

Elastic Cloud Server

Service Overview

Issue 04
Date 2024-12-23



Copyright © Huawei Cloud Computing Technologies Co., Ltd. 2024. All rights reserved.

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Cloud Computing Technologies Co., Ltd.

Trademarks and Permissions



HUAWEI and other Huawei trademarks are the property of Huawei Technologies Co., Ltd.

All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

The purchased products, services and features are stipulated by the contract made between Huawei Cloud and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Contents

1 What Is ECS?	1
2 ECS Advantages	4
3 ECS Application Scenarios	6
4 ECS Types and Specifications	8
4.1 ECS Overview.....	8
4.2 ECS Lifecycle.....	8
4.3 ECS Types.....	10
4.4 CPU Credits.....	12
4.5 x86 ECS Specifications.....	16
4.5.1 A Summary List of x86 ECS Specifications.....	16
4.5.2 General Computing ECSs.....	39
4.5.3 General Computing-plus ECSs.....	41
4.5.4 General Computing-Basic ECSs.....	51
4.5.5 Memory-optimized ECSs.....	55
4.5.6 Large-Memory ECSs.....	61
4.5.7 Disk-intensive ECSs.....	63
4.5.8 GPU-accelerated ECSs.....	68
4.5.9 AI-accelerated ECSs.....	74
5 Images	77
5.1 Image Types.....	77
5.2 Cloud-Init.....	78
6 EVS Disks	80
7 Network	82
8 Security	85
8.1 Identity Authentication and Access Control.....	85
8.1.1 Access Control for ECS.....	85
8.2 Data Protection.....	87
8.2.1 Hypervisor Security.....	87
8.2.2 User Encryption.....	88
8.2.3 HSS.....	90

8.3 Fault Recovery.....	91
8.4 License Types.....	92
9 Billing.....	93
10 Constraints.....	99
11 ECS and Other Services.....	101
12 Permissions Management.....	104
13 Region and AZ.....	110

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.

You can create an ECS by specifying its vCPUs, memory, OS, and login mode. After creating an ECS, you can use it on the cloud like using your local PC or physical server. You can also modify its specifications if necessary. ECS lets your applications run in a reliable, secure, efficient computing environment.

- For details about vCPUs, memory, and specifications of an ECS, see [A Summary List of x86 ECS Specifications](#).
- For details about the operating systems supported by an ECS, see [Image Types](#).
- For details about the login authentication modes, see [Logging In to an ECS](#).

Why ECS

- Rich specifications: A variety of ECS types with custom specifications are available for different scenarios.
- Various image types: Public, private, and shared images are available for you to choose from.
- A broad range of disk types: High I/O, extreme SSD, and ultra-high I/O disks are provided to meet the requirements of different service scenarios.
- Reliable data: High-throughput virtual block storage uses the distributed architecture to ensure high availability and it can be scaled out as needed.
- Security protection: The network is isolated and protected using security group rules. Security services, such as Anti-DDoS, Web Application Firewall (WAF), and Vulnerability Scan Service (VSS) can also be used to further enhance ECS security.
- Auto scaling: Elastic computing resources can be automatically adjusted to suit your needs.
- Efficient O&M: ECSs can be efficiently managed through the management console, remote terminals, or APIs with full rights.
- Cloud monitoring: Cloud Eye samples monitored metrics in real time, generates alarms when detecting problems, and immediately notifies related personnel of the alarms.

- Load balancing: Elastic Load Balance (ELB) evenly distributes incoming traffic across ECSs to prevent overload on an individual ECS. Applications are more tolerant of errors and bursty traffic.

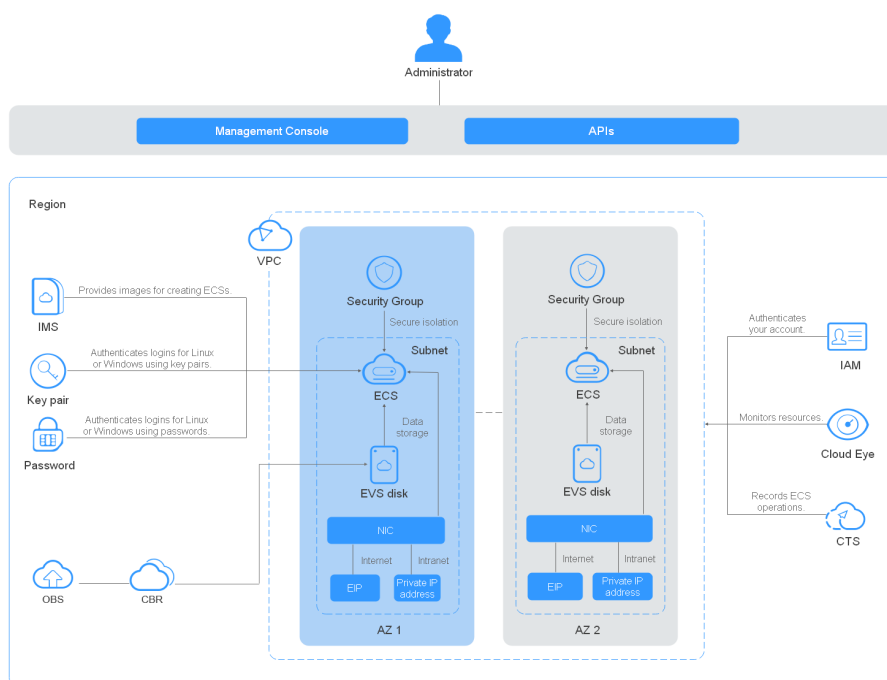
For more details, see [ECS Advantages](#) and [ECS Application Scenarios](#).

System Architecture

ECS works with other products and services to provide computing, storage, and network resources.

- You can deploy ECSs across different availability zones (AZs) that are connected over an intranet. If one AZ becomes unavailable, ECSs in other AZs can continue to provide services.
- Virtual Private Cloud (VPC) helps you build your own dedicated network on the cloud. You can set subnets and security groups within your VPC for further isolation. You can also bind an EIP to your ECSs for Internet access.
- With the Image Management Service (IMS), you can use an image to create ECSs. You can also use an existing ECS to create a private image and use the private image to create the same ECSs for rapid service deployment.
- Elastic Volume Service (EVS) provides storage space. Volume Backup Service (VBS) provides data backup and restoration.
- Cloud Eye lets you keep a close eye on the performance and resource utilization of ECSs, ensuring ECS reliability and availability.
- Cloud Backup and Recovery (CBR) backs up data for EVS disks and ECSs, and uses snapshots and backups to restore the EVS disks and ECSs.

Figure 1-1 System architecture



Access Methods

You can access ECS through the web-based management console or HTTPS-based application programming interfaces (APIs).

- Accessing ECSs through APIs

Use this method if you intend to integrate ECSs into a third-party system for secondary development. For details, 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.

Log in to the management console with your account and choose **Elastic Cloud Server** on the homepage.

2 ECS Advantages

ECS supports automated scaling of compute resources based on traffic changes and predefined scaling policies. You can customize ECS specifications including vCPUs, memory, and bandwidth to let your applications run in a flexible, efficient environment.

Reliability

- A broad range of EVS disk types

You can choose from the following disk types to meet your specific demand:

- High I/O EVS disks: They provide high performance, scalability, and reliability. If your applications require high performance, high read/write speed, and instant data storage, you can store data on high I/O EVS disks.
- General purpose SSDs: Such disks are cost-effective. You can use them for high-throughput, low-latency applications.
- Ultra-high I/O EVS disks: They provide high performance and low latency. They are ideal for read/write-intensive applications requiring extremely high performance and low latency.
- Extreme SSDs: They use RDMA that controls congestion at low latency. They are ideal for applications that require ultra-large bandwidth and ultra-low latency.

For more information about EVS disk specifications and performance, see [Elastic Volume Service User Guide](#).

- Distributed architecture

ECS provides scalable, reliable, and high-throughput virtual block storage on a 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

You can set automatic backup policies to back up in-service ECSs and EVS disks. You can also configure policies on the management console or use an API to back up the data of ECSs and EVS disks at a specified time.

Security

- Multi-dimensional protection
A number of security services, such as Web Application Firewall (WAF) and Vulnerability Scan Service (VSS) are available.
- Security evaluation
Cloud security evaluation and security configuration check help you identify security vulnerabilities and threats, reducing or eliminating your loss from viruses or attacks.
- Intelligent process management
You can customize an allowlist to automatically prohibit the execution of unauthorized programs.
- Vulnerability scan
Comprehensive scan services are available, including general web vulnerability scan, third-party application vulnerability scan, port detection, and fingerprint identification.

Hardware and Software

- Professional hardware devices
You can deploy ECSs on professional hardware devices that allow in-depth virtualization optimization, delivering superior virtual server performance.
- Virtual resources accessible anytime, anywhere
You can obtain scalable, dedicated resources from the virtual resource pool anytime, anywhere, so your applications can run in reliable, secure, flexible, and efficient environments. You can use your ECS like the way you are using your local computer.

Scalability

- Automated scaling 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 at a regular interval or a specified time based on the predicted load or a pre-set plan.
- Flexible adjustment of ECS specifications
ECS specifications and bandwidth can be flexibly adjusted based on service requirements.

3 ECS Application Scenarios

Internet

- No special requirements on CPUs, memory, disk space, or bandwidth
- High security and reliability standards
- Deploying an application on one or only a few ECSs to minimize upfront investment and maintenance costs, such as website development and testing, and small 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.

For details, see [General Computing ECSs](#) and [General Computing-plus ECSs](#).

E-Commerce

- Large amount of memory
- Quick processing of large volumes of data
- Large incoming traffic

Use memory-optimized ECSs, which provide a large memory, ultra-high I/O EVS disks, and the needed bandwidths. This ECS type is suitable for precision marketing, E-Commerce, and mobile apps.

For details, see [Memory-optimized ECSs](#).

Graphics Rendering

- High-quality graphics and video
- Large amount of memory and rapid processing of large volumes of data
- Fast network with high I/O
- High GPU performance for graphics rendering and engineering drawing

Use GPU-accelerated ECSs, which adopt NVIDIA Tesla M60 hardware virtualization and provide cost-effective graphics acceleration. These ECSs support DirectX and OpenGL, and provide up to 1 GiB of GPU memory and 4096 x 2160 resolution.

For details, see [GPU-accelerated ECSs](#).

Data Analytics

- 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 use hard disk drives (HDDs) and a default network bandwidth of 10GE, providing high packets per second (PPS) and low network latency. Each disk-intensive ECS supports up to 24 local disks, 48 vCPUs, and 384 GiB of memory.

For details, see [Disk-intensive ECSs](#).

4 ECS Types and Specifications

4.1 ECS 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. This ensures a reliable, secure, efficient computing environment.

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 [ECS Types](#).
- For details about all ECS statuses in a lifecycle, see [ECS Lifecycle](#).
- For details about ECS specifications, see [A Summary List of x86 ECS Specifications](#).

4.2 ECS Lifecycle

The ECS lifecycle refers to the entire journey an ECS goes through, from creation to deletion (or release).

Table 4-1 ECS statuses

Status	Status Attribute	Description
Creating	Intermediate	The ECS is being created.
Starting	Intermediate	The ECS is being started.

Status	Status Attribute	Description
Running	Stable	The ECS is running properly.
Stopping	Intermediate	The ECS is being stopped.
Stopped	Stable	The ECS has been stopped.
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 new size.
Deleting	Intermediate	The ECS is being deleted. If the ECS remains in this state for a long time, exceptions may have occurred. In such a case, contact technical support.
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. Contact technical support for assistance.
Reinstalling	Intermediate	The ECS has received a request to reinstall the OS and has begun the reinstallation.
Reinstalling failed	Stable	The ECS received a request to reinstall the OS, but the reinstallation failed. Contact technical support for assistance.
Changing OS	Intermediate	The ECS received a request to change the OS and has begun implementing the changes.
Failed to change the OS	Stable	The ECS has received a request to change the OS, but due to exceptions, the change attempt failed. Contact technical support for assistance.
Forcibly restarting	Intermediate	The ECS is being forcibly restarted.
Reverting resizing	Intermediate	The ECS is rolling back a resizing operation.

Status	Status Attribute	Description
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.

4.3 ECS Types

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

- General computing
- Dedicated general-purpose
- Memory-optimized
- Disk-intensive
- GPU-accelerated

ECS Flavor Naming Rules

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

Example: s6.medium.4

The format is defined as follows:

- **A** specifies the ECS type. For example, **s** indicates a general-purpose ECS, **c** a general computing-plus ECS, and **m** a memory-optimized ECS.
- **B** specifies the type ID. For example, **6** in **s6** indicates the sixth-generation general-purpose ECS.
- **C** specifies the flavor size, such as medium, large, xlarge, 2xlarge, 4xlarge, or 8xlarge.
- **D** specifies the ratio of memory to vCPUs expressed in a digit. For example, value **4** indicates that the ratio of memory to vCPUs is 4.

vCPU

ECS supports hyper-threading, which enables two threads to run concurrently on a single CPU core. Each thread is represented as a virtual CPU (vCPU) and a CPU core contains two vCPUs (logical cores).

Hyper-threading is enabled for most ECS flavors by default. If hyper-threading is disabled during the ECS creation or flavor change, the number of vCPUs queried from the ECS is half of the number of vCPUs defined by the ECS flavor.

For example, a 2-core physical CPU contains 4 vCPUs (threads).

Network QoS

Network QoS uses basic technologies to improve the quality of network communication. A network with QoS enabled offers predictable network performance and effectively allocates network bandwidth to use network resources.

The intranet bandwidth and PPS of an ECS are determined by the ECS flavor.

- **Assured intranet bandwidth:** indicates the guaranteed bandwidth allocated to an ECS when there is a network bandwidth contention in the entire network.
- **Maximum intranet bandwidth:** indicates the maximum bandwidth that can be allocated to an ECS when the ECS does not compete for network bandwidth (other ECSs on the host do not have high requirements on network bandwidth).
- **Maximum intranet PPS:** indicates the maximum ECS capability in sending and receiving packets.

PPS: packets per second, indicates the number of packets received and sent per second. It is usually used to measure the network performance.

- **NIC multi-queues:** allocates NIC interrupt requests to multiple vCPUs for higher PPS performance and bandwidth
- **Maximum NICs:** indicates the maximum number of NICs that can be attached to an ECS.
- **Maximum supplementary NICs:** indicates the maximum number of supplementary NICs that can be attached to an ECS.

NOTE

- For instructions about how to test packet transmit and receive, see [How Can I Test Network Performance?](#)
- For instructions about how to enable NIC multi-queue, see [Enabling NIC Multi-Queue.](#)
- The maximum bandwidth is the total bandwidth allocated to an ECS. If an ECS has multiple NICs, the sum of the maximum bandwidths allocated to all NICs cannot exceed the maximum bandwidth allocated to the ECS.
- A NIC refers to an elastic network interface. You can create and configure network interfaces and attach them to your ECSs for flexible and highly available network configurations.
For details, see [Elastic Network Interface.](#)
- A supplementary NIC is a supplement to NICs. If the number of NICs that can be attached to your ECSs cannot meet your requirements, you can use supplementary NICs.
For details, see [Supplementary Network Interface.](#)

Dedicated and Shared ECSs

Table 4-2 Differences between dedicated and shared ECSs

Dimension	Dedicated ECS	Shared ECS
CPU Allocation	CPUs are exclusively used and there is no CPU contention.	CPUs are shared and CPU contention may occur.

Dimension	Dedicated ECS	Shared ECS
Feature	<ul style="list-style-type: none">• High performance• Dedicated and stable computing, storage, and network resources• High costs	<ul style="list-style-type: none">• Unstable performance when loads are high• Shared computing, storage, and network resources• Low costs
Application Scenario	For enterprises that have high requirements on service stability	For small- and medium-sized websites or individuals that have requirements on cost-effectiveness
ECS Specifications	Specifications except general computing	x86 computing: <ul style="list-style-type: none">• General computing

4.4 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 vCPU performance generally and burstable performance in case of traffic bursts.

Currently, general computing S6 and general computing-basic T6 ECSs run based on CPU credits (no additional billing is incurred). For details about the specifications, see [General Computing S6](#) and [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.
- General computing S6 ECS burstable performance is neither restricted by CPU credits nor billed for additional credits.
- General computing-basic T6 ECSs are shared ECSs restricted by CPU credits. If the actual computing performance is higher than the baseline performance, the accrued credits are spent. When they are used up, the actual computing performance cannot exceed the baseline performance.

Related Terms

Table 4-3 Terms related to CPU credits (taking a T6 ECS as an 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.

Term	Description	Example
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 credits spent per minute is as follows: Number of CPU credits spent per minute = 1 CPU credit x Actual computing performance	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
Accrued CPU credits	<ul style="list-style-type: none">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. <p>The formula for calculating the number of CPU credits accrued per minute is as follows:</p> <p>Number of CPU credits accrued per minute = 1 CPU credit x (CPU baseline – Actual computing performance)</p>	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 4-4 Billing modes and CPU credits

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.

4.5 x86 ECS Specifications

4.5.1 A Summary List of x86 ECS Specifications

For details about how to purchase an x86 ECS, see [Purchasing an ECS](#).

For details about specifications, see [ECS Types](#).

General Computing ECSs

For more details, see [General Computing ECSs](#).

Table 4-5 General computing ECS features

Series	Compute	Disk Type	Network
S6	<ul style="list-style-type: none"> vCPU to memory ratio: 1:1, 1:2, or 1:4 Number of vCPUs: 1 to 8 2nd Generation Intel® Xeon® Scalable Processor Basic/Turbo frequency: 2.6 GHz/3.5 GHz 	<ul style="list-style-type: none"> High I/O General Purpose SSD Ultra-high I/O Extreme SSD General Purpose SSD V2 	<ul style="list-style-type: none"> An ECS with higher specifications has better network performance. Maximum PPS: 500,000 Maximum intranet bandwidth: 3 Gbit/s Maximum number of network connections: 300,000

Table 4-6 S6 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Virtualization
s6.small.1	1	1	0.8/0.1	10	30	1	2	KVM
s6.medium.2	1	2	0.8/0.1	10	30	1	2	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Virtualization
s6.large.2	2	4	1.5/0.2	15	30	1	2	KVM
s6.xlarge.2	4	8	2/0.35	25	30	1	2	KVM
s6.2xlarge.2	8	16	3/0.75	50	30	2	2	KVM
s6.medium.4	1	4	0.8/0.1	10	30	1	2	KVM
s6.large.4	2	8	1.5/0.2	15	30	1	2	KVM
s6.xlarge.4	4	16	2/0.35	25	30	1	2	KVM
s6.2xlarge.4	8	32	3/0.75	50	30	2	2	KVM

General Computing-plus ECSs

For more details, see [General Computing-plus ECSs](#).

Table 4-7 General computing-plus ECS features

Series	Compute	Disk Type	Network
C7n	<ul style="list-style-type: none"> vCPU to memory ratio: 1:2 or 1:4 Number of vCPUs: 2 to 96 3rd Generation Intel® Xeon® Scalable Processor Basic/Turbo frequency: 2.6 GHz/3.4 GHz 	<ul style="list-style-type: none"> High I/O General Purpose SSD Ultra-high I/O Extreme SSD General Purpose SSD V2 	<ul style="list-style-type: none"> Ultra-high PPS throughput An ECS with higher specifications has better network performance. Maximum PPS: 8,500,000 Maximum intranet bandwidth: 40 Gbit/s Maximum number of network connections: 8 million
C7t	<ul style="list-style-type: none"> vCPU to memory ratio: 1:4 Number of vCPUs: 2 to 128 3rd Generation Intel® Xeon® Scalable Processor Basic/Turbo frequency: 2.6 GHz/3.4 GHz and 2.8 GHz/3.5 GHz 	<ul style="list-style-type: none"> General Purpose SSD Ultra-high I/O Extreme SSD General Purpose SSD V2 Extreme SSD V2 	<ul style="list-style-type: none"> Ultra-high PPS throughput An ECS with higher specifications has better network performance. Maximum PPS: 30 million PPS Maximum intranet bandwidth: 90 Gbit/s Maximum number of network connections: 10 million
C6s	<ul style="list-style-type: none"> vCPU to memory ratio: 1:2 or 1:4 Number of vCPUs: 2 to 64 2nd Generation Intel® Xeon® Scalable Processor Basic/Turbo frequency: 2.6 GHz/3.5 GHz 	<ul style="list-style-type: none"> High I/O General Purpose SSD Ultra-high I/O Extreme SSD General Purpose SSD V2 	<ul style="list-style-type: none"> Ultra-high PPS throughput An ECS with higher specifications has better network performance. Maximum PPS: 8,500,000 Maximum intranet bandwidth: 30 Gbit/s Maximum number of network connections: 5 million

Series	Compute	Disk Type	Network
C6	<ul style="list-style-type: none"> vCPU to memory ratio: 1:2 or 1:4 Number of vCPUs: 2 to 64 2nd Generation Intel® Xeon® Scalable Processor Basic/Turbo frequency: 3.0 GHz/3.4 GHz 	<ul style="list-style-type: none"> High I/O General Purpose SSD Ultra-high I/O Extreme SSD General Purpose SSD V2 	<ul style="list-style-type: none"> Ultra-high PPS throughput An ECS with higher specifications has better network performance. Maximum PPS: 10,000,000 Maximum intranet bandwidth: 40 Gbit/s Maximum number of network connections: 10 million

Table 4-8 C7n ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Max. Supplementarily NICs	EV S Basic Bandwidth/Burst Bandwidth (Gbit/s)	Virtualization
c7n.large.2	2	4	4/0.8	40	50	2	2	16	1.5/6	KVM
c7n.xlarge.2	4	8	8/1.6	80	50	2	3	32	2/6	KVM
c7n.2xlarge.2	8	16	15/3	150	100	4	4	64	3/6	KVM
c7n.3xlarge.2	12	24	17/5	200	150	4	6	96	4/6	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Max. SupplementarY NICs	EV S Basic Bandwidth/Burst Bandwidth (Gbit/s)	Virtualization
c7n.4xlarge.2	16	32	20/6	280	150	8	8	128	5/6	KVM
c7n.6xlarge.2	24	48	25/9	400	200	8	8	192	6/None	KVM
c7n.8xlarge.2	32	64	30/12	550	300	16	8	256	8/None	KVM
c7n.12xlarge.2	48	96	35/18	750	400	16	8	256	12/None	KVM
c7n.16xlarge.2	64	128	36/24	800	500	28	8	256	16/None	KVM
c7n.24xlarge.2	96	192	40/36	850	800	32	8	256	24/None	KVM
c7n.large.4	2	8	4/0.8	40	50	2	2	16	1.5/6	KVM
c7n.xlarge.4	4	16	8/1.6	80	50	2	3	32	2/6	KVM
c7n.2xlarge.4	8	32	15/3	150	100	4	4	64	3/6	KVM
c7n.3xlarge.4	12	48	17/5	200	150	4	6	96	4/6	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Max. Supplementar y NICs	EV S Basic Bandwidth/Burst Bandwidth (Gbit/s)	Virtualization
c7n.4xlarge.4	16	64	20/6	280	150	8	8	128	5/6	KVM
c7n.6xlarge.4	24	96	25/9	400	200	8	8	192	6/None	KVM
c7n.8xlarge.4	32	128	30/12	550	300	16	8	256	8/None	KVM
c7n.12xlarge.4	48	192	35/18	750	400	16	8	256	12/None	KVM
c7n.16xlarge.4	64	256	36/24	800	500	28	8	256	16/None	KVM
c7n.24xlarge.4	96	384	40/36	850	800	32	8	256	24/None	KVM

Table 4-9 C7t ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Max. Supplementary NICs	EV S Basic Bandwidth/Burst Bandwidth (Gbit/s)	Virtualization
c7t.large.4	2	8	10/1.28	80	50	2	2	16	1.5/6	QingTian
c7t.xlarge.4	4	16	16/2.56	150	50	2	3	32	2/6	
c7t.2xlarge.4	8	32	20/4	200	100	4	4	64	3/6	
c7t.3xlarge.4	12	48	34/6.4	300	150	4	6	96	4/6	
c7t.4xlarge.4	16	64	40/8	400	150	8	8	128	5/6	
c7t.6xlarge.4	24	96	50/14.4	600	200	8	8	192	6/None	
c7t.8xlarge.4	32	128	60/16	800	300	16	8	256	8/None	
c7t.12xlarge.4	48	192	70/28.8	1,200	400	16	8	256	10/None	
c7t.16xlarge.4	64	256	72/32	1500	500	28	8	256	16/None	

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Max. Supplementary NICs	EV S Basic Bandwidth/Burst Bandwidth (Gbit/s)	Virtualization
c7t.24xlarge.4	96	384	80/40	2400	800	32	8	256	20/None	

Table 4-10 C6s ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Virtualization
c6s.large.2	2	4	1/1	30	50	2	2	KVM
c6s.xlarge.2	4	8	2/2	60	50	2	3	KVM
c6s.2xlarge.2	8	16	4/4	120	100	4	4	KVM
c6s.3xlarge.2	12	24	5.5/5.5	180	150	4	6	KVM
c6s.4xlarge.2	16	32	7.5/7.5	240	150	8	8	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Virtualization
c6s.6xlarge.2	24	48	11/11	350	200	8	8	KVM
c6s.8xlarge.2	32	64	15/15	450	300	16	8	KVM
c6s.12xlarge.2	48	96	22/22	650	400	16	8	KVM
c6s.16xlarge.2	64	128	30/30	850	500	32	8	KVM
c6s.large.4	2	8	1/1	30	50	2	2	KVM
c6s.xlarge.4	4	16	2/2	60	50	2	3	KVM
c6s.2xlarge.4	8	32	4/4	120	100	4	4	KVM
c6s.3xlarge.4	12	48	5.5/5.5	180	150	4	6	KVM
c6s.4xlarge.4	16	64	7.5/7.5	240	150	8	8	KVM
c6s.6xlarge.4	24	96	11/11	350	200	8	8	KVM
c6s.8xlarge.4	32	128	15/15	450	300	16	8	KVM
c6s.12xlarge.4	48	192	22/22	650	400	16	8	KVM
c6s.16xlarge.4	64	256	30/30	850	500	32	8	KVM

Table 4-11 C6 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	EVS Basic Bandwidth/Burst Bandwidth (Gbit/s)	Virtualization
c6.large.2	2	4	4/1.2	40	50	2	2	1/5	KVM
c6.xlarge.2	4	8	8/2.4	80	50	2	3	1.5/5	KVM
c6.2xlarge.2	8	16	15/4.5	150	100	4	4	2/5	KVM
c6.3xlarge.2	12	24	17/7	200	150	4	6	2.5/5	KVM
c6.4xlarge.2	16	32	20/9	280	150	8	8	3.5/5	KVM
c6.6xlarge.2	24	48	25/14	400	200	8	8	4/5	KVM
c6.8xlarge.2	32	64	30/18	550	300	16	8	7/10	KVM
c6.12xlarge.2	48	96	35/27	750	400	16	8	10/15	KVM
c6.16xlarge.2	64	128	40/36	1,000	500	32	8	20/None	KVM
c6.22xlarge.2	88	176	44/40	1,200	1,000	32	8	20/None	KVM
c6.large.4	2	8	4/1.2	40	50	2	2	1/5	KVM
c6.xlarge.4	4	16	8/2.4	80	50	2	3	1.5/5	KVM
c6.2xlarge.4	8	32	15/4.5	150	100	4	4	2/5	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	EVS Basic Bandwidth/Burst Bandwidth (Gbit/s)	Virtualization
c6.3xlarge.4	12	48	17/7	200	150	4	6	2.5/5	KVM
c6.4xlarge.4	16	64	20/9	280	150	8	8	3.5/5	KVM
c6.6xlarge.4	24	96	25/14	400	200	8	8	4/5	KVM
c6.8xlarge.4	32	128	30/18	550	300	16	8	7/10	KVM
c6.12xlarge.4	48	192	35/27	750	400	16	8	10/15	KVM
c6.16xlarge.4	64	256	40/36	1,000	500	32	8	20/None	KVM
c6.22xlarge.4	88	352	44/40	1,200	1,000	28	8	20/None	KVM
c6.22xlarge.4. physical	88	384	44/40	1,000	1,000	16	33	20/None	Bare metal

General Computing-Basic ECSs

For more details, see [General Computing-Basic ECSs](#).

Table 4-12 General computing-basic ECS features

Series	Compute	Disk Type	Network
T6	<ul style="list-style-type: none"> vCPU to memory ratio: 1:1, 1:2, or 1:4 Number of vCPUs: 1 to 16 Intel® Xeon® Scalable Processor Basic/Turbo frequency: 2.2 GHz/3.0 GHz 	<ul style="list-style-type: none"> High I/O General Purpose SSD Ultra-high I/O Extreme SSD General Purpose SSD V2 	<ul style="list-style-type: none"> An ECS with higher specifications has better network performance. Maximum packets per second (PPS): 600,000 Maximum intranet bandwidth: 3 Gbit/s Maximum number of network connections: 500,000

Table 4-13 T6 ECS specifications

Flavor	vCPUs	Memory (GiB)	CPU Baseline (%)	Average CPU Baseline (%)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Virtualization
t6.small.1	1	1	10	10	0.3/0.05	6	30	1	1	KVM
t6.large.1	2	2	40	20	0.5/0.1	10	30	1	1	KVM
t6.xlarge.1	4	4	80	20	1/0.2	20	30	2	2	KVM
t6.2xlarge.1	8	8	120	15	2/0.4	40	30	4	2	KVM

Flavor	vCPUs	Memory (GiB)	CPU Baseline (%)	Average CPU Baseline (%)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Virtualization
t6.4xlarge.1	16	16	240	15	3/0.8	60	50	8	2	KVM
t6.medium.2	1	2	10	10	0.3/0.05	6	30	1	1	KVM
t6.large.2	2	4	40	20	0.5/0.1	10	30	1	1	KVM
t6.xlarge.2	4	8	80	20	1/0.2	20	30	2	2	KVM
t6.2xlarge.2	8	16	120	15	2/0.4	40	30	4	2	KVM
t6.4xlarge.2	16	32	240	15	3/0.8	60	50	8	2	KVM
t6.large.4	2	8	40	20	0.5/0.1	10	30	1	1	KVM
t6.xlarge.4	4	16	80	20	1/0.2	20	30	2	2	KVM
t6.2xlarge.4	8	32	120	15	2/0.4	40	30	4	2	KVM

Memory-optimized ECSs

For more details, see [Memory-optimized ECSs](#).

Table 4-14 Memory-optimized ECS features

Series	Compute	Disk Type	Network
M7n	<ul style="list-style-type: none"> vCPU to memory ratio: 1:8 Number of vCPUs: 2 to 96 3rd Generation Intel® Xeon® Scalable Processor Basic/Turbo frequency: 2.6 GHz/3.4 GHz 	<ul style="list-style-type: none"> High I/O General Purpose SSD Ultra-high I/O Extreme SSD General Purpose SSD V2 	<ul style="list-style-type: none"> Support for IPv6 Ultra-high PPS throughput An ECS with higher specifications has better network performance. Maximum PPS: 8,500,000 Maximum intranet bandwidth: 40 Gbit/s Maximum number of network connections: 8 million
M6s	<ul style="list-style-type: none"> vCPU to memory ratio: 1:8 Number of vCPUs: 2 to 64 2nd Generation Intel® Xeon® Scalable Processor Basic/Turbo frequency: 3.0 GHz/3.4 GHz 	<ul style="list-style-type: none"> High I/O General Purpose SSD Ultra-high I/O Extreme SSD General Purpose SSD V2 	<ul style="list-style-type: none"> Ultra-high PPS throughput An ECS with higher specifications has better network performance. Maximum PPS: 10 million PPS Maximum intranet bandwidth: 34 Gbit/s

Series	Compute	Disk Type	Network
M6	<ul style="list-style-type: none"> vCPU to memory ratio: 1:8 Number of vCPUs: 2 to 64 2nd Generation Intel® Xeon® Scalable Processor Basic/Turbo frequency: 3.0 GHz/3.4 GHz 	<ul style="list-style-type: none"> High I/O General Purpose SSD Ultra-high I/O Extreme SSD General Purpose SSD V2 	<ul style="list-style-type: none"> Ultra-high PPS throughput An ECS with higher specifications has better network performance. Maximum PPS: 10,000,000 Maximum intranet bandwidth: 40 Gbit/s Maximum number of network connections: 10 million

Table 4-15 M7n ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	EVS Basic Bandwidth/Burst Bandwidth (Gbit/s)	Virtualization
m7n.large.8	2	16	4/0.8	40	50	2	2	1.5/6	KVM
m7n.xlarge.8	4	32	8/1.6	80	50	2	3	2/6	KVM
m7n.2xlarge.8	8	64	15/3	150	100	4	4	3/6	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	EV S Basic Bandwidth/Burst Bandwidth (Gbit/s)	Virtualization
m7n.3xlarge.8	12	96	17/5	200	150	4	6	4/6	KVM
m7n.4xlarge.8	16	128	20/6	280	150	8	8	5/6	KVM
m7n.6xlarge.8	24	192	25/9	400	200	8	8	6/None	KVM
m7n.8xlarge.8	32	256	30/12	550	300	16	8	8/None	KVM
m7n.12xlarge.8	48	384	35/18	750	400	16	8	12/None	KVM
m7n.16xlarge.8	64	512	36/24	800	500	28	8	16/None	KVM
m7n.24xlarge.8	96	768	40/36	850	800	32	8	24/None	KVM

Table 4-16 M6s ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Virtualization
m6s.large.8	2	16	3/1	30	50	2	2	KVM
m6s.xlarge.8	4	32	6/2	60	50	2	3	KVM
m6s.2xlarge.8	8	64	12/4	120	100	4	4	KVM
m6s.3xlarge.8	12	96	14/5.5	160	150	4	6	KVM
m6s.4xlarge.8	16	128	16/7.5	220	150	8	8	KVM
m6s.6xlarge.8	24	192	20/11	320	200	8	8	KVM
m6s.8xlarge.8	32	256	25/15	450	300	16	8	KVM
m6s.16xlarge.8	64	512	34/30	1,000	500	32	8	KVM

Table 4-17 M6 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	EVS Basic Bandwidth/Burst Bandwidth (Gbit/s)	Virtualization
m6.large.8	2	16	4/1.2	40	50	2	2	1/5	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	EVS Basic Bandwidth/Burst Bandwidth (Gbit/s)	Virtualization
m6.xlarge.8	4	32	8/2.4	80	50	2	3	1.5/5	KVM
m6.2xlarge.8	8	64	15/4.5	150	100	4	4	2/5	KVM
m6.3xlarge.8	12	96	17/7	200	150	4	6	2.5/5	KVM
m6.4xlarge.8	16	128	20/9	280	150	8	8	3.5/5	KVM
m6.6xlarge.8	24	192	25/14	400	200	8	8	4/5	KVM
m6.8xlarge.8	32	256	30/18	550	300	16	8	7/10	KVM
m6.16xlarge.8	64	512	40/36	1,000	500	32	8	20/None	KVM

Large-Memory ECSs

For more details, see [Large-Memory ECSs](#).

Table 4-18 Large-memory ECS features

Series	Compute	Disk Type	Network
E7	<ul style="list-style-type: none"> vCPU to memory ratio: 1:20 Number of vCPUs: 48 to 96 Next-generation Intel® Xeon® Scalable Processor Basic/Turbo frequency: 2.6 GHz/3.5 GHz 	<ul style="list-style-type: none"> High I/O General Purpose SSD Ultra-high I/O Extreme SSD General Purpose SSD V2 	<ul style="list-style-type: none"> Ultra-high packets per second (PPS) throughput An ECS with higher specifications has better network performance. Maximum PPS: 10,000,000 Maximum intranet bandwidth: 44 Gbit/s Maximum number of network connections: 10 million

Table 4-19 E7 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Max. Supplementary NICs	Virtualization
e7.12xlarge.20	48	960	30/20	550	500	16	8	10 24	KVM
e7.24xlarge.20	96	1920	44/40	1,000	1,000	32	8	20 48	KVM

Disk-intensive ECSs

For more details, see [Disk-intensive ECSs](#).

Table 4-20 Disk-intensive ECS features

Series	Compute	Disk Type	Network
D6	<ul style="list-style-type: none"> vCPU to memory ratio: 1:4 Number of vCPUs: 4 to 72 Intel® Xeon® Scalable Processor Basic/Turbo frequency: 2.6 GHz/3.5 GHz 	<ul style="list-style-type: none"> High I/O General Purpose SSD Ultra-high I/O Extreme SSD General Purpose SSD V2 	<ul style="list-style-type: none"> Ultra-high PPS throughput An ECS with higher specifications has better network performance. Maximum PPS: 9,000,000 Maximum intranet bandwidth: 44 Gbit/s Maximum number of network connections: 7 million

Table 4-21 D6 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Local Disks (GiB)	Virtualization
d6.xlarge.4	4	16	5/2	60	50	2	3	2 × 3,600	KVM
d6.2xlarge.4	8	32	10/4	120	100	4	4	4 × 3,600	KVM
d6.4xlarge.4	16	64	20/7.5	240	150	8	8	8 × 3,600	KVM
d6.6xlarge.4	24	96	25/11	350	200	8	8	12 × 3,600	KVM
d6.8xlarge.4	32	128	30/15	450	300	16	8	16 × 3,600	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Local Disks (GiB)	Virtualization
d6.1 2xlarge.4	48	192	40/22	650	400	16	8	24 × 3,600	KVM
d6.1 6xlarge.4	64	256	42/30	850	500	32	8	32 × 3,600	KVM

GPU-accelerated ECSs

For more details, see [GPU-accelerated ECSs](#).

Table 4-22 GPU-accelerated ECSs

Type	Series	GPU	CUDA Cores per GPU	Single-GPU Performance	Application
Computing-accelerated	P2s	NVIDIA V100	5,120	<ul style="list-style-type: none"> 14 TFLOPS of single-precision floating-point computing 7 TFLOPS of double-precision floating-point computing 112 TFLOPS Tensor Cores for deep learning acceleration 	AI deep learning training, scientific computing, computational fluid dynamics, computational finance, seismic analysis, molecular modeling, and genomics.

Type	Series	GPU	CUDA Cores per GPU	Single-GPU Performance	Application
Inference-accelerated	Pi2	NVIDIA T4 (GPU passthrough)	2,560	<ul style="list-style-type: none"> 8.1 TFLOPS of single-precision floating-point computing 130 INT8 TOPS 260 INT4 TOPS 	Machine learning, deep learning, inference training, scientific computing, seismic analysis, computing finance, rendering, multimedia encoding and decoding

Table 4-23 P2s ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	GPUs	GPU Connection	GPU Memory (GiB)	Virtualization
p2s.2xlarge.8	8	64	10/4	50	4	4	1 × V100	PCIe Gen3	1 × 32 GiB	KVM
p2s.4xlarge.8	16	128	15/8	100	8	8	2 × V100	PCIe Gen3	2 × 32 GiB	KVM
p2s.8xlarge.8	32	256	25/15	200	16	8	4 × V100	PCIe Gen3	4 × 32 GiB	KVM
p2s.16xlarge.8	64	512	30/30	400	32	8	8 × V100	PCIe Gen3	8 × 32 GiB	KVM

Table 4-24 Pi2 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	GPUs	GPU Memory (GiB)	Local Disks	Virtualization
pi2.2xlarge.4	8	32	10/4	50	4	4	1 × T4	1 × 16	N/A	KVM
pi2.4xlarge.4	16	64	15/8	100	8	8	2 × T4	2 × 16	N/A	KVM
pi2.8xlarge.4	32	128	25/15	200	16	8	4 × T4	4 × 16	N/A	KVM
pi2.16xlarge.4	64	256	30/30	400	32	8	8 × T4	8 × 16	N/A	KVM

AI-accelerated ECSs

Table 4-25 Ai2 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth	Max. PPS (10,000)	Ascend 310 series	Ascend RAM (GiB)	Max. NIC Queues	Max. NICs	Virtualization
ai2.4xlarge.4	16	64	20/6	280	1	24	8	8	KVM
ai2.8xlarge.4	32	128	30/12	550	2	48	8	8	KVM

Flavor	vCPUs	Memory (GiB)	Max. / Assured Bandwidth	Max. PPS (10,000)	Ascend 310 series	Ascend RAM (GiB)	Max. NIC Queues	Max. NICs	Virtualization
ai2.16xlarge.e4	64	256	36/24	800	4	96	16	8	KVM
ai2.24xlarge.e4	96	384	40/36	850	6	144	32	8	KVM

4.5.2 General Computing ECSs

Overview

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

Available now: S6

Table 4-26 General computing ECS features

Series	Compute	Disk Type	Network
S6	<ul style="list-style-type: none"> vCPU to memory ratio: 1:1, 1:2, or 1:4 Number of vCPUs: 1 to 8 2nd Generation Intel® Xeon® Scalable Processor Basic/Turbo frequency: 2.6 GHz/3.5 GHz 	<ul style="list-style-type: none"> High I/O General Purpose SSD Ultra-high I/O Extreme SSD General Purpose SSD V2 	<ul style="list-style-type: none"> An ECS with higher specifications has better network performance. Maximum PPS: 500,000 Maximum intranet bandwidth: 3 Gbit/s Maximum number of network connections: 300,000

General Computing S6

Overview

General computing S6 ECSs use the 2nd generation Intel® Xeon® Scalable processors and self-developed 25GE high-speed intelligent NICs to provide high network bandwidth and packets per second (PPS). For details, see [Table 4-26](#).

Notes

S6 ECS burstable performance is neither restricted by CPU credits nor billed for additional credits.

Scenarios

- Websites and web applications that have high requirements on PPS throughput
- Lightweight databases and cache servers
- Light- and medium-load enterprise applications

Specifications

Table 4-27 S6 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Virtualization
s6.small.1	1	1	0.8/0.1	10	30	1	2	KVM
s6.medium.2	1	2	0.8/0.1	10	30	1	2	KVM
s6.large.2	2	4	1.5/0.2	15	30	1	2	KVM
s6.xlarge.2	4	8	2/0.35	25	30	1	2	KVM
s6.2xlarge.2	8	16	3/0.75	50	30	2	2	KVM
s6.medium.4	1	4	0.8/0.1	10	30	1	2	KVM
s6.large.4	2	8	1.5/0.2	15	30	1	2	KVM
s6.xlarge.4	4	16	2/0.35	25	30	1	2	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Virtualization
s6.2xlarge.4	8	32	3/0.75	50	30	2	2	KVM

4.5.3 General Computing-plus ECSs

Overview

General computing-plus ECSs use dedicated vCPUs to deliver powerful performance. In addition, the ECSs use the latest-generation network acceleration engines and DPDK to provide high network performance.

Available now: C7n, C7t, C6s, and C6

Table 4-28 General computing-plus ECS features

Series	Compute	Disk Type	Network
C7n	<ul style="list-style-type: none"> vCPU to memory ratio: 1:2 or 1:4 Number of vCPUs: 2 to 96 3rd Generation Intel® Xeon® Scalable Processor Basic/Turbo frequency: 2.6 GHz/3.4 GHz 	<ul style="list-style-type: none"> High I/O General Purpose SSD Ultra-high I/O Extreme SSD General Purpose SSD V2 	<ul style="list-style-type: none"> Ultra-high PPS throughput An ECS with higher specifications has better network performance. Maximum PPS: 8,500,000 Maximum intranet bandwidth: 40 Gbit/s Maximum number of network connections: 8 million

Series	Compute	Disk Type	Network
C7t	<ul style="list-style-type: none"> vCPU to memory ratio: 1:4 Number of vCPUs: 2 to 128 3rd Generation Intel® Xeon® Scalable Processor Basic/Turbo frequency: 2.6 GHz/3.4 GHz and 2.8 GHz/3.5 GHz 	<ul style="list-style-type: none"> General Purpose SSD Ultra-high I/O Extreme SSD General Purpose SSD V2 Extreme SSD V2 	<ul style="list-style-type: none"> Ultra-high PPS throughput An ECS with higher specifications has better network performance. Maximum PPS: 30 million PPS Maximum intranet bandwidth: 90 Gbit/s Maximum number of network connections: 10 million
C6s	<ul style="list-style-type: none"> vCPU to memory ratio: 1:2 or 1:4 Number of vCPUs: 2 to 64 2nd Generation Intel® Xeon® Scalable Processor Basic/Turbo frequency: 2.6 GHz/3.5 GHz 	<ul style="list-style-type: none"> High I/O General Purpose SSD Ultra-high I/O Extreme SSD General Purpose SSD V2 	<ul style="list-style-type: none"> Ultra-high PPS throughput An ECS with higher specifications has better network performance. Maximum PPS: 8,500,000 Maximum intranet bandwidth: 30 Gbit/s Maximum number of network connections: 5 million
C6	<ul style="list-style-type: none"> vCPU to memory ratio: 1:2 or 1:4 Number of vCPUs: 2 to 64 2nd Generation Intel® Xeon® Scalable Processor Basic/Turbo frequency: 3.0 GHz/3.4 GHz 	<ul style="list-style-type: none"> High I/O General Purpose SSD Ultra-high I/O Extreme SSD General Purpose SSD V2 	<ul style="list-style-type: none"> Ultra-high PPS throughput An ECS with higher specifications has better network performance. Maximum PPS: 10,000,000 Maximum intranet bandwidth: 40 Gbit/s Maximum number of network connections: 10 million

C7t

Overview

C7t ECSs use the third-generation Intel® Xeon® scalable processor and the latest QingTian architecture to meet trusted computing and Internet connection requirements.

Scenarios

- Finance and government
- Internet communication that requires high PPS, such as on-screen comments and real-time voice

Specifications

Table 4-29 C7t ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Max. Supplementary NICs	EV S Basic Bandwidth/Burst Bandwidth (Gbit/s)	Virtualization
c7t.large.4	2	8	10/1.28	80	50	2	2	16	1.5/6	QingTian
c7t.xlarge.4	4	16	16/2.56	150	50	2	3	32	2/6	
c7t.2xlarge.4	8	32	20/4	200	100	4	4	64	3/6	
c7t.3xlarge.4	12	48	34/6.4	300	150	4	6	96	4/6	
c7t.4xlarge.4	16	64	40/8	400	150	8	8	128	5/6	

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Max. Supplementary NICs	EV S Basic Bandwidth/Burst Bandwidth (Gbit/s)	Virtualization
c7t.6xlarge.4	24	96	50/14.4	600	200	8	8	192	6/None	
c7t.8xlarge.4	32	128	60/16	800	300	16	8	256	8/None	
c7t.12xlarge.4	48	192	70/28.8	1,200	400	16	8	256	10/None	
c7t.16xlarge.4	64	256	72/32	1500	500	28	8	256	16/None	
c7t.24xlarge.4	96	384	80/40	2400	800	32	8	256	20/None	

C7n

Overview

C7n ECSs use the third-generation Intel® Xeon® Scalable processors to provide enhanced computing, security, and stability. A C7n ECS can have a maximum number of 96 vCPUs and a memory speed of 3,200 MHz. C7n ECSs support secure boot, providing a secure and trusted cloud environment.

Scenarios

Medium- and heavy-load enterprise applications with strict requirements on computing and network performance, such as web applications, e-commerce platforms, short video platforms, online games, and insurance and finance.

Specifications

Table 4-30 C7n ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Max. Supplementary NICs	EV S Basic Bandwidth/Burst Bandwidth (Gbit/s)	Virtualization
c7n.large.2	2	4	4/0.8	40	50	2	2	16	1.5/6	KVM
c7n.xlarge.2	4	8	8/1.6	80	50	2	3	32	2/6	KVM
c7n.2xlarge.2	8	16	15/3	150	100	4	4	64	3/6	KVM
c7n.3xlarge.2	12	24	17/5	200	150	4	6	96	4/6	KVM
c7n.4xlarge.2	16	32	20/6	280	150	8	8	128	5/6	KVM
c7n.6xlarge.2	24	48	25/9	400	200	8	8	192	6/None	KVM
c7n.8xlarge.2	32	64	30/12	550	300	16	8	256	8/None	KVM
c7n.12xlarge.2	48	96	35/18	750	400	16	8	256	12/None	KVM
c7n.16xlarge.2	64	128	36/24	800	500	28	8	256	16/None	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Max. SupplementarY NICs	EV S Basic Bandwidth/Burst Bandwidth (Gbit/s)	Virtualization
c7n.24xlarge.2	96	192	40/36	850	800	32	8	256	24/None	KVM
c7n.large.4	2	8	4/0.8	40	50	2	2	16	1.5/6	KVM
c7n.xlarge.4	4	16	8/1.6	80	50	2	3	32	2/6	KVM
c7n.2xlarge.4	8	32	15/3	150	100	4	4	64	3/6	KVM
c7n.3xlarge.4	12	48	17/5	200	150	4	6	96	4/6	KVM
c7n.4xlarge.4	16	64	20/6	280	150	8	8	128	5/6	KVM
c7n.6xlarge.4	24	96	25/9	400	200	8	8	192	6/None	KVM
c7n.8xlarge.4	32	128	30/12	550	300	16	8	256	8/None	KVM
c7n.12xlarge.4	48	192	35/18	750	400	16	8	256	12/None	KVM
c7n.16xlarge.4	64	256	36/24	800	500	28	8	256	16/None	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Max. Supplementarily NICs	EV S Basic Bandwidth/Burst Bandwidth (Gbit/s)	Virtualization
c7n.24xlarge.4	96	384	40/36	850	800	32	8	256	24/None	KVM

C6s

Overview

C6s ECSs use the second-generation Intel® Xeon® Scalable processors that feature high performance, stability, low latency, and cost-effectiveness. They are suitable for Internet, gaming, and rendering scenarios, especially those that require high computing and network stability.

Scenarios

Internet, gaming, and rendering scenarios, especially those with strict requirements on computing and network stability

- Gaming: C6s ECSs meet requirements of high performance and stability.
- Rendering: C6s ECSs provide high-quality rendering at the most optimal cost.
- Other scenarios: C6s ECSs can be used in gaming acceleration, video bullet screen, website building, and app development.

Specifications

Table 4-31 C6s ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Virtualization
c6s.large.2	2	4	1/1	30	50	2	2	KVM
c6s.xlarge.2	4	8	2/2	60	50	2	3	KVM
c6s.2xlarge.2	8	16	4/4	120	100	4	4	KVM
c6s.3xlarge.2	12	24	5.5/5.5	180	150	4	6	KVM
c6s.4xlarge.2	16	32	7.5/7.5	240	150	8	8	KVM
c6s.6xlarge.2	24	48	11/11	350	200	8	8	KVM
c6s.8xlarge.2	32	64	15/15	450	300	16	8	KVM
c6s.12xlarge.2	48	96	22/22	650	400	16	8	KVM
c6s.16xlarge.2	64	128	30/30	850	500	32	8	KVM
c6s.large.4	2	8	1/1	30	50	2	2	KVM
c6s.xlarge.4	4	16	2/2	60	50	2	3	KVM
c6s.2xlarge.4	8	32	4/4	120	100	4	4	KVM
c6s.3xlarge.4	12	48	5.5/5.5	180	150	4	6	KVM
c6s.4xlarge.4	16	64	7.5/7.5	240	150	8	8	KVM
c6s.6xlarge.4	24	96	11/11	350	200	8	8	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Virtualization
c6s.8xlarge.4	32	128	15/15	450	300	16	8	KVM
c6s.12xlarge.4	48	192	22/22	650	400	16	8	KVM
c6s.16xlarge.4	64	256	30/30	850	500	32	8	KVM

C6

Overview

C6 ECSs use second-generation Intel® Xeon® Scalable processors to provide powerful and stable computing performance. By using 25GE high-speed intelligent NICs, C6 ECSs offer ultra-high network bandwidth and PPS.

Scenarios

- Websites and web applications that require high computing and network performance
- General databases and cache servers
- Medium- and heavy-load enterprise applications
- Gaming and rendering

Specifications

Table 4-32 C6 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	EVS Basic Bandwidth/Burst Bandwidth (Gbit/s)	Virtualization
c6.large.2	2	4	4/1.2	40	50	2	2	1/5	KVM
c6.xlarge.2	4	8	8/2.4	80	50	2	3	1.5/5	KVM
c6.2xlarge.2	8	16	15/4.5	150	100	4	4	2/5	KVM
c6.3xlarge.2	12	24	17/7	200	150	4	6	2.5/5	KVM
c6.4xlarge.2	16	32	20/9	280	150	8	8	3.5/5	KVM
c6.6xlarge.2	24	48	25/14	400	200	8	8	4/5	KVM
c6.8xlarge.2	32	64	30/18	550	300	16	8	7/10	KVM
c6.12xlarge.2	48	96	35/27	750	400	16	8	10/15	KVM
c6.16xlarge.2	64	128	40/36	1,000	500	32	8	20/None	KVM
c6.22xlarge.2	88	176	44/40	1,200	1,000	32	8	20/None	KVM
c6.large.4	2	8	4/1.2	40	50	2	2	1/5	KVM
c6.xlarge.4	4	16	8/2.4	80	50	2	3	1.5/5	KVM
c6.2xlarge.4	8	32	15/4.5	150	100	4	4	2/5	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	EVS Basic Bandwidth/Burst Bandwidth (Gbit/s)	Virtualization
c6.3xlarge.4	12	48	17/7	200	150	4	6	2.5/5	KVM
c6.4xlarge.4	16	64	20/9	280	150	8	8	3.5/5	KVM
c6.6xlarge.4	24	96	25/14	400	200	8	8	4/5	KVM
c6.8xlarge.4	32	128	30/18	550	300	16	8	7/10	KVM
c6.12xlarge.4	48	192	35/27	750	400	16	8	10/15	KVM
c6.16xlarge.4	64	256	40/36	1,000	500	32	8	20/None	KVM
c6.22xlarge.4	88	352	44/40	1,200	1,000	28	8	20/None	KVM
c6.22xlarge.4.physical	88	384	44/40	1,000	1,000	16	33	20/None	Bare metal

4.5.4 General Computing-Basic ECSs

General Computing-Basic T6

Overview

General computing-basic ECSs are suitable for scenarios that require moderate CPU performance generally but burstable high performance occasionally while keeping costs low. The performance of such ECSs is constrained by the baseline performance and CPU credits.

 NOTE

- General computing-basic ECSs are available only in LA-Mexico City1 and LA-Mexico City2 regions.
- Before using general computing-basic ECSs, learn about the concepts related to CPU credits by referring to [CPU Credits of T6 ECSs](#).
- CPU credits do not incur additional costs. For more information about CPU credits, see [CPU Credits](#).

Table 4-33 General computing-basic ECS features

Series	Compute	Disk Type	Network
T6	<ul style="list-style-type: none"> • vCPU to memory ratio: 1:1, 1:2, or 1:4 • Number of vCPUs: 1 to 16 • Intel® Xeon® Scalable Processor • Basic/Turbo frequency: 2.2 GHz/3.0 GHz 	<ul style="list-style-type: none"> • High I/O • General Purpose SSD • Ultra-high I/O • Extreme SSD • General Purpose SSD V2 	<ul style="list-style-type: none"> • An ECS with higher specifications has better network performance. • Maximum packets per second (PPS): 600,000 • Maximum intranet bandwidth: 3 Gbit/s • Maximum number of network connections: 500,000

Notes

General computing-basic T6 ECSs are suitable for scenarios that require moderate CPU performance generally but burstable high performance occasionally. The performance of such ECSs is constrained by the baseline performance and CPU credits. After a T6 ECS is created, you can view its CPU credits by choosing **More > Manage Credits** in the **Operation** column.

Scenarios

General computing-basic ECSs are suitable for applications that require moderate CPU performance generally but burstable high performance occasionally, such as web application servers, light-load applications, and microservices.

Specifications

Table 4-34 T6 ECS specifications

Flavor	vCPUs	Memory (GiB)	CPU Baseline (%)	Average CPU Baseline (%)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Virtualization
t6.small.1	1	1	10	10	0.3/0.05	6	30	1	1	KVM
t6.large.1	2	2	40	20	0.5/0.1	10	30	1	1	KVM
t6.xlarge.1	4	4	80	20	1/0.2	20	30	2	2	KVM
t6.2xlarge.1	8	8	120	15	2/0.4	40	30	4	2	KVM
t6.4xlarge.1	16	16	240	15	3/0.8	60	50	8	2	KVM
t6.medium.2	1	2	10	10	0.3/0.05	6	30	1	1	KVM
t6.large.2	2	4	40	20	0.5/0.1	10	30	1	1	KVM
t6.xlarge.2	4	8	80	20	1/0.2	20	30	2	2	KVM
t6.2xlarge.2	8	16	120	15	2/0.4	40	30	4	2	KVM

Flavor	vCPUs	Memory (GiB)	CPU Baseline (%)	Average CPU Baseline (%)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Virtualization
t6.xlarge.2	16	32	240	15	3/0.8	60	50	8	2	KVM
t6.large.4	2	8	40	20	0.5/0.1	10	30	1	1	KVM
t6.xlarge.4	4	16	80	20	1/0.2	20	30	2	2	KVM
t6.xlarge.4	8	32	120	15	2/0.4	40	30	4	2	KVM

CPU Credits of T6 ECSs

- **Initial CPU credits:** After a T6 ECS is created, it automatically obtains initial CPU credits. Initial CPU credits are allocated only after an ECS is created.
For example, after a t6.large.1 ECS is created, it obtains 60 initial CPU credits.
- **Maximum CPU credits:** If the CPU usage is below the baseline, credits earned (accrued credits) are greater than credits spent. The accrued credits will not expire on a running ECS. When the credits reach the maximum value allowed (depending on the ECS flavor), no more credits will be earned.
For example, the maximum number of CPU credits for a t6.large.1 ECS is 576. When the credits reach 576, no more credits will be earned. When the credits are below 576, credits can be earned again.
- **CPU credits earned per hour:** The number of CPU credits earned by an ECS per hour reflects the CPU baseline. One CPU credit is equal to one vCPU running at 100% usage for one minute.
For example, a t6.large.1 ECS can earn 24 CPU credits per hour.

For more information about CPU usage calculations, see [CPU Credits](#).

Table 4-35 CPU credits of T6 ECSs

Flavor	Initial CPU Credits	Maximum CPU Credits	CPU Credits Earned per Hour
t6.small.1	30	144	6
t6.large.1	60	576	24
t6.xlarge.1	120	1,152	48
t6.2xlarge.1	120	1,728	72
t6.4xlarge.1	160	3,456	144
t6.medium.2	30	144	6
t6.large.2	60	576	24
t6.xlarge.2	120	1,152	48
t6.2xlarge.2	120	1,728	72
t6.4xlarge.2	160	3,456	144
t6.large.4	60	576	24
t6.xlarge.4	120	1,152	48
t6.2xlarge.4	120	1,728	72

4.5.5 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 marketing, e-commerce, and IoT big data analysis.

Available now: M7n, M6, and M6s

Table 4-36 Memory-optimized ECS features

Series	Compute	Disk Type	Network
M7n	<ul style="list-style-type: none"> vCPU to memory ratio: 1:8 Number of vCPUs: 2 to 96 3rd Generation Intel® Xeon® Scalable Processor Basic/Turbo frequency: 2.6 GHz/3.4 GHz 	<ul style="list-style-type: none"> High I/O General Purpose SSD Ultra-high I/O Extreme SSD General Purpose SSD V2 	<ul style="list-style-type: none"> Support for IPv6 Ultra-high PPS throughput An ECS with higher specifications has better network performance. Maximum PPS: 8,500,000 Maximum intranet bandwidth: 40 Gbit/s Maximum number of network connections: 8 million
M6s	<ul style="list-style-type: none"> vCPU to memory ratio: 1:8 Number of vCPUs: 2 to 64 2nd Generation Intel® Xeon® Scalable Processor Basic/Turbo frequency: 3.0 GHz/3.4 GHz 	<ul style="list-style-type: none"> High I/O General Purpose SSD Ultra-high I/O Extreme SSD General Purpose SSD V2 	<ul style="list-style-type: none"> Ultra-high PPS throughput An ECS with higher specifications has better network performance. Maximum PPS: 10 million PPS Maximum intranet bandwidth: 34 Gbit/s

Series	Compute	Disk Type	Network
M6	<ul style="list-style-type: none"> vCPU to memory ratio: 1:8 Number of vCPUs: 2 to 64 2nd Generation Intel® Xeon® Scalable Processor Basic/Turbo frequency: 3.0 GHz/3.4 GHz 	<ul style="list-style-type: none"> High I/O General Purpose SSD Ultra-high I/O Extreme SSD General Purpose SSD V2 	<ul style="list-style-type: none"> Ultra-high PPS throughput An ECS with higher specifications has better network performance. Maximum PPS: 10,000,000 Maximum intranet bandwidth: 40 Gbit/s Maximum number of network connections: 10 million

Memory-optimized M7n

M7n ECSs use the third-generation Intel® Xeon® Scalable processors to provide enhanced computing, security, and stability. Each M7n ECS can have a maximum number of 96 vCPUs and a memory speed of 3,200 MHz, and provide a secure and trusted cloud environment for memory-intensive computing applications.

Scenarios

- Massively parallel processing (MPP) of data warehouse
- MapReduce and Hadoop distributed computing
- Distributed file systems
- Network file system, log, or data processing applications

Specifications

Table 4-37 M7n ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	EVS Basic Bandwidth/Burst Bandwidth (Gbit/s)	Virtualization
m7n.large.8	2	16	4/0.8	40	50	2	2	1.5/6	KVM
m7n.xlarge.8	4	32	8/1.6	80	50	2	3	2/6	KVM
m7n.2xlarge.8	8	64	15/3	150	100	4	4	3/6	KVM
m7n.3xlarge.8	12	96	17/5	200	150	4	6	4/6	KVM
m7n.4xlarge.8	16	128	20/6	280	150	8	8	5/6	KVM
m7n.6xlarge.8	24	192	25/9	400	200	8	8	6/None	KVM
m7n.8xlarge.8	32	256	30/12	550	300	16	8	8/None	KVM
m7n.12xlarge.8	48	384	35/18	750	400	16	8	12/None	KVM
m7n.16xlarge.8	64	512	36/24	800	500	28	8	16/None	KVM
m7n.24xlarge.8	96	768	40/36	850	800	32	8	24/None	KVM

Memory-optimized M6s

Overview

M6s ECSs use the second-generation Intel® Xeon® Scalable processors with technologies optimized to offer powerful and stable computing performance. Using 25GE high-speed intelligent NICs, M6s ECSs provide a maximum memory size of 512 GiB based on DDR4 for memory-intensive applications with high requirements on network bandwidth and Packets Per Second (PPS).

Scenarios

- Massively parallel processing (MPP) of data warehouse
- MapReduce and Hadoop distributed computing
- Distributed file systems
- Network file system, log, or data processing applications

Specifications

Table 4-38 M6s ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Virtualization
m6s.large.8	2	16	3/1	30	50	2	2	KVM
m6s.xlarge.8	4	32	6/2	60	50	2	3	KVM
m6s.2xlarge.8	8	64	12/4	120	100	4	4	KVM
m6s.3xlarge.8	12	96	14/5.5	160	150	4	6	KVM
m6s.4xlarge.8	16	128	16/7.5	220	150	8	8	KVM
m6s.6xlarge.8	24	192	20/11	320	200	8	8	KVM
m6s.8xlarge.8	32	256	25/15	450	300	16	8	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Virtualization
m6s.16xlarge.8	64	512	34/30	1,000	500	32	8	KVM

Memory-optimized M6

Overview

M6 ECSs use the second-generation Intel® Xeon® Scalable processors with technologies optimized to offer powerful and stable computing performance. Using 25GE high-speed intelligent NICs, M6 ECSs provide a maximum memory size of 512 GiB based on DDR4 for memory-intensive applications with high requirements on network bandwidth and Packets Per Second (PPS).

Scenarios

- Massively parallel processing (MPP) database
- MapReduce and Hadoop distributed computing
- Distributed file systems
- Network file system, log, or data processing applications

Specifications

Table 4-39 M6 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	EVS Basic Bandwidth/Burst Bandwidth (Gbit/s)	Virtualization
m6.large.8	2	16	4/1.2	40	50	2	2	1/5	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	EVS Basic Bandwidth/Burst Bandwidth (Gbit/s)	Virtualization
m6.xlarge.8	4	32	8/2.4	80	50	2	3	1.5/5	KVM
m6.2xlarge.8	8	64	15/4.5	150	100	4	4	2/5	KVM
m6.3xlarge.8	12	96	17/7	200	150	4	6	2.5/5	KVM
m6.4xlarge.8	16	128	20/9	280	150	8	8	3.5/5	KVM
m6.6xlarge.8	24	192	25/14	400	200	8	8	4/5	KVM
m6.8xlarge.8	32	256	30/18	550	300	16	8	7/10	KVM
m6.16xlarge.8	64	512	40/36	1,000	500	32	8	20/None	KVM

4.5.6 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 capability of processing large volumes of data. Large-memory ECSs provide large memory and high computing, storage, and network performance.

Available now: E7

Table 4-40 Large-memory ECS features

Series	Compute	Disk Type	Network
E7	<ul style="list-style-type: none">vCPU to memory ratio: 1:20Number of vCPUs: 48 to 96Next-generation Intel® Xeon® Scalable ProcessorBasic/Turbo frequency: 2.6 GHz/3.5 GHz	<ul style="list-style-type: none">High I/OGeneral Purpose SSDUltra-high I/OExtreme SSDGeneral Purpose SSD V2	<ul style="list-style-type: none">Ultra-high packets per second (PPS) throughputAn ECS with higher specifications has better network performance.Maximum PPS: 10,000,000Maximum intranet bandwidth: 44 Gbit/sMaximum number of network connections: 10 million

Large-Memory E7

Overview

Large-memory E7 ECSs use next-generation Intel® Xeon® Scalable processors to provide powerful and stable compute. By using high-speed intelligent NICs, E7 ECSs offer ultra-high network bandwidth and packets per second (PPS).

Notes

For notes about using E7 ECSs, see [Notes](#).

Scenarios

- Online transaction processing (OLTP) and online analytical processing (OLAP) scenarios
- SAP HANA in-memory databases, such as SAP Business Suite on HANA (SoH), SAP S/4HANA (S4H), SAP Business Warehouse on HANA (BWoH), and SAP BW/4HANA (B4H)
- High-performance databases
- Distributed cache
- Big data processing engines and data mining applications

Specifications

Table 4-41 E7 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Max. Supplementary NICs	Virtualization
e7.12xlarge.20	48	960	30/20	550	500	16	8	1024	KVM
e7.24xlarge.20	96	1920	44/40	1,000	1,000	32	8	2048	KVM

Notes

- Large-memory ECSs do not support NIC hot swapping.
- Affected by the memory loading speed, large-memory ECSs may take longer to start.
- The primary and extension NICs of a large-memory ECS can only be used in the scenarios listed in [Table 4-42](#).

Table 4-42 Application scenarios of the NICs of a large-memory ECS

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

- An ECS can have a maximum of 60 attached disks, including the system disk. For details about constraints, see [Can I Attach Multiple Disks to an ECS?](#) For example, an E7 ECS can have one system disk and 59 data disks.

4.5.7 Disk-intensive ECSs

Overview

Disk-intensive ECSs are delivered with local disks for high storage bandwidth and IOPS. In addition, local disks are more cost-effective in massive data storage scenarios. Disk-intensive ECSs have the following features:

- They use local disks to provide high sequential read/write performance and low latency, improving file read/write performance.

- They provide powerful and stable computing capabilities, ensuring efficient data processing.
- They provide high intranet performance, including robust intranet bandwidth and packets per second (PPS), for data exchange between ECSs during peak hours.

Available flavors

Available now: D6

Table 4-43 Disk-intensive ECS features

Series	Compute	Disk Type	Network
D6	<ul style="list-style-type: none">• vCPU to memory ratio: 1:4• Number of vCPUs: 4 to 72• Intel® Xeon® Scalable Processor• Basic/Turbo frequency: 2.6 GHz/3.5 GHz	<ul style="list-style-type: none">• High I/O• General Purpose SSD• Ultra-high I/O• Extreme SSD• General Purpose SSD V2	<ul style="list-style-type: none">• Ultra-high PPS throughput• An ECS with higher specifications has better network performance.• Maximum PPS: 9,000,000• Maximum intranet bandwidth: 44 Gbit/s• Maximum number of network connections: 7 million

Disk-intensive D6

Overview

D6 ECSs, with a vCPU/memory ratio of 1:4, use 2nd Generation Intel® Xeon® Scalable processors to offer powerful and stable computing performance. Equipped with 25GE high-speed intelligent NICs and local SATA disks, D6 ECSs offer ultra-high network bandwidth, PPS, and local storage. The capacity of a single SATA disk is up to 3,600 GiB, and an ECS can have up to 36 such disks attached.

Notes

For details, see [Notes on Using D6 ECSs](#).

Scenarios

Disk-intensive D6 ECSs are suitable for applications that need to process large volumes of data and require high I/O performance and rapid data switching and processing, including massively parallel processing (MPP) databases, MapReduce and Hadoop distributed computing, and big data computing, distributed file systems, network file systems, and logs and data processing applications.

Specifications

Table 4-44 D6 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Network Connections (10,000)	Max. NIC Queues	Max. NICs	Local Disks (GiB)	Virtualization
d6.xlarge.4	4	16	5/2	60	50	2	3	2 × 3,600	KVM
d6.2xlarge.4	8	32	10/4	120	100	4	4	4 × 3,600	KVM
d6.4xlarge.4	16	64	20/7.5	240	150	8	8	8 × 3,600	KVM
d6.6xlarge.4	24	96	25/11	350	200	8	8	12 × 3,600	KVM
d6.8xlarge.4	32	128	30/15	450	300	16	8	16 × 3,600	KVM
d6.12xlarge.4	48	192	40/22	650	400	16	8	24 × 3,600	KVM
d6.16xlarge.4	64	256	42/30	850	500	32	8	32 × 3,600	KVM

Performance of a Single SATA HDD Disk Attached to a D6 ECS

Table 4-45 Performance of a single SATA HDD disk attached to a D6 ECS

Metric	Performance
Disk capacity	3,600 GiB
Maximum throughput	198 Mbit/s

Metric	Performance
Access latency	Millisecond-level

Notes on Using D6 ECSs

- Currently, the following operating systems are supported (subject to the information displayed on the console):
 - CentOS 6.10 64bit
 - EulerOS 2.9 64bit
 - Ubuntu 20.04 64bit
- If the host where a D6 ECS is deployed is faulty, the ECS cannot be restored through live migration.
 - If the host is faulty or subhealthy and needs to be repaired, you need to stop the ECS.
 - In case of system maintenance or hardware faults, the ECS will be redeployed (to ensure HA) and cold migrated to another host. The local disk data of the ECS will not be retained.
- D6 ECSs do not support specifications modification.
- D6 ECSs do not support local disk snapshots or backups.
- D6 ECSs can use both local disks and EVS disks to store data. Restrictions on using the two types of storage media are as follows:
 - Only an EVS disk can be used as the system disk of a D6 ECS.
 - Both EVS disks and local disks can be used as data disks of a D6 ECS.
 - A maximum of 60 disks (including VBD, SCSI, and local disks) can be attached to a D6 ECS. Among the 60 disks, the maximum number of SCSI disks is 30, and the VBD disks (including the system disk) is 24. For details, see [Can I Attach Multiple Disks to an ECS?](#)

NOTE

The maximum number of disks attached to an existing D6 ECS remains unchanged.

- You can modify the **fstab** file to set automatic disk mounting at ECS start. For details, see [Configuring Automatic Mounting at System Start](#).
- The local disk data of a D6 ECS may be lost if an exception occurs, such as physical server breakdown or local disk damage. If your application does not use the data reliability architecture, it is a good practice to use EVS disks to build your ECS.
- When a D6 ECS is deleted, its local disk data will also be automatically deleted, which can take some time. As a result, a D6 ECS takes a longer time than other ECSs to be deleted. Back up the data before deleting such an ECS.
- Do not store service data in local disks for a long time. Instead, store it in EVS disks. To improve data security, use a high availability architecture and back up data in a timely manner.
- Local disks can only be purchased during ECS creation. They cannot be separately purchased after the ECS has been created. The quantity and

capacity of your local disks are determined according to the specifications of your ECS.

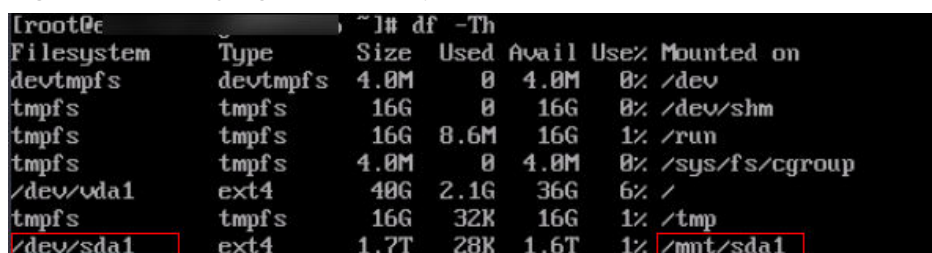
Handling Damaged Local Disks Attached to an ECS of D Series

If a local disk attached to an ECS is damaged, perform the following operations to handle this issue:

For a Linux ECS:

1. Detach the faulty local disk.
 - a. Run the following command to query the mount point of the faulty disk:
df -Th

Figure 4-1 Querying the mount point



```
[root@ec ~]# df -Th
Filesystem      Type      Size  Used Avail Use% Mounted on
devtmpfs        devtmpfs  4.0M   0    4.0M  0%  /dev
tmpfs           tmpfs     16G    0    16G   0%  /dev/shm
tmpfs           tmpfs     16G   8.6M  16G   1%  /run
tmpfs           tmpfs     4.0M   0    4.0M  0%  /sys/fs/cgroup
/dev/vda1       ext4      48G   2.1G  36G   6%  /
tmpfs           tmpfs     16G   32K   16G   1%  /tmp
/dev/sda1       ext4      1.7T   28K   1.6T  1%  /mnt/sda1
```

- b. Run the following command to detach the faulty local disk:
umount *Mount point*
In the example shown in [Figure 4-1](#), the mount point of `/dev/sda1` is `/mnt/sda1`. Run the following command:
umount /mnt/sda1
 2. Check whether the mount point of the faulty disk is configured in `/etc/fstab` of the ECS. If yes, comment out the mount point to prevent the ECS from entering the maintenance mode upon ECS startup after the faulty disk is replaced.
 - a. Run the following command to obtain the partition UUID:
blkid *Disk partition*
In this example, run the following command to obtain the UUID of the `/dev/sda1` partition:
blkid /dev/sda1
Information similar to the following is displayed:
`/dev/sda1: UUID="b9a07b7b-9322-4e05-ab9b-14b8050cd8cc" TYPE="ext4"`
 - b. Run the following command to check whether `/etc/fstab` contains the automatic mounting information about the disk partition:
cat /etc/fstab
Information similar to the following is displayed:
`UUID=b9a07b7b-9322-4e05-ab9b-14b8050cd8cc /mnt ext4 defaults 0 0`
 - c. If the mounting information exists, perform the following steps to delete it.
 - i. Run the following command to edit `/etc/fstab`:
vi /etc/fstab

Use the UUID obtained in 2.a to check whether the mounting information of the local disk is contained in `/etc/fstab`. If yes, comment out the information. This prevents the ECS from entering the maintenance mode upon ECS startup after the local disk is replaced.

- ii. Press **i** to enter editing mode.
- iii. Delete or comment out the automatic mounting information of the disk partition.

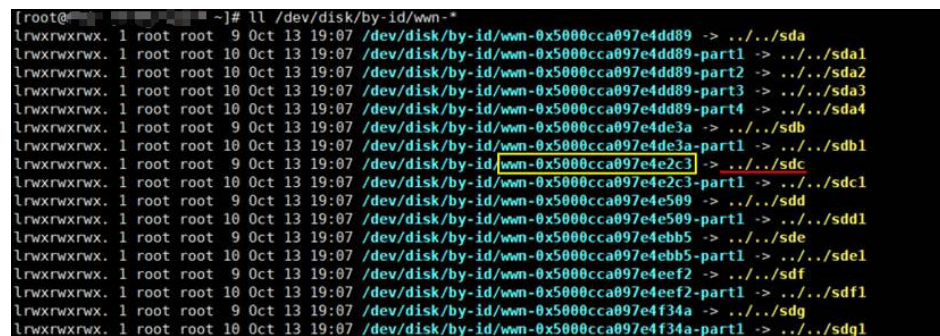
For example, add a pound sign (**#**) at the beginning of the following command line to comment out the automatic mounting information:

```
# UUID=b9a07b7b-9322-4e05-ab9b-14b8050cd8cc /mnt ext4 defaults 0 0
```

- iv. Press **Esc** to exit editing mode. Enter **:wq** and press **Enter** to save the settings and exit.
3. Run the following command to obtain the WWN of the local disk:
For example, if the sdc disk is faulty, obtain the WWN of the sdc disk.

```
ll /dev/disk/by-id/ | grep wwn-
```

Figure 4-2 Querying the WWN of the faulty local disk



```
[root@ ~]# ll /dev/disk/by-id/wwn-*
lrwxrwxrwx. 1 root root 9 Oct 13 19:07 /dev/disk/by-id/wwn-0x5000cca097e4dd89 -> ../../sda
lrwxrwxrwx. 1 root root 10 Oct 13 19:07 /dev/disk/by-id/wwn-0x5000cca097e4dd89-part1 -> ../../sda1
lrwxrwxrwx. 1 root root 10 Oct 13 19:07 /dev/disk/by-id/wwn-0x5000cca097e4dd89-part2 -> ../../sda2
lrwxrwxrwx. 1 root root 10 Oct 13 19:07 /dev/disk/by-id/wwn-0x5000cca097e4dd89-part3 -> ../../sda3
lrwxrwxrwx. 1 root root 10 Oct 13 19:07 /dev/disk/by-id/wwn-0x5000cca097e4dd89-part4 -> ../../sda4
lrwxrwxrwx. 1 root root 9 Oct 13 19:07 /dev/disk/by-id/wwn-0x5000cca097e4de3a -> ../../sdb
lrwxrwxrwx. 1 root root 10 Oct 13 19:07 /dev/disk/by-id/wwn-0x5000cca097e4de3a-part1 -> ../../sdb1
lrwxrwxrwx. 1 root root 9 Oct 13 19:07 /dev/disk/by-id/wwn-0x5000cca097e4e2c3 -> ../../sdc
lrwxrwxrwx. 1 root root 10 Oct 13 19:07 /dev/disk/by-id/wwn-0x5000cca097e4e2c3-part1 -> ../../sdc1
lrwxrwxrwx. 1 root root 9 Oct 13 19:07 /dev/disk/by-id/wwn-0x5000cca097e4e509 -> ../../sdd
lrwxrwxrwx. 1 root root 10 Oct 13 19:07 /dev/disk/by-id/wwn-0x5000cca097e4e509-part1 -> ../../sdd1
lrwxrwxrwx. 1 root root 9 Oct 13 19:07 /dev/disk/by-id/wwn-0x5000cca097e4ebb5 -> ../../sde
lrwxrwxrwx. 1 root root 10 Oct 13 19:07 /dev/disk/by-id/wwn-0x5000cca097e4ebb5-part1 -> ../../sde1
lrwxrwxrwx. 1 root root 9 Oct 13 19:07 /dev/disk/by-id/wwn-0x5000cca097e4eef2 -> ../../sdf
lrwxrwxrwx. 1 root root 10 Oct 13 19:07 /dev/disk/by-id/wwn-0x5000cca097e4eef2-part1 -> ../../sdf1
lrwxrwxrwx. 1 root root 9 Oct 13 19:07 /dev/disk/by-id/wwn-0x5000cca097e4f34a -> ../../sdg
lrwxrwxrwx. 1 root root 10 Oct 13 19:07 /dev/disk/by-id/wwn-0x5000cca097e4f34a-part1 -> ../../sdg1
```

4. Stop the ECS and provide the WWN of the faulty disk to technical support personnel to replace the local disk.

After the local disk is replaced, restart the ECS to synchronize the new local disk information to the virtualization layer.

4.5.8 GPU-accelerated ECSs

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

GPU-accelerated ECSs

- P series
 - [Computing-accelerated P2s](#) (recommended)
 - [Inference-accelerated Pi2](#) (recommended)

Helpful links:

- [Images Supported by GPU-accelerated ECSs](#)
- [Installing a GRID Driver on a GPU-accelerated ECS](#)

- [Installing a NVIDIA GPU Driver and CUDA Toolkit on a GPU-accelerated ECS](#)

Table 4-46 GPU-accelerated ECSs

Type	Series	GPU	CUDA Cores per GPU	Single-GPU Performance	Application
Computing-accelerated	P2s	NVIDIA V100	5,120	<ul style="list-style-type: none"> • 14 TFLOPS of single-precision floating-point computing • 7 TFLOPS of double-precision floating-point computing • 112 TFLOPS Tensor Cores for deep learning acceleration 	AI deep learning training, scientific computing, computational fluid dynamics, computational finance, seismic analysis, molecular modeling, and genomics.
Inference-accelerated	Pi2	NVIDIA T4 (GPU passthrough)	2,560	<ul style="list-style-type: none"> • 8.1 TFLOPS of single-precision floating-point computing • 130 INT8 TOPS • 260 INT4 TOPS 	Machine learning, deep learning, inference training, scientific computing, seismic analysis, computing finance, rendering, multimedia encoding and decoding

Images Supported by GPU-accelerated ECSs

Table 4-47 Images supported by GPU-accelerated ECSs

Type	Series	Supported Image
Computing-accelerated	P2s	<ul style="list-style-type: none"> CentOS 8.2 64bit CentOS 7.6 64bit Ubuntu 20.04 Server 64bit Ubuntu 18.04 Server 64bit
Inference-accelerated	Pi2	<ul style="list-style-type: none"> CentOS 8.2 64bit CentOS 7.6 64bit Ubuntu 20.04 Server 64bit Ubuntu 18.04 Server 64bit

Computing-accelerated P2s

Overview

P2s ECSs use NVIDIA Tesla V100 GPUs to provide flexibility, high-performance computing, and cost-effectiveness. P2s 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 4-48 P2s ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	GPUs	GPU Connection	GPU Memory (GiB)	Virtualization
p2s.2xlarge.8	8	64	10/4	50	4	4	1 × V100	PCIe Gen 3	1 × 32 GiB	KVM
p2s.4xlarge.8	16	128	15/8	100	8	8	2 × V100	PCIe Gen 3	2 × 32 GiB	KVM
p2s.8xlarge.8	32	256	25/15	200	16	8	4 × V100	PCIe Gen 3	4 × 32 GiB	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	GPUs	GPU Connection	GPU Memory (GiB)	Virtualization
p2s.16xlarge.8	64	512	30/30	400	32	8	8 × V100	PCIe Gen 3	8 × 32 GiB	KVM

P2s ECS Features

- CPU: 2nd Generation Intel® Xeon® Scalable 6278 processors (2.6 GHz of base frequency and 3.5 GHz of turbo frequency), or Intel® Xeon® Scalable 6151 processors (3.0 GHz of base frequency and 3.4 GHz of turbo frequency)
- 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
- 14 TFLOPS of single-precision computing and 7 TFLOPS of double-precision computing
- NVIDIA Tensor cores with 112 TFLOPS of single- and double-precision computing for deep learning
- Up to 30 Gbit/s of network bandwidth on a single ECS
- 32 GiB of HBM2 GPU memory with a bandwidth of 900 Gbit/s
- Comprehensive basic capabilities
 - User-defined network with flexible subnet division and network access policy configuration
 - Mass storage, elastic expansion, and backup and restoration
 - Elastic scaling
- Flexibility

Similar to other types of ECSs, P2s 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 P2s ECSs.

Supported Common Software

P2s ECSs are used in computing acceleration scenarios, such as deep learning training, inference, scientific computing, molecular modeling, and seismic analysis. If the software is required to support GPU CUDA, use P2s ECSs. P2s 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 3ds Max and V-Ray for 3ds Max

- Agisoft PhotoScan
- MapD

Notes

- After a P2s ECS is stopped, basic resources (including vCPUs, memory, image, and GPUs) are not billed, but its system disk is billed based on the disk capacity. If other products, such as EVS disks, EIP, and bandwidth are associated with the ECS, these products are billed separately.

NOTE

Resources will be released after a P2s ECS is stopped. If resources are insufficient at the next start, the start may fail. If you want to use such an ECS for a long period of time, do not stop the ECS or change its billing mode to yearly/monthly.

- By default, P2s ECSs created using a public image have the Tesla driver installed.
- If a P2s ECS is created using a private image, make sure that the Tesla driver was installed during the private image creation. If not, install the driver for computing acceleration after the ECS is created. For details, see [Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS](#).
- GPU-accelerated ECSs differ greatly in general-purpose and heterogeneous computing power. Their specifications can only be changed to other specifications of the same instance type.
- GPU-accelerated ECSs do not support live migration.

Inference-accelerated Pi2

Overview

Pi2 ECSs use NVIDIA Tesla T4 GPUs dedicated for real-time AI inference. These ECSs use the T4 INT8 calculator for up to 130 TOPS of INT8 computing. The Pi2 ECSs can also be used for light-load training.

Specifications

Table 4-49 Pi2 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	GPUs	GPU Memory (GiB)	Local Disks	Virtualization
pi2.2xlarge.4	8	32	10/4	50	4	4	1 × T4	1 × 16	N/A	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (10,000)	Max. NIC Queues	Max. NICs	GPUs	GPU Memory (GiB)	Local Disks	Virtualization
pi2.4xlarge.4	16	64	15/8	100	8	8	2 × T4	2 × 16	N/A	KVM
pi2.8xlarge.4	32	128	25/15	200	16	8	4 × T4	4 × 16	N/A	KVM
pi2.16xlarge.4	64	256	30/30	400	32	8	8 × T4	8 × 16	N/A	KVM

Pi2 ECS Features

- CPU: 2nd Generation Intel® Xeon® Scalable 6278 processors (2.6 GHz of base frequency and 3.5 GHz of turbo frequency), or Intel® Xeon® Scalable 6151 processors (3.0 GHz of base frequency and 3.4 GHz of turbo frequency)
- Up to four NVIDIA Tesla T4 GPUs on an ECS
- GPU hardware passthrough
- Up to 8.1 TFLOPS of single-precision computing on a single GPU
- Up to 130 TOPS of INT8 computing on a single GPU
- 16 GiB of GDDR6 GPU memory with a bandwidth of 320 GiB/s on a single GPU
- One NVENC engine and two NVDEC engines embedded

Supported Common Software

Pi2 ECSs are used in GPU-based inference computing scenarios, such as image recognition, speech recognition, and natural language processing. The Pi2 ECSs can also be used for light-load training.

Pi2 ECSs support the following commonly used software:

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

Notes

- After a Pi2 ECS is stopped, basic resources (including vCPUs, memory, image, and GPUs) are not billed, but its system disk is billed based on the disk capacity. If other products, such as EVS disks, EIP, and bandwidth are associated with the ECS, these products are billed separately.

NOTE

Resources will be released after a Pi2 ECS is stopped. If resources are insufficient at the next start, the start may fail. If you want to use such an ECS for a long period of time, do not stop the ECS or change its billing mode to yearly/monthly.

- Pi2 ECSs support automatic recovery when the hosts accommodating such ECSs become faulty.
- By default, Pi2 ECSs created using a public image have the Tesla driver installed.
- If a Pi2 ECS is created using a private image, make sure that the Tesla driver was installed during the private image creation. If not, install the driver for computing acceleration after the ECS is created. For details, see [Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS](#).
- GPU-accelerated ECSs differ greatly in general-purpose and heterogeneous computing power. Their specifications can only be changed to other specifications of the same instance type.
- GPU-accelerated ECSs do not support live migration.

4.5.9 AI-accelerated ECSs

AI-accelerated ECSs, powered by Ascend processors and software stacks, are dedicated for accelerating AI applications.

AI inference-accelerated ECSs use self-developed Ascend 310 processors for AI inference acceleration.

Public Images Supported by AI-accelerated ECSs

Table 4-50 Public images

Type	Series	Public Images
Enhanced AI inference-accelerated (type II)	Ai2	Ubuntu Server 18.04 64bit CentOS 7.8 64bit CentOS 7.6 64bit

Enhanced AI Inference-accelerated Ai2 (Type II)

Overview

Ai2 ECSs use Ascend 310 series processors for AI acceleration. Featuring low power consumption and high computing power, Ascend 310 series processors have significantly improved energy efficiency ratio (EER), promoting the wide application of AI inference. AI inference-accelerated Ai2 ECSs enable you to easily use the powerful processing capability of Ascend 310 series processors.

Ai2 ECSs are ideal for computer vision, smart campus, smart city, smart transportation, smart retail, Internet-based real-time communication, and video encoding and decoding scenarios.

Specifications

Table 4-51 Ai2 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max. / Assured Bandwidth	Max. PPS (10,000)	Ascend 310 series	Ascend RAM (GiB)	Max. NIC Queues	Max. NICs	Virtualization
ai2.4xlarge.4	16	64	20/6	280	1	24	8	8	KVM
ai2.8xlarge.4	32	128	30/12	550	2	48	8	8	KVM
ai2.16xlarge.4	64	256	36/24	800	4	96	16	8	KVM
ai2.24xlarge.4	96	384	40/36	850	6	144	32	8	KVM

Features

Ai2 ECSs have the following features:

- 1:4 ratio of vCPUs to memory
- CPU: 3rd Generation Intel® Xeon® Scalable 6348 processors (2.6 GHz of base frequency and 3.5 GHz of turbo frequency), or Intel® Xeon® Scalable 8378A processors (3.0 GHz of base frequency and 3.4 GHz of turbo frequency)
- Each Atlas 300I Pro accelerator card is equipped with one Ascend 310 series processor, including eight DaVinci AI cores and eight self-developed CPU cores.
- 140 TeraOPS of integer-precision (INT8) on a single GPU
- 24 GiB of GPU memory with a bandwidth of 204.8 GiB/s on a single GPU
- Built-in hardware video codec engine for HD video decoder (H.264/265) and JPEG image encoding and decoding

Notes

1. Ai2 ECSs support the following public images:
 - Ubuntu Server 18.04 64bit
 - CentOS 7.8 64bit
 - CentOS 7.6 64bit
2. Ai2 ECSs do not support modification of specifications.
3. Ai2 ECSs support automatic recovery when the hosts accommodating such ECSs become faulty.

Using an AI-accelerated ECS

Perform the following steps:

1. Create an ECS. For details, see [Step 1: Configure Basic Settings](#).
 - In the **Specifications** field, select AI-accelerated specifications.
 - In the **Image** field, select **Public image** or **Private image**.
 - **Public image:** The CANN 3.1.0 development kit has been included and environment variables have been configured in public images by default. You need to verify the environment availability.
 - **Private image:** You need to install the driver, firmware, and development kit, and configure environment variables by yourself. For details, see the *CANN Software Installation Guide* of the corresponding version in [Ascend Documentation](#).
2. Remotely log in to the ECS.

If your Ai1 ECS runs Linux, use an SSH password to log in to the ECS.
3. Verify the environment availability.

Use a sample for compilation and running. For details, see "Sample Overview" in the *Model Development Learning Map* of the corresponding CANN edition in [Ascend Documentation](#).

The sample shows how to classify images (decode, resize, and infer images) based on the Caffe ResNet-50 network.

5 Images

5.1 Image Types

What Is Image?

An image is an ECS template that contains an OS. It may also contain proprietary software and application software, such as database software. You can use images to create ECSs.

Images can be public or private. Public images are provided by the system by default, and private images are manually created. You can use a public or private image to create an ECS. You can also create a private image using an existing ECS. This provides you with a simple and fast way to create ECSs tailored to your needs. For example, if you use web services, your image can contain web server configurations, static configurations, and dynamic page code. After you use this image to create an ECS, the web server will run on the created ECS.

Image Types

Image Type	Description
Public	<p>A public image is a standard, widely used image. It contains an OS and preinstalled public applications and is available to all users. Public images are very stable and their OS and any included software have been officially authorized for use. If a public image does not contain the environment or software you need, you can use a public image to create an ECS and then deploy the required environment or software on it.</p> <p>Public images include the following OSs to choose from: Ubuntu and CentOS</p> <p>For more information, see Managing Public Images.</p>

Image Type	Description
Private	<p>A private image contains an OS or service data, preinstalled public applications, and a user's personal applications. Private images are only available to the users who created them.</p> <p>A private image can be a system disk image, data disk image, ISO image, or full-ECS image.</p> <ul style="list-style-type: none">• A system disk image contains an OS and preinstalled software for various services. You can use a system disk image to create ECSs and migrate your services to the cloud.• A data disk image contains only service data. You can use a data disk image to create EVS disks and use them to migrate your service data to the cloud.• An ISO image is created from an external ISO image file. It is a special image that is not available on the ECS console.• A full-ECS image contains an OS, preinstalled software, and service data. A full-ECS image is created using differential backups and the creation takes less time than creating a system or data disk image that has the same disk capacity.
Shared	<p>A shared image is a private image another user has shared with you.</p> <p>For more information, see Sharing Images.</p>

5.2 Cloud-Init

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

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

Impact on IMS

To ensure that ECSs that are created using a private image support custom configurations, you must install Cloud-Init on the ECSs before using them to create private images.

After being installed in an image, Cloud-Init automatically configures initial attributes for the ECSs created using this image.

For more information, see [Installing Cloud-Init](#).

Impact on ECS

- When creating an ECS, if the selected image supports Cloud-Init, you can use the **User Data** function to specify custom configuration, such as ECS login password to the ECS. Such custom settings will take effect upon ECS initialization.
- If 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 the **User Data** function 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 based on the specified ECS 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 may fail to take effect. To resolve this issue, see [How Does Cloud-Init Affect the AS Service?](#) in *Auto Scaling User Guide*.

Notes

- When using Cloud-Init, enable DHCP in the VPC which the ECS belongs to.
- When using Cloud-Init, ensure that security group rules for the outbound direction meet the following requirements:
 - **Protocol: TCP**
 - **Port: 80**
 - **Destination: 169.254.0.0/16**

NOTE

If you use the default security group rules for the outbound direction, the metadata can be accessed because the default rules meet the preceding requirements. For details about the default security group rules for the outbound direction, see [Security Group](#).

6 EVS Disks

What Is Elastic Volume Service?

Elastic Volume Service (EVS) offers scalable block storage for ECSs. With high reliability, high performance, and rich specifications, EVS disks can be used for distributed file systems, development and test environments, data warehouses, and high-performance computing (HPC) scenarios to meet diverse service requirements.

Disk Types

EVS disk types differ in performance. Choose a disk type based on your requirements.

For more information about EVS disk specifications and performance, see [Elastic Volume Service User Guide](#).

Device Types

EVS disks have two device types, Virtual Block Device (VBD) and Small Computer System Interface (SCSI).

- VBD
When you create an EVS disk on the management console, **Device Type** of the EVS disk is VBD by default. VBD EVS disks support only simple SCSI read/write commands.
- SCSI
You can create EVS disks whose **Device Type** is SCSI on the management console. These EVS disks support transparent SCSI command transmission, allowing ECS OS to directly access underlying storage media. SCSI EVS disks support both basic and advanced SCSI commands.

NOTE

For more information about how to use SCSI EVS disks, for example, how to install a driver for SCSI EVS disks, see [Device Types and Usage Instructions](#).

Helpful Links

- [Attaching an EVS Disk to an ECS](#)
- [Initializing an EVS Data Disk](#)
- [What Are the Requirements for Attaching an EVS Disk to an ECS?](#)

7 Network

VPC

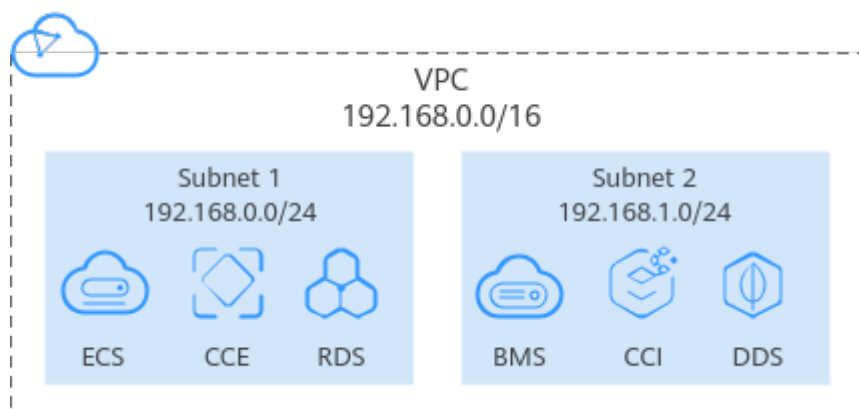
Virtual Private Cloud (VPC) allows you to create customized virtual networks in your logically isolated AZ. Such networks are dedicated zones that are logically isolated, providing secure network environments for your ECSs. You can define security groups, virtual private networks (VPNs), IP address segments,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

For more information about VPC, see [Virtual Private Cloud Service Overview](#).

Subnets

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 7-1 Subnets



By default, ECSs in all subnets of the same VPC can communicate with each other, 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 that are mutually trusted. By adding an ECS to a security group, you apply all the rules defined for this security group to this ECS.

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

Figure 7-2 Default security group




Table 7-1 describes the rules in the default security group.

Table 7-1 Default security group rules

Direction	Action	Type	Protocol & Port	Source/Destination	Description
Inbound	Allow	IPv4	All	Source: default security group (default)	Allows IPv4 instances in the security group to communicate with each other using any protocol over any port.
Inbound	Allow	IPv6	All	Source: default security group (default)	Allows IPv6 instances in the security group to communicate with each other using any protocol over any port.
Outbound	Allow	IPv4	All	Destination: 0.0.0.0/0	Allows all traffic from the instances in the security group to any IPv4 address over any port.

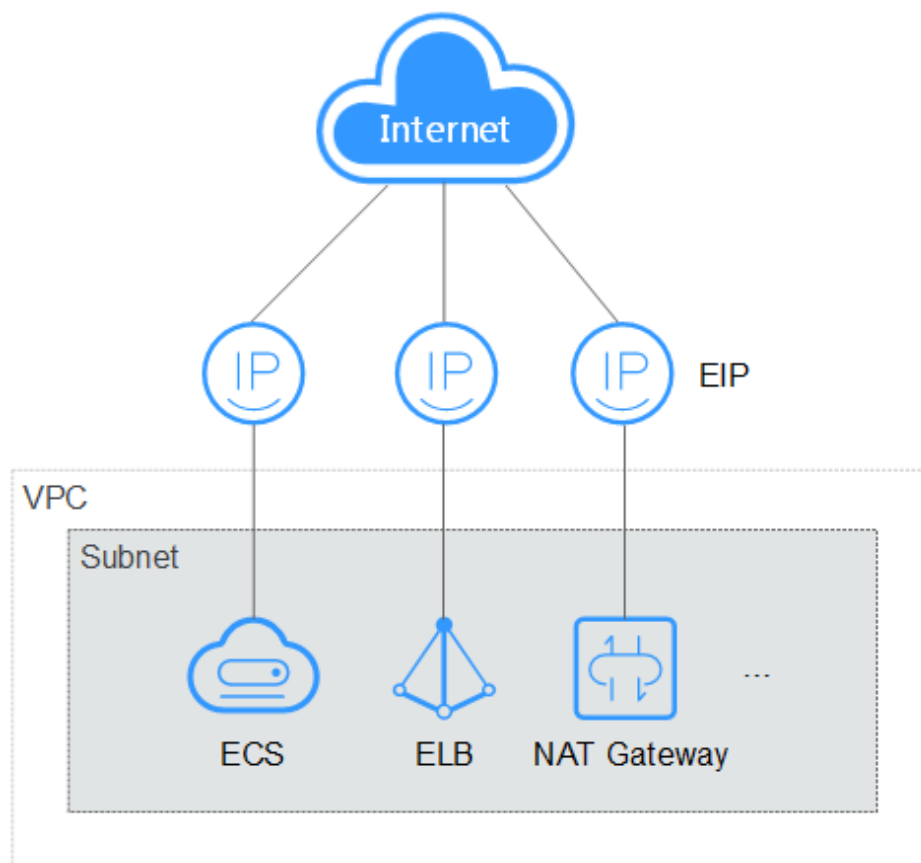
Direction	Action	Type	Protocol & Port	Source/Destination	Description
Outbound	Allow	IPv6	All	Destination: : :/0	Allows all traffic from the instances in the security group to any IPv6 address over any port.

EIP

The Elastic IP (EIP) service enables your cloud resources to communicate with the Internet using static public IP addresses and scalable bandwidths. EIPs can be bound to or unbound from ECSs, BMSs, virtual IP addresses, NAT gateways or load balancers. Various billing modes are provided to meet different service requirements.

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

Figure 7-3 Accessing the Internet using an EIP



8 Security

8.1 Identity Authentication and Access Control

8.1.1 Access Control for ECS

IAM Identity Authentication

IAM provides fine-grained permissions management, user identity authentication, and resource access control.

You can use your account to create IAM users, and assign permissions to the IAM users to control their access to specific resources. IAM permissions define which actions on your cloud resources are allowed or denied.

- For details about permissions management, see [Permissions Management](#).
- For details about how to grant ECS permissions, see [Creating a User and Granting ECS Permissions](#).
- For details about custom policies, see [ECS Custom Policies](#).
- For details about policies and supported actions, see [Permissions and Supported Actions](#).

Project and Enterprise Project

You can group, manage, and isolate resources by project or enterprise project to control resource access and manage permissions by organization like enterprise, department, or project team.

- Project
Projects in IAM are used to group and isolate OpenStack resources (computing resources, storage resources, and network resources). Resources in your account must be mounted under projects. A project can be a department or a project team. Multiple projects can be created for one account.
- Enterprise Project
Enterprise projects are used to categorize and manage multiple resources. Resources in different regions can belong to one enterprise project. An

enterprise can classify resources based on department or project group and put relevant resources into one enterprise project for management. Resources can be migrated between enterprise projects.

- Differences Between Projects and Enterprise Projects


- IAM project

Projects are used to categorize and physically isolate resources in a region. Resources in an IAM project cannot be transferred. They can only be deleted and then rebuilt.



- Enterprise project

Enterprise projects provide advanced functions than IAM projects and can be used to group and manage resources of different IAM projects of an enterprise. An enterprise project can contain resources across multiple regions, and resources can be added to or removed from enterprise projects. If you have enabled enterprise management, you cannot create IAM projects anymore and can only manage existing projects. In the future, enterprise projects will replace IAM projects.



Both projects and enterprise projects can be managed by one or more user groups. Users who manage enterprise projects are in user groups. By attaching a policy to a user group, users in the group can obtain the permissions defined in the policy in the project or enterprise project.

For details about how to create a project, create an enterprise project, and assign permissions, see [Project and Enterprise Project](#).

Access Control

- 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, providing secure network environments 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 and allows you to change your network in a secure and convenient network manner. You can also customize the ECS access rules within a security group and between security groups to improve ECS security.

- Security Group

A security group is a collection of access control rules for ECSs that have the same security protection requirements and that are mutually trusted. By adding an ECS to a security group, you apply all the rules defined for this security group to this ECS.

For details about how to set a VPC and security group, see [Step 2: Configure Network](#).

8.2 Data Protection

8.2.1 Hypervisor Security

Hypervisor

The hypervisor isolates ECSs running on a physical server to prevent data theft and attacks. You can only access the ECSs allocated to you, including hardware and software resources and data.

CPU Isolation

The x86 architecture supports four privilege levels ranging from ring 0 to ring 3 with privileges in descending order.

- Ring 0: OS kernel
- Ring 1: OS services
- Ring 2: privileged code (user applications with I/O access permissions)
- Ring 3: applications

Each privilege level separately restricts the commands that can be executed.

- The hypervisor runs at ring 0.
- ECS OSs run at ring 1. This prevents the Guest OS from executing privileged instructions.
- Applications run at ring 3, ensuring that they are isolated from the OS.

The introduction of CPU hardware-assisted virtualization further isolates the hypervisor from ECS OSs.

Memory Isolation

The hypervisor uses memory virtualization to isolate the virtual memory of different ECSs. Memory virtualization has introduced layer 3 addressing (virtual guest address into physical guest address and into machine address) based on traditional layer 2 addressing (virtual address into machine address). The process is as follows: An ECS translates its virtual guest address into a physical guest address. Then, the hypervisor translates the physical guest address into a machine address, and sends the machine address to the CPU for processing. This allows the ECS to access only the physical memory allocated it and prevents the ECS from obtaining the machine address.

I/O Isolation

The hypervisor uses a separate device model for I/O virtualization. The frontend transfers the I/O requests from an ECS to the hypervisor backend. The backend parses the requests and sends them to the target device to finish I/O operations. The hypervisor ensures that the ECS can access only the I/O resources allocated to it.

8.2.2 User Encryption

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

Image Encryption

Image encryption supports encrypting private images. When creating an ECS, if you select an encrypted image, the system disk of the created ECS is automatically encrypted, improving data security.

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

- Use an external image file.
- Use an existing encrypted ECS.

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, 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.
- When creating an ECS, you can encrypt added data disks.

For more information about EVS disk encryption, see [EVS 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 uses a customer master key (CMK) to encrypt the DEK.

Figure 8-1 Data encryption process

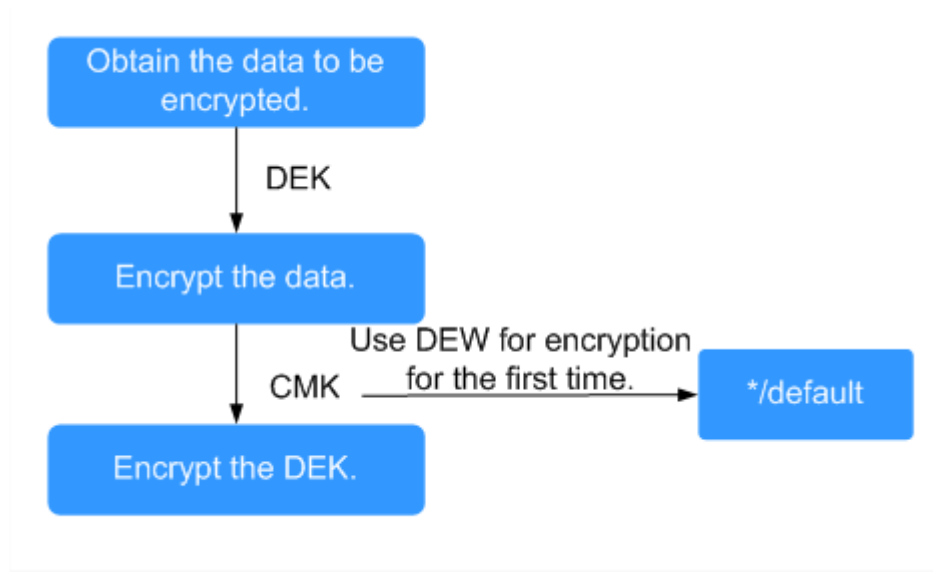


Table 8-1 describes the keys involved in the data encryption process.

Table 8-1 Keys

Name	Description	Function
DEK	An encryption key that is used for encrypting data.	Encrypts specific data.
Custom key	An encryption key created using DEW for encrypting DEKs. A custom key can encrypt multiple DEKs.	Supports CMK disabling and scheduled deletion.
Default key	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 .	<ul style="list-style-type: none"> Supports query of the default key on the DEW 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, so the EVS disk becomes unavailable.

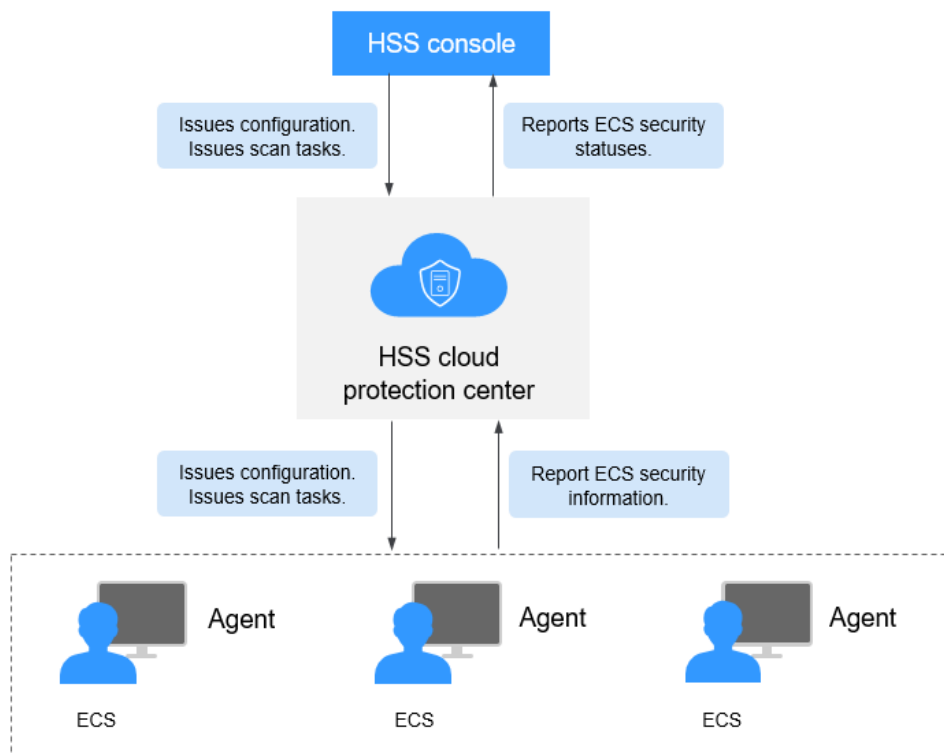
For details about DEW, see [Data Encryption Workshop User Guide](#).

8.2.3 HSS

Host Security Service (HSS) is designed to improve the overall security for ECSs. It reduces intrusion risks with asset management, vulnerability management, intrusion detection, and baseline inspection functions.

After installing the HSS agent on your ECSs, you will be able to check the ECS security status and risks in a region on the HSS console.

Figure 8-2 HSS working rules



The following table describes HSS components.

- **HSS console**
A visualized management platform, where you can centrally apply configurations and view the defense status and scan results of ECSs in a region.
- **HSS cloud protection center**
 - Receives configurations and scan tasks sent from the console and forwards them to agents on the ECSs.
 - Receives ECS information reported by agents, analyzes security risks and exceptions on ECSs, and displays the analysis results on the console.
- **Agent**
 - Receives detection tasks forwarded by the HSS cloud protection center.
 - Scans ECSs based on task requirements and reports scanning results to the HSS cloud protection center.

Before using the HSS service, install the HSS agent on your ECS. For details, see [Host Security Service](#).

8.3 Fault Recovery

Cloud Backup and Recovery (CBR) lets you back up and restore data in case of a failure. If an ECS or EVS disk is faulty or data is deleted accidentally, you can use data backups to quickly restore data.

What Is CBR?

CBR enables you to back up ECSs and EVS disks with ease. If any exceptions occur, such as virus attack, accidental deletion, or software or hardware fault, you can restore data to any point in the past when the data was backed up.

CBR secures your services by ensuring the security and consistency of your data.

Differences Between Cloud Server Backup and Cloud Disk Backup

You can back up ECS data using Cloud Server Backup or Cloud Disk Backup.

- **Cloud Server Backup (recommended):** Use this backup function if you want to back up the data of all EVS disks (system and data disks) on an ECS. This prevents data inconsistency caused by time difference in creating a backup.
- **Cloud Disk Backup:** Use this backup function if you want to back up the data of one or more EVS disks (system or data disk) on an ECS. This minimizes backup costs on the basis of data security.

Table 8-2 Differences between cloud server backup and cloud disk backup

Item	Cloud Server Backup	Cloud Disk Backup
Resources to be backed up or restored	All disks (system and data disks) on a server	One or more specified disks (system or data disks)
Recommended scenario	An entire cloud server needs to be protected.	Only data disks need to be backed up, because the system disk does not contain users' application data.
Advantages	All disks on a server are backed up at the same time, ensuring data consistency.	Backup cost is reduced without compromising data security.

Helpful Links

- [Purchasing a Server Backup Vault](#)
- [Purchasing a Disk Backup Vault](#)
- [Restoring from a Cloud Server Backup](#)
- [Using a Backup to Create an Image](#)

- [Restoring from a Cloud Disk Backup](#)

8.4 License Types

BYOL

What Is BYOL?

Bring your own license (BYOL) allows you to use your existing OS license. In such a case, you do not need to apply for a license again.

How to Use BYOL?

If you select the BYOL license type, you are required to manage licenses by yourself. If you have obtained an OS license, you do not need to apply for a license.

The OSs supporting BYOL are as follows:

- Windows

Application Scenarios

The system does not support dynamic license type changing. ECSs support BYOL in the following scenarios:

- Creating an ECS
After creating an ECS, you cannot change its license type. If the license type must be changed, reinstall or change the ECS OS.
- Reinstalling an ECS OS
When reinstalling an ECS OS, you can set the license type for the ECS.
- Changing an ECS OS
When changing an ECS OS, you can set the license type for the ECS.
- Attaching a system disk
The license type of a system disk is determined by the ECS license type after the ECS is created, the ECS OS is reinstalled, or the ECS OS is changed. If the system disk is detached and then attached to a new ECS or the original ECS, ensure that the ECS license type is the same as the system disk license type.

9 Billing

Billing Items

ECSs are billed based on ECS specifications and the service duration.

Table 9-1 ECS billing

Billing Item	Description
ECS	Pricing is based on the ECS type, flavor (including vCPUs and memory), service duration, and the number of purchased ECSs. For pricing details, see Elastic Cloud Server Pricing Details .
Image	Public images of the community edition, such as Linux, are free of charge.
EVS disk	EVS disks are mandatory. A system disk is 40 GB by default. EVS disks can be billed on a pay-per-use or yearly/monthly basis. For pricing details, see Elastic Volume Service Pricing Details . The usage duration of EVS disks should be the same as the associated ECS.
EIP	A public IP address is required for public accessibility. For pricing details, see Elastic IP Pricing Details .
Bandwidth	An EIP can be billed by bandwidth or traffic. For pricing details, see Billed By Bandwidth .

Billing Modes

An ECS can be billed on a pay-per-use, spot pricing, or yearly/monthly basis.

- **Yearly/Monthly:** The ECS will be billed based on the service duration. This cost-effective mode is ideal when the duration of ECS usage is predictable.
- **Pay-per-use:** The ECS will be billed based on usage frequency and duration. This mode is ideal when you want more flexibility and control on ECS usage.

- Spot pricing: The ECS will be billed based on the price that is effective for the time it is being used. This mode is more cost-effective than pay-per-use, and the price will be adjusted based on supply-and-demand changes.

Table 9-2 lists the differences between the billing modes.

Table 9-2 Billing modes

Billing Mode	Yearly/Monthly	Pay-per-Use	Spot Pricing
Payment Method	Prepaid Billed by the purchased duration specified in the order.	Postpaid Billed by service duration.	Postpaid Billed at the market price, which varies according to the changes in supply and demand. The start price of the bill is the market price when the ECS was purchased, and then the ECS is billed at the market price on the hour. Learn more about spot pricing
Billing Period	Billed by the purchased duration specified in the order.	Billed by the second and settled by the hour.	Billed by the second and settled by the hour.

<p>Billing for Stopped ECS</p>	<p>Billed by the purchased duration specified in the order, regardless of whether the ECS is stopped or not.</p>	<ul style="list-style-type: none"> • After a common ECS (without local disks or FPGAs attached) is stopped, basic resources (vCPUs, memory, image, and GPUs) are not billed. Other resources (EVS disks, EIPs, and bandwidth) associated with the ECS will continue to be billed. • An ECS with local disks attached, an FPGA-based ECS, or a bare metal ECS will continue to generate costs even after being stopped. Its vCPUs and memory are retained. To stop the ECS from being billed, delete it and its associated resources. 	<ul style="list-style-type: none"> • After a common ECS (without local disks or FPGAs attached) is stopped, basic resources (vCPUs, memory, image, and GPUs) are not billed. Other resources (EVS disks, EIPs, and bandwidth) associated with the ECS will continue to be billed. • An ECS with local disks attached, an FPGA-based ECS, or a bare metal ECS will continue to generate costs even after being stopped. Its vCPUs and memory are retained. To avoid being billed for such an ECS, delete it and its associated resources.
<p>Billing Mode Change</p>	<p>Can be changed to pay-per-use. Changing Yearly/Monthly to Pay-per-Use</p>	<p>Can be changed to yearly/monthly. Changing the Billing Mode from Pay-per-Use to Yearly/Monthly</p>	<p>Cannot be changed to pay-per-use or yearly/monthly.</p>
<p>Specification Modification</p>	<p>Supported</p>	<p>Supported</p>	<p>Not supported</p>

Applications Scenarios	This cost-effective mode is ideal when the duration of ECS usage is predictable. The yearly/monthly mode is recommended for long-term users.	This mode is ideal when you want more flexibility and control on ECS usage.	Spot ECSs can be reclaimed at any time and are suitable for stateless, fault-tolerant instances that are not sensitive to interruptions.
-------------------------------	--	---	--

- **Yearly/Monthly:** This mode provides a larger discount than pay-per-use and is recommended for long-term users. A yearly/monthly ECS is billed based on the purchased duration specified in the order.
- **Pay-per-use:** a flexible mode with the billing accurately down to the second. An ECS is billed from the time when it is provisioned to the time when it is deleted.

Common ECSs refer to ECSs without local disks or FPGAs attached. After a common ECS is stopped, it is billed as follows:

- ECS basic resources (vCPUs, memory, image, and GPUs) no longer generate costs. Its associated resources such as its EVS disks, EIPs, and bandwidth will continue to be billed.
- When you try to start the ECS the next time, the system will allocate vCPUs and memory again, but if resources are insufficient, the startup may fail. In this case, you can try again later or resize the ECS specifications first before trying to start it.

Special pay-per-use ECSs will continue to be billed after being stopped and its resources such as vCPUs and memory are still retained.

 **NOTE**

Special ECSs include:

- Bare metal ECSs
- ECSs attached with local disks, such as disk-intensive ECSs and ultra-high I/O ECSs
- FPGA-based ECSs

To stop billing for special ECSs, delete them and their associated resources.

- **Spot pricing**

Price: Spot ECSs are billed based on the market price, which varies according to the changes in supply and demand. The maximum price you are willing to pay is not used as a billing basis. A higher price ensures a greater success rate for you to purchase such an ECS. A spot ECS can be used only when the market price is lower than the maximum price you are willing to pay and inventory resources are sufficient. When the market price exceeds the maximum price, the ECS will be reclaimed.

Billing period: A spot ECS is billed by the second. A bill is generated on the hour. The start price of the bill is the market price when the ECS was purchased, and the market price of the hour is used for billing.

Associated services: Spot pricing applies only to vCPUs and memory. The OS, system disk, data disk, bandwidth, and IP address are billed following the billing rules for these items in pay-per-use billing mode. A system disk is created and released with the ECS to which the system disk is attached. A data disk must be manually deleted.

Coupons: Coupons are not applicable to spot ECSs.

Billing Examples

In both pay-per-use and spot pricing billing modes, ECSs are billed by the second. The price per second of each type of ECS can be obtained by dividing their hourly price by 3600. Obtain the hourly price on the **Product Pricing Details** page.

For example, if you purchase a pay-per-use ECS priced \$0.68 USD/hour, the ECS will be billed based on the usage duration by the second.

- If you use the ECS for 30 minutes, you need to pay for \$0.34 USD ($0.68/3600 \times 30 \times 60$).
- If you use the ECS for 1 hour and 30 minutes, you need to pay for \$1.02 USD ($0.68/3600 \times 90 \times 60$).

Configuration Changes

- Changing the billing mode
 - Changing from pay-per-use to yearly/monthly: After the billing mode of an ECS is changed from pay-per-use to yearly/monthly, a new order is then generated for you, and the new billing mode takes effect immediately after you pay for the order.
 - Changing from yearly/monthly to pay-per-use: The pay-per-use billing mode will take effect after the original yearly/monthly subscription has expired.
 - A spot ECS cannot be changed to a pay-per-use or yearly/monthly ECS.

- Modifying ECS specifications

For an ECS billed on a yearly/monthly or pay-per-use basis, you can modify its specifications (vCPUs and memory). However, a spot ECS does not support specifications modification.

Notes

- Vouchers will not be refunded if the specifications of the ECS purchased with those vouchers are downgraded.
- If ECS specifications are upgraded, the price difference between the original and new specifications must be returned according to the in-service duration.
 - For pay-per-use ECSs: The specification modifications take effect immediately.
 - For yearly/monthly ECSs: The specification modifications take effect immediately within the original validity period. For details about the price difference, see [Pricing of a Changed Specification](#).

- ECS specifications (vCPU or memory) degrade deteriorates the ECS performance.
- The price difference must be reimbursed if a downgraded ECS needs to be upgraded back to its original specifications.

Helpful Links

- [What Are the Differences Between Yearly/Monthly and Pay-per-Use Billing Modes?](#)
- [Will Am I Continue to Be Billed After ECSs Are Stopped?](#)
- [Can I Switch Between Yearly/Monthly and Pay-per-Use Payments?](#)
- [FAQs About ECS Frozen, Deletion, and Unsubscription](#)
- [How Can I Stop an ECS from Being Billed?](#)
- [FAQs About Spot ECSs](#)

10 Constraints

This section describes notes and constraints on using ECSs.

Precautions

- Do not use ECSs as unauthorized servers for any illegal or violation activities, such as gambling or cross-border VPN.
- Do not use ECSs for fraudulent 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 to spread viruses and Trojan horses.
- Do not use ECSs for traffic transit.
- Do not use ECSs for web crawling.
- Do not use ECSs to detect other systems like scanning or penetration unless otherwise being authorized.
- Do not deploy any illegal websites or applications on ECSs.
- Do not use ECSs to send spams or engage in activities that violate personal privacy.

Restrictions on using ECSs

- Do not uninstall drivers on the ECS hardware.
- Do not install external hardware devices, such as encryption dongles, USB flash drives, external hard disks, or bank USB security keys on ECSs.
- Do not change the MAC address of NICs.
- Do not install virtualization software on ECSs for nested virtualization.
- Do not associate software licenses with the physical server hosting an ECS. Once an ECS is migrated from one physical server to another, the associated licenses may become invalid.
- Do not deploy applications on a single ECS if you require high availability. Set up auto start for your ECSs or deploy applications in cluster or active/standby mode.
- Data on ECSs running core applications needs to be backed up.
- Monitoring needs to be configured for ECSs.

- Do not change the default DNS server address. If you need to configure a public DNS address, configure both a public and a private DNS address on your ECS.
- The system disk can boot from Basic Input Output System (BIOS) or Unified Extensible Firmware Interface (UEFI) according to the boot mode in the image file.
 - You can change the OS to convert the boot mode of the ECS.
 - You can create a UEFI or BIOS private image and use it to create an ECS.

Precautions for Using Windows ECSs

- Do not stop system processes if you are not sure about the consequences. Otherwise, BSOD or a restart may occur on the ECS.
- Ensure that there is at least 2 GiB of idle memory. Otherwise, BSOD, freezing, or service failures may occur.
- Do not modify the registry. Otherwise, the system startup 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 disable virtual memory. Otherwise, system performance may deteriorate, or system exceptions may occur.
- Do not delete the VMTool program, or an exception may occur on the ECS.

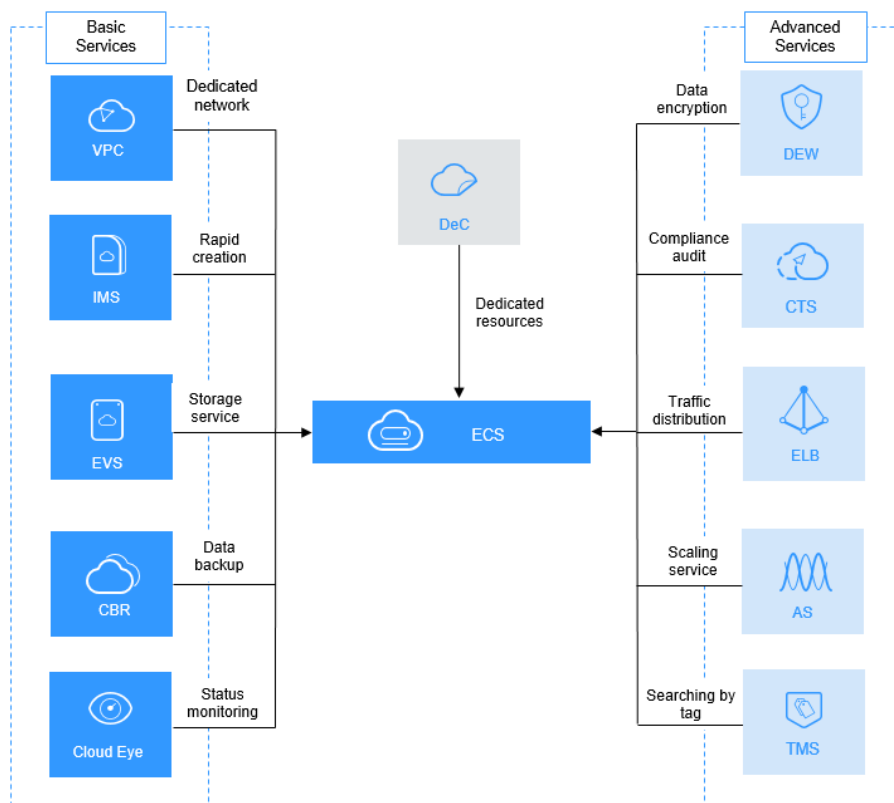
Precautions for Using Linux ECSs

- Do not modify the `/etc/issue` file. Otherwise, the OS distribution will not be identified.
- Do not delete system directories or files. Otherwise, the system may fail to run or start.
- Do not change the permissions for or names of system directories. Otherwise, the system may fail to run or start.
- Do not upgrade the kernel of the Linux unless necessary.
When you have to upgrade the Linux kernel, follow the instructions provided in [How Can I Upgrade the Kernel of a Linux ECS?](#)
- Do not change the default `/etc/resolv.conf` of the DNS server. Otherwise, software sources and NTP may be unavailable.
- Do not modify default intranet configurations, such as the IP address, subnet mask, or gateway address of an ECS. Otherwise, network exceptions may occur.
- Manually specified IP addresses for Linux ECSs are generally static IP addresses. To avoid network exceptions caused by conflicts between NetworkManager and internal network services, do not enable NetworkManager when not required, such as when installing Kubernetes.

11 ECS and Other Services

Figure 11-1 shows the relationships between ECS and other services.

Figure 11-1 Relationships between ECS and other services



ECS-related Services

Table 11-1 ECS-related services

Service	Function	Related Operation
Auto Scaling (AS)	Automatically adjusts ECS resources to keep up with changes in demand based on the configured AS policies. This improves resource utilization and reduces costs.	<ul style="list-style-type: none">• Using an Existing ECS to Create an AS Configuration• Using a New Specifications Template to Create an AS Configuration
Elastic Load Balance (ELB)	Automatically distributes traffic to multiple ECSs. This improves fault tolerance.	Adding Backend Servers
Elastic Volume Service (EVS)	Allows you to attach EVS disks to an ECS and expand their capacity.	<ul style="list-style-type: none">• Attaching a Non-Shared Disk• Attaching a Shared Disk
Virtual Private Cloud (VPC)	Allows you to create your own dedicated network on the cloud by customizing security groups, VPNs, IP address ranges, and bandwidths. This simplifies network management. You can also customize the ECS access rules within a security group and between security groups to improve ECS security.	<ul style="list-style-type: none">• Assigning an EIP and Binding It to an ECS• Adding a Security Group Rule
Image Management Service (IMS)	Allows 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.	<ul style="list-style-type: none">• Creating a Data Disk Image Using an ECS Data Disk• Creating a Full-ECS Image Using an ECS
Cloud Eye	Allows you to check the status of monitored ECS resources. This can be done without requiring additional plug-ins.	<ul style="list-style-type: none">• Basic ECS Metrics• ECS Metrics Under OS Monitoring (with Agent Installed)
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 need to use the key provided by DEW to improve data security.	<ul style="list-style-type: none">• EVS Disk Encryption• Encrypting Images

Service	Function	Related Operation
Cloud Trace Service (CTS)	Records ECS-related operations for later query, auditing, and backtracking.	<ul style="list-style-type: none">• Key Operations Supported by CTS
Cloud Backup and Recovery (CBR)	Backs up EVS disks and ECSs for fault recovery.	Purchasing a Server Backup Vault
Tag Management Service (TMS)	Identifies ECSs to help classification and search.	<ul style="list-style-type: none">• Adding Tags• Searching for Resources by Tag

12 Permissions Management

If you need to assign different permissions to employees in your enterprise to access your ECS resources, IAM is a good choice for fine-grained permissions management. IAM provides identity authentication, permissions management, and access control, helping you secure access to your resources.

With IAM, you can use your account to create IAM users, and assign permissions to the users to control their access to specific resources. For example, some software developers in your enterprise need to use ECS resources but should not be allowed to delete the resources or perform any other high-risk operations. In this scenario, you can create IAM users for the software developers and grant them only the permissions required for using ECS resources.

If your account does not need individual IAM users for permissions management, skip this section.

IAM is a free service. You pay only for the resources in your account. For more information about IAM, see [IAM Service Overview](#).

ECS Permissions

By default, new IAM users do not have permissions assigned. You need to add a user to one or more groups, and attach permissions policies or roles to these groups. Users inherit permissions from the groups to which they are added and can perform specified operations on cloud services based on the permissions.

ECS is a project-level service deployed and accessed in specific physical regions. To assign ECS permissions to a user group, specify the scope as region-specific projects and select projects (such as **eu-west-101** for the **EU-Dublin** region) for the permissions to take effect. If you select **All projects**, the permissions will take effect for user groups in all region-specific projects. When accessing ECS, the users need to switch to a region where they have got permissions to use this service.

You can grant users permissions by using roles and policies.

- **Roles:** A type of coarse-grained authorization mechanism that defines permissions related to user responsibilities. This mechanism provides only a limited number of service-level roles for authorization. When using roles to grant permissions, you also need to assign other roles which the permissions depend on to take effect. However, roles are not an ideal choice for fine-grained authorization and secure access control.

- **Policies:** A fine-grained authorization strategy that defines permissions required to perform operations on specific cloud resources under certain conditions. This mechanism allows for more flexible policy-based authorization, meeting requirements for secure access control. For example, you can grant ECS users only the permissions for managing a certain type of ECSs.

Most policies define permissions based on APIs. For the API actions supported by ECS, see [Permissions and Supported Actions](#).

[Table 12-1](#) and [Table 12-2](#) list all ECS system-defined policies and roles.

Table 12-1 ECS system-defined policies (recommended)

Policy/Role Name	Description	Policy Content
ECS FullAccess	Administrator permissions for ECS. Users granted these permissions can perform all operations on ECSs, including creating, deleting, and viewing ECSs, and modifying ECS specifications.	ECS FullAccess Policy Content
ECS CommonOperations	Common user permissions for ECS. Users granted these permissions can start, stop, restart, and query ECSs.	ECS CommonOperations Policy Content
ECS ReadOnlyAccess	Read-only permissions for ECS. Users granted these permissions can only view ECS data.	ECS ReadOnlyAccess Policy Content

Table 12-2 ECS system-defined roles

Role Name	Description	Role Content
Server Administrator	<p>Full permissions for ECS. This role must be used together with the Tenant Guest role in the same project.</p> <p>If a user needs to create, delete, or change resources of other services, the user must also be granted administrator permissions of the corresponding services in the same project.</p> <p>For example, if a user needs to create a new VPC when creating an ECS, the user must also be granted permissions with the VPC Administrator role.</p>	Server Administrator Role Content

Table 12-3 lists the common operations supported by each system-defined policy of ECS. Select the policies as required.

Table 12-3 Common operations supported by each system-defined policy

Operation	ECS FullAccess	ECS CommonOperations	ECS ReadOnlyAccess
Creating an ECS	Supported	Not supported	Not supported
Remotely logging in to an ECS on the management console	Supported	Supported	Not supported (VNC login not supported)
Querying an ECS list	Supported	Supported	Supported
Querying ECS details	Supported	Supported	Supported
Modifying ECS details	Supported	Not supported	Not supported
Starting an ECS	Supported	Supported	Not supported
Stopping an ECS	Supported	Supported	Not supported
Restarting an ECS	Supported	Supported	Not supported
Deleting an ECS	Supported	Not supported	Not supported
Reinstalling an ECS OS	Supported	Not supported	Not supported
Changing an ECS OS	Supported	Not supported	Not supported
Attaching a disk to an ECS	Supported	Not supported	Not supported
Detaching a disk from an ECS	Supported	Not supported	Not supported
Querying a disk list	Supported	Supported	Supported
Attaching a NIC to an ECS	Supported	Not supported	Not supported
Detaching a NIC from an ECS	Supported	Not supported	Not supported
Querying a NIC list	Supported	Supported	Supported
Adding tags to an ECS	Supported	Supported	Not supported
Modifying ECS specifications	Supported	Not supported	Not supported
Querying the ECS flavor list	Supported	Supported	Supported
Querying ECS groups	Supported	Supported	Supported

Helpful Links

- [IAM Service Overview](#)
- [Creating a User and Granting ECS Permissions](#)
- [Permissions and Supported Actions](#) in *Elastic Cloud Server API Reference*

ECS FullAccess Policy Content

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ecs:*",
        "evs:get",
        "evs:list",
        "evs:volumes:create",
        "evs:volumes:delete",
        "evs:volumes:attach",
        "evs:volumes:detach",
        "evs:volumes:manage",
        "evs:volumes:update",
        "evs:volumes:use",
        "evs:volumes:uploadImage",
        "evs:snapshots:create",
        "vpc:get",
        "vpc:list",
        "vpc:networks:create",
        "vpc:networks:update",
        "vpc:subnets:update",
        "vpc:subnets:create",
        "vpc:ports:*",
        "vpc:routers:get",
        "vpc:routers:update",
        "vpc:securityGroups:*",
        "vpc:securityGroupRules:*",
        "vpc:floatingIps:*",
        "vpc:publicIps:*",
        "ims:images:create",
        "ims:images:delete",
        "ims:images:get",
        "ims:images:list",
        "ims:images:update",
        "ims:images:upload"
      ]
    }
  ]
}
```

ECS CommonOperations Policy Content

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ecs:get*",
        "ecs:list*",
        "ecs:start",
        "ecs:stop",
        "ecs:reboot",
        "ecs:blockDevice:use",
        "ecs:cloudServerFpgaImages:relate",
        "ecs:cloudServerFpgaImages:register",
        "ecs:cloudServerFpgaImages:delete",

```

```
"ecs:cloudServerFpgaImags:unrelate",
"ecs:cloudServers:setAutoRecovery",
"ecs:cloudServerPasswords:reset",
"ecs:cloudServerPorts:modify",
"ecs:cloudServers:vnc",
"ecs:diskConfigs:use",
"ecs:securityGroups:use",
"ecs:serverGroups:manage",
"ecs:serverFloatingIps:use",
"ecs:serverKeypairs:*",
"ecs:serverPasswords:manage",
"ecs:servers:createConsole",
"ecs:servers:createImage",
"ecs:servers:setMetadata",
"ecs:servers:setTags",
"ecs:serverVolumes:use",
"evs:*:get*",
"evs:*:list*",
"evs:snapshots:create",
"evs:volumes:uploadImage",
"evs:volumes:delete",
"evs:volumes:update",
"evs:volumes:attach",
"evs:volumes:detach",
"evs:volumes:manage",
"evs:volumes:use",
"vpc:*:get*",
"vpc:*:list*",
"vpc:floatingIps:create",
"vpc:floatingIps:update",
"vpc:floatingIps:delete",
"vpc:publicIps:update",
"vpc:publicIps:delete",
"ims:images:create",
"ims:images:delete",
"ims:images:get",
"ims:images:list",
"ims:images:update",
"ims:images:upload"
    ]
  }
}
```

ECS ReadOnlyAccess Policy Content

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ecs:*:get*",
        "ecs:*:list*",
        "ecs:serverGroups:manage",
        "ecs:serverVolumes:use",
        "evs:*:get*",
        "evs:*:list*",
        "vpc:*:get*",
        "vpc:*:list*",
        "ims:*:get*",
        "ims:*:list*"
      ]
    }
  ]
}
```


Server Administrator Role Content

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Action": [
        "ecs:*",
        "evs:*:get",
        "evs:*:list",
        "evs:volumes:create",
        "evs:volumes:delete",
        "evs:volumes:attach",
        "evs:volumes:detach",
        "evs:volumes:manage",
        "evs:volumes:update",
        "evs:volumes:uploadImage",
        "evs:snapshots:create",
        "vpc:*:get",
        "vpc:*:list",
        "vpc:networks:create",
        "vpc:networks:update",
        "vpc:subnets:update",
        "vpc:subnets:create",
        "vpc:routers:get",
        "vpc:routers:update",
        "vpc:ports:*",
        "vpc:privatelps:*",
        "vpc:securityGroups:*",
        "vpc:securityGroupRules:*",
        "vpc:floatingIps:*",
        "vpc:publicIps:*",
        "vpc:bandwidths:*",
        "vpc:firewalls:*",
        "ims:images:create",
        "ims:images:delete",
        "ims:images:get",
        "ims:images:list",
        "ims:images:update",
        "ims:images:upload"
      ],
      "Effect": "Allow"
    }
  ]
}
```

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.

Selecting a Region

If your target users are in Europe, select the **EU-Dublin** region.

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 lower network latency, deploy resources in the same AZ.