, and c3.15xlarge.2 flavors in 7.5 General Computing-plus ECSs. • Added the newly released d2.2xlarge.8, d2.4xlarge.8, d2.6xlarge.8, d2.8xlarge.8, and d2.12xlarge.8 flavors in 7.9 Disk-intensive ECSs.

Released On	Description
2019-04-09	This issue is the third official release. Modified the following content: • Added the newly released m3.15xlarge.8 flavor in 7.7 Memory-optimized ECSs.
2019-03-04	This issue is the second official release. • Added reserved instances. • Modified 14 ECS and Other Services.
2018-11-19	This issue is the first official release.