Elastic Volume Service

Service Overview

Issue 01

Date 2025-09-19





Copyright © Huawei Technologies Co., Ltd. 2025. 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 Technologies Co., Ltd.

Trademarks and Permissions

HUAWEI and other Huawei trademarks are trademarks 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 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.

Security Declaration

Vulnerability

Huawei's regulations on product vulnerability management are subject to the *Vul. Response Process.* For details about this process, visit the following web page:

https://www.huawei.com/en/psirt/vul-response-process

For vulnerability information, enterprise customers can visit the following web page:

https://securitybulletin.huawei.com/enterprise/en/security-advisory

Contents

1 What Is EVS?	1
2 Disk Types and Performance	6
3 Device Types	13
4 Billing	15
4.1 Billing for EVS Disks	
4.2 Impacts and Usage Suggestions on Yearly/Monthly Disks Before and After Expiration	17
4.3 Impacts and Usage Suggestions on Pay-per-Use Disks Before and After Account Arrears	20
5 Permissions	22
6 Notes and Constraints	25
7 EVS and Other Services	32
8 Basic Concepts	34
8.1 EVS Concepts	34
8.2 Region and AZ	35
8.3 Three-Copy Redundancy	
8.4 Disk Sharing	38
8.5 Disk Encryption	30

1 What Is EVS?

Overview

Elastic Volume Service (EVS) offers scalable block storage for cloud servers. EVS disks provide high reliability, high performance, and come with a variety of disk types. They can be used for distributed file systems, development and test environments, data warehouses, and high-performance computing (HPC) applications. Cloud servers that EVS supports include Elastic Cloud Servers (ECSs) and Bare Metal Servers (BMSs).

Just like the physical disks in local PC need to be installed before they can be used, EVS disks need to be attached to servers before they can be used. They cannot be used alone. You also need to partition and create file systems on them before they can be used for persistent data storage.

In this document, EVS disks are sometimes just referred to as "disks".

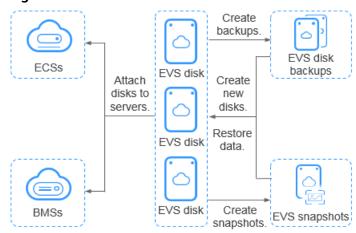


Figure 1-1 EVS architecture

EVS Advantages

EVS has the following advantages.

Table 1-1 EVS advantages

Advantage	Description	Reference
Various disk types	EVS provides a variety of disk types for you to choose from. They can be used as data disks or system disks for cloud servers. You can select whichever disk type that has the specifications best suited to your budget and service requirements.	Disk Types and Performance
Elastic scalability	The EVS disk capacity ranges from 10 GiB to 32 TiB. You can start with 10 GiB, and if, later on, that no longer meets your needs, you can expand the disk capacity to up to 32 TiB in increments of 1 GiB, without interrupting your applications.	Expansion Overview
	In addition to the disk capacity limit, there is an EVS capacity quota. The additional space you add cannot exceed the remaining quota. However, if this happens, you can apply for a higher quota.	Querying EVS Resource Quotas
High security and reliability	Both system disks and data disks support data encryption to ensure data security.	Managing Encrypted EVS Disks
High security and reliability	Data protection functions, such as backups, safeguard the disk data. If your data is ever damaged by a software exception or online attack, you can restore your data from backups.	Managing EVS Disk Backups
High security and reliability	Data protection functions, such as snapshots, safeguard the disk data. If your data is ever damaged by a software exception or online attack, you can restore your data from snapshots.	Managing EVS Snapshots
Real-time monitoring	On Cloud Eye, you can monitor the disk health and operating status at any time.	Viewing EVS Monitoring Data

Differences Among EVS, SFS, and OBS

There are three types of storage available for you to choose from: EVS, Scalable File Service (SFS), and Object Storage Service (OBS). Their differences are described in the following table.

Table 1-2 Comparison among EVS, SFS, and OBS

Dimensio n	SFS	OBS	EVS
Concept	SFS provides on- demand high- performance file storage that can be shared by multiple servers. SFS can be used like a remote directory for Windows or Linux servers.	OBS provides massive, secure, reliable, and cost-effective data storage for you to store data of any type and size.	EVS provides scalable, high-performance, high-reliability, block storage that can be used to meet a wide variety of service requirements. EVS disks are like physical disks on PCs.
Data storage logic	Stores files. Data is sorted and displayed in files and folders.	Stores objects. Files are saved directly to OBS. The files automatically generate corresponding system metadata. You can also customize the metadata if needed.	Stores binary data and cannot directly store files. To store files, you need to format the disk with a file system first.
Access method	SFS file systems need to be mounted to cloud servers through the Network File System (NFS) or Common Internet File System (CIFS) protocol before they can be accessed. A network address must be specified or mapped to a local directory for access.	OBS buckets can be accessed through the Internet or Direct Connect. The bucket address must be specified for access, and transfer protocols HTTP and HTTPS are used.	EVS disks can only be used and accessed from applications after being attached to cloud servers and initialized.

Dimensio n	SFS	OBS	EVS
Applicatio n scenarios	High-performance computing, media processing, file sharing, content management, and web services NOTE Mainly suitable for high-performance computing workloads like gene sequencing and image rendering that require high bandwidth for file sharing.	Big data analytics, static website hosting, online video on demand (VoD), gene sequencing, and intelligent video surveillance	High-performance computing, enterprise critical clustered applications, enterprise application systems, and development and testing NOTE Mainly suitable for high-performance workloads like industrial design and energy exploration that require high speed and high IOPS for high-performance storage.
Capacity	PiB-level	EiB-level	TiB-level
Latency	3–10 ms	Milliseconds	1 ms
IOPS/TPS	10,000 per file system	Tens of millions	33,000 per disk
Bandwidth	GiB/s	TiB/s	MiB/s
Data sharing	Supported	Supported	Supported
Remote access	Supported	Supported	Not supported
Used independe ntly	Supported	Supported	Not supported

Access Methods

EVS provides a web-based console and HTTPS-based APIs that you can use to access the EVS service.

APIs

Use APIs if you need to integrate EVS into a third-party system for secondary development. For details, see **Elastic Volume Service API Reference**.

Console

Use the console if you do not need to integrate EVS with a third-party system. Sign in to the console with your account and choose **Elastic Volume Service**

from the service list. If you do not have an account, **sign up with Huawei Cloud**.

2 Disk Types and Performance

EVS disks are classified based on the disk I/O performance. EVS disks differ in performance and price. You can choose whichever disk type that is the best fit for your applications.

EVS Performance

EVS performance metrics include:

- IOPS: The number of read/write operations performed by an EVS disk per second.
- Throughput: The amount of data read from and written into an EVS disk per second.
- Read/write I/O latency: The minimum interval between two consecutive read/write operations on an EVS disk.

■ NOTE

- The performance of an EVS disk is not only restricted by its specifications. It is also limited by the storage QoS of the instance that disk is attached to. For details, see Instance QoS.
- The max. IOPS, max. throughput, and burst IOPS limit provided in **Table 2-1** are subject to factors like instance specifications, workload, network environment, and data block size. They have no Service Level Agreement (SLA) commitments. You are advised to test disk performance based on service scenarios.
- You can submit a service ticket to increase the maximum capacity of a data disk to 65,536 GiB.

Table 2-1 EVS performance data

Paramet er	Extreme SSD	Ultra-high I/O	General Purpose SSD	High I/O
Max. capacity	• System disk: 1,024	• System disk: 1,024	• System disk: 1,024	• System disk: 1,024
(GiB)	• Data disk: 32,768	• Data disk: 32,768	• Data disk: 32,768	• Data disk: 32,768

Paramet er	Extreme SSD	Ultra-high I/O	General Purpose SSD	High I/O
Short descripti on	Superfast disks for workloads demanding ultra-high bandwidth and ultra-low latency	High- performance disks excellent for enterprise mission-critical services as well as workloads demanding high throughput and low latency	Cost-effective disks designed for enterprise applications with medium performance requirements	Disks suitable for non- frequently accessed, latency- insensitive workloads ^f
Typical use cases	 Databases Oracle SQL Server ClickHou se AI workloads 	 Transcoding services I/O-intensive workloads NoSQL Oracle SQL Server PostgreSQ L Latency-sensitive applications Redis Memcach e 	 Enterprise OA Medium- scale developmen t and test environmen ts Small- and medium- sized databases Web applications System disks 	Common development and test environments
Max. IOPS ^a	128,000	50,000	20,000	5,000
Max. throughp ut ^a (MiB/s)	1,000	350	250	150
Burst IOPS limit ^a	64,000	16,000	8,000	5,000

Paramet er	Extreme SSD	Ultra-high I/O	General Purpose SSD	High I/O
Disk IOPS calculati on formula ^c . Capacity unit in the formula: GiB	IOPS limit = Min. (128,000, 1,800 + 50 × Capacity)	IOPS limit = Min. (50,000, 1,800 + 50 × Capacity)	IOPS limit = Min. (20,000, 1,800 + 12 × Capacity)	IOPS limit = Min. (5,000, 1,800 + 8 × Capacity)
Disk throughp ut calculati on formula ^b (MiB/s). Capacity unit in the formula: GiB	Throughput limit = Min. (1,000, 120 + 0.5 × Capacity)	Throughput limit = Min. (350, 120 + 0.5 × Capacity)	Throughput limit = Min. (250, 100 + 0.5 × Capacity)	Throughput limit = Min. (150, 100 + 0.15 × Capacity)
Single- queue access latency ^d	Sub- milliseconds	1 ms	1 ms	1–3 ms
API name ^e	ESSD	SSD	GPSSD	SAS

□ NOTE

a: The maximum IOPS, maximum throughput, and burst IOPS limit all include both read and write operations. So, maximum IOPS = read IOPS + write IOPS.

b: Take ultra-high I/O for example: The baseline throughput is 120 MiB/s. The throughput increases by 0.5 MiB/s for every one GiB added until it reaches the maximum throughput 350 MiB/s. Min in the formula indicates that the smaller value between the two values is used.

c: Take ultra-high I/O for example: The baseline IOPS is 1,800. The IOPS increases by 50 for every one GiB added until it reaches the maximum IOPS 50,000. Min in the formula indicates that the smaller value between the two values is used.

d: A single queue indicates that the queue depth or concurrency is 1. The single-queue access latency is the I/O latency when all I/O requests are processed sequentially. The values in the table are calculated with 4 KiB data blocks.

e: This API name is the value of the **volume_type** parameter in the EVS API. It does not represent the type of the underlying hardware device.

f: High I/O disks (except for those created in dedicated storage pools) are HDD-backed disks. They are suitable for applications with commonly accessed, latency-insensitive workloads. The baseline throughput of a High I/O disk is 40 MiB/s per TiB, and the maximum throughput of a High I/O disk is 150 MiB/s. If your applications have high workloads and are latency-sensitive, it is recommended that you choose SSD-backed disks which have higher specifications.

EVS disk performance is closely related with the data block size:

- If data blocks are all the same size, a disk can achieve either the maximum IOPS or maximum throughput depending on which one is reached first.
- If data blocks are of different sizes, the maximum performance metric that a disk can achieve varies:
 - For small data blocks, such as 4 KiB or 8 KiB, a disk can reach the maximum IOPS.
 - For data blocks of a large size, 16 KiB or greater, a disk can reach the maximum throughput.

Table 2-2 uses an ultra-high I/O disk as an example. In theory, when the size of an ultra-high I/O disk is at least 964 GiB, the disk theoretically can reach either the maximum IOPS 50,000 or the maximum throughput 350 MiB/s. However, this is not the case in practice. The maximum IOPS and maximum throughput that a disk can reach also vary with the data block size.

Table 2-2 Maximum performance of an ultra-high I/O EVS disk

Data Block Size (KiB)	Max. IOPS	Max. Throughput (MiB/s)
4	About 50,000	About 195
8	About 44,800	About 350
16	About 22,400	About 350
32	About 11,200	About 350

Disk IOPS Calculation Formula

Disk IOPS limit = Min. (Maximum IOPS, Baseline IOPS + IOPS per GiB × Capacity)

Take an ultra-high I/O EVS disk with a maximum IOPS of 50,000 for example.

- If the disk capacity is 100 GiB, the disk IOPS limit is calculated as follows: Disk IOPS limit = Min. (50,000, 1,800 + 50 × 100)
 - The disk IOPS limit is 6,800, the smaller of the two values (50,000 and 6,800).
- If the disk capacity is 1,000 GiB, the disk IOPS limit is calculated as follows: Disk IOPS limit = Min. $(50,000, 1,800 + 50 \times 1,000)$
 - The disk IOPS limit is 50,000, the smaller of the two values (50,000 and 51,800).

Disk Burst Capability and Principles

EVS disks have a burst capability. A small-capacity disk can surpass its official maximum IOPS for a short period of time. This IOPS applies to each disk individually.

Disks with burst capability are well-suited for speeding up server startup. In most cases, system disks are fairly small, so their basic IOPS is fairly low. For example, the IOPS of a 50-GiB ultra-high I/O disk without burst can only reach up to 4,300 IOPS (Min. $(50,000, 1,800 + 50 \times \text{Capacity})$). But with burst capability, its IOPS can burst up to 16,000.

The following example uses an ultra-high I/O EVS disk with the IOPS burst limit of 16.000.

- If the disk capacity is 100 GiB, the disk has a maximum IOPS of 6,800, but it can temporarily burst to 16,000 IOPS.
- If the disk capacity is 1,000 GiB, the disk has a maximum IOPS of 50,000. The
 disk maximum IOPS already exceeds its burst IOPS 16,000, and the disk does
 not use the burst capability.

The following describes the burst IOPS consumption and reservation.

A token bucket is used to handle burst I/O operations. The number of initial tokens in the bucket is calculated as follows:

Number of initial tokens = Burst duration × IOPS burst limit

In the following example, a 100-GiB ultra-high I/O EVS disk is used, and the fixed burst duration is 1800 seconds. Therefore, the number of initial tokens is $28,800,000 \ (1,800 \times 16,000)$.

- Token production rate: This rate equals the disk maximum IOPS, which is 6,800 tokens/s.
- Token consumption rate: This rate is based on the I/O usage. Each I/O request consumes a token. The maximum consumption rate is 16,000 tokens/s, which is the larger value of the disk burst IOPS and the maximum IOPS.

Consumption principles

When tokens are consumed faster than they are produced, the number of tokens decreases accordingly, and eventually the disk IOPS will be consistent with the token production rate (the maximum IOPS). In this example, the disk can burst for approximately 3,130 seconds (28,800,000/(16,000 - 6,800)).

Reservation principles

When tokens are consumed more slowly than they are produced, the number of tokens increases accordingly, and the disk regains burst capability. In this example, if the disk is suspended for approximately 4,235 seconds (28,800,000/6,800), the token bucket will be filled up with tokens.

NOTE

As long as there are tokens in the token bucket, the disk has the burst capability.

Figure 2-1 shows the token consumption and reservation principles. The blue bars indicate the disk IOPS usage, the green dashed line represents the maximum IOPS, the red dashed line indicates the IOPS burst limit, and the black curve indicates the changes of the number of tokens.

- As long as there are tokens, the disk IOPS can exceed 6,800 and can burst up to 16,000, the IOPS burst limit.
- When there are no more tokens, the disk loses the burst capability, and the disk IOPS can reach up to 6,800.
- Anytime the disk IOPS is less than 6,800, the number of tokens starts to increase, and the disk regains the burst capability.

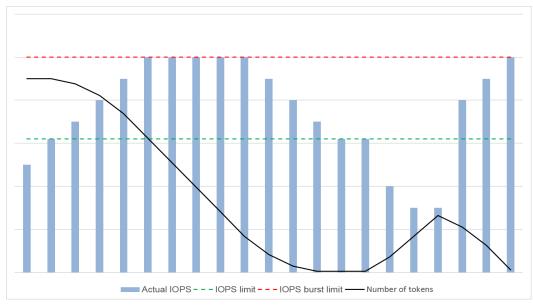
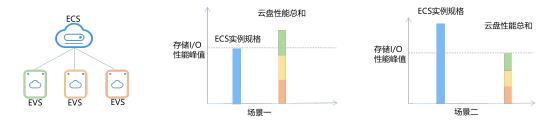


Figure 2-1 Burst capability diagram

Instance QoS

ECSs support instance-level storage I/O isolation. Each instance (ECS) enjoys their storage bandwidth and IOPS exclusively. This prevents resource contention between the ECSs during peak hours.



Scenario 1: If the total performance of EVS disks on an instance exceeds the storage I/O capabilities of the instance flavor, the instance flavor's storage I/O capabilities are used. If the instance has multiple EVS disks, the storage I/O performance is dynamically allocated to the EVS disks based on a contention mechanism.

Example 1: A customer purchases an ECS that offers up to 2,688 MiB/s (21 Gbit/s) storage bandwidth. The ECS has ten 1,000 GiB Ultra-high I/O EVS disks, each of which offers a maximum bandwidth of 350 MiB/s. So, the total maximum bandwidth of the ten disks is 3,500 MiB/s, which already exceeds the maximum storage bandwidth of the instance (2,688 MiB/s). In this case, the 3,500 MiB/s bandwidth cannot be reached.

Example 2: A customer purchases an ECS that offers up to 130,000 IOPS. The ECS has three 8 TiB Ultra-high I/O EVS disks, each of which offers a maximum IOPS of 50,000. So, the total maximum IOPS of the three disks is 150,000, which already exceeds the maximum IOPS of the instance (130,000). In this case, the maximum storage IOPS of the disks can only reach 130,000, not 150,000.

Scenario 2: If the total performance of EVS disks on an instance is less than the storage I/O capabilities of the instance flavor, the disks' storage performance is used.

Example: A customer purchases an ECS that offers up to 2,688 MiB/s (21 Gbit/s) storage bandwidth. The ECS has three 100 GiB Ultra-high I/O EVS disks, each of which offers a maximum bandwidth of 170 MiB/s. So, the total maximum bandwidth of the three disks is 510 MiB/s, which is less than the maximum storage bandwidth of the instance (2,688 MiB/s). In this case, the 2,688 MiB/s bandwidth cannot be reached.

Performance Testing

For details about how to test the EVS disk performance, see **How Do I Test My Disk Performance?**

3 Device Types

What Device Types Are Available?

There are two EVS device types: Virtual Block Device (VBD) and Small Computer System Interface (SCSI).

- VBD is the default EVS device type. VBD EVS disks support only basic read/ write SCSI commands.
- SCSI EVS disks support transparent SCSI command transmission and allow the server OS to directly access the underlying storage media. Besides basic read/ write SCSI commands, SCSI disks support advanced SCSI commands.

Device type is configured during purchase. It cannot be changed after the disk has been purchased.

Common Application Scenarios and Usage Instructions of SCSI EVS Disks

- BMSs support only SCSI EVS disks.
- Shared SCSI EVS disks: Shared SCSI EVS disks must be used together with a
 distributed file system or cluster software. Because most cluster applications,
 such as Windows MSCS, Veritas VCS, and Veritas CFS, require SCSI
 reservations, you are advised to use shared EVS disks with SCSI.

SCSI reservations take effect only when shared SCSI EVS disks are attached to ECSs in the same ECS group.

Do I Need to Install a Driver for SCSI EVS Disks?

To use SCSI EVS disks, a cloud server must have a SCSI driver installed. If the SCSI driver is not pre-installed, you need to install it manually.

Check whether you need to manually install the driver based on the server type.

- Bare Metal Server (BMS)
 Both the Windows and Linux images for BMSs are pre-installed with the required SDI card driver. Therefore, no driver needs to be installed.
- KVM ECS

You are advised to use SCSI EVS disks with KVM ECSs. Linux images and Windows images for KVM ECSs already have the required driver. Therefore, no driver needs to be installed for KVM ECSs.

ECS virtualization types are categorized into KVM and Xen. For details, see ECS Types.

Xen ECS

Due to driver limitations, you are advised not to use SCSI EVS disk with Xen ECSs.

However, a few Windows and Linux images support SCSI EVS disks on Xen ECSs. For the supported images, see **Table 3-1**.

□ NOTE

After confirming that the OS images of Xen ECSs support SCSI EVS disks, determine whether you need to install the driver:

- Public Windows images are preinstalled with the Paravirtual SCSI (PVSCSI) driver.
 Therefore, no driver needs to be installed.
- Private Windows images are not preinstalled with the PVSCSI driver. You need to download and install it explicitly.
 - For details, see section "(Optional) Optimizing Windows Private Images" in the *Image Management Service User Guide*.
- Linux images are not preinstalled with the PVSCSI driver. You need to obtain the source code of the open-source Linux driver at https://github.com/UVP-Tools/SAP-HANA-Tools, compile the code, and then install the driver.

Table 3-1 OSs supporting SCSI EVS disks

Virtualizatio n Type	os	
Xen	Windows	See the Windows images listed on the Public Images page. Sign in to the console, choose Image
		Management Service, click the Public Images tab, and select ECS image and Windows from the drop-down lists, respectively.
	Linux	SUSE Linux Enterprise Server 11 SP4 64bit (The kernel version is 3.0.101-68-default or 3.0.101-80-default.)
		SUSE Linux Enterprise Server 12 64bit (The kernel version is 3.12.51-52.31-default.)
		SUSE Linux Enterprise Server 12 SP1 64bit (The kernel version is 3.12.67-60.64.24-default.)
		• SUSE Linux Enterprise Server 12 SP2 64bit (The kernel version is 4.4.74-92.35.1-default.)

4 Billing

4.1 Billing for EVS Disks

Billing Items

EVS disks are billed based on the disk type, size, and usage duration.

- Billing starts: You will be billed for the EVS disks right after you have purchased them, regardless of whether they are attached or not.
- Billing stops:
 - For a yearly/monthly disk, the billing stops after the disk is successfully unsubscribed from, and the refund is calculated as follows: Refund = Your actual payment Amount due Handling fees. For more information, see How Do I View the Refund for My Resource Unsubscription?
 - For a pay-per-use disk, the billing stops after the disk is successfully deleted.

□ NOTE

Billing items marked with an asterisk (*) are mandatory.

Billing Modes

EVS disks can be billed on a yearly/monthly or pay-per-use basis.

- Yearly/Monthly is a prepaid payment method.
- Pay-per-use is a postpaid payment model. Your disk is billed by the second, and you are billed for a minimum of 60 minutes each time. If the usage is less than an hour, you are billed based on the actual usage period consumed.

Billing Involved in Configuration Modifications

Item	Yearly/Monthly	Pay-per-Use
Capacity change	 EVS does not support the reduction of disk capacities. EVS supports the expansion of disk capacities. Additional capacities need to be paid. NOTE The expiration time of an EVS disk remains unchanged after the capacity expansion. 	 EVS does not support the reduction of disk capacities. EVS supports the expansion of disk capacities. Multiple pieces of billing records will be generated within a billing cycle (an hour) when an expansion succeeded. For example, if you expand the capacity of an EVS disk from 100 GiB to 200 GiB at 01:30:01, two billing records will be generated in the billing cycle from 01:00:00 to 02:00:00. One is the billing record generated for the 100 GiB from 01:00:00 to 01:30:00, and the other is the billing record generated for the 200 GiB from 01:30:01 to 02:00:00.
Billing mode change	EVS supports the billing mode change from pay-per-use to yearly/monthly. For details, see From Pay-per-Use to Yearly/Monthly. NOTE Non-shared, pay-per-use EVS disks cannot be changed to yearly/monthly billing separately. You need to attach them to servers and then change the billing mode: If you attach them to pay-per-use server billing mode to yearly/monthly, the disks will have the same expiration times as their servers. If you attach them to yearly/monthly servers and change the disk billing mode to yearly/monthly, you can specify the disk expiration times as required. You are advised to specify the expiration times the same as those of servers.	EVS supports the billing mode change from yearly/monthly to pay-per-use. For details, see From Yearly/Monthly to Pay-per-Use.

Expiration

Before a yearly/monthly disk expires, if you do not renew the disk or auto renewal is enabled but fails, the disk will enter the retention period after expiration. For details, see Impacts and Usage Suggestions on Yearly/Monthly Disks Before and After Expiration.

- During the retention period, if you renew the disk, the disk will be unfrozen.
- During the retention period, if you do not renew the disk, the disk will be released after the retention period ends.

Renewal

You can renew a yearly/monthly disk anytime before it expires or during the retention period.

You can renew your resources on the **Renewals** page of the console. For details, see **Renewal Management**.

Overdue Payment

If your account is not topped up after the account balance falls below zero, your account is in arrears and your pay-per-use disk will enter the retention period. For details, see Impacts and Usage Suggestions on Pay-per-Use Disks Before and After Account Arrears.

- During the retention period, if you top up your account, the disk will be unfrozen.
- During the retention period, if you do not top up your account, the disk will be released after the retention period ends.

4.2 Impacts and Usage Suggestions on Yearly/Monthly Disks Before and After Expiration

Introduction to Retention Period of Yearly/Monthly Resources

Yearly/Monthly is a prepaid billing mode, of which resource charges are paid in advance. You can choose yearly/monthly billing when purchasing disks.

Before a yearly/monthly disk expires, if you do not renew the disk or auto renewal is enabled but fails, the disk will enter the retention period after expiration.

- During the retention period, if you renew the disk, the disk will be unfrozen.
- During the retention period, if you do not renew the disk, the disk will be released after the retention period ends.

Ⅲ NOTE

The duration of the retention period varies depending on user levels. For more information, see **Resource Suspension and Release**.

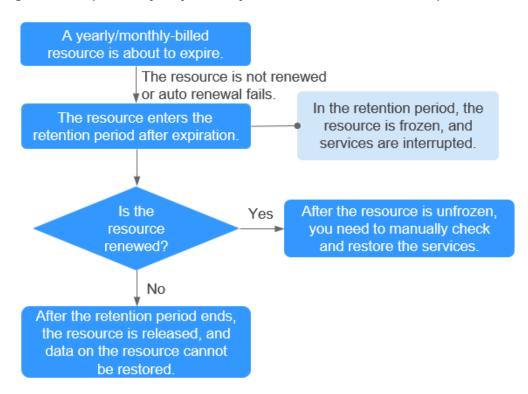


Figure 4-1 Impacts on yearly/monthly resources before and after expiration

Impact on Services When Resources Are Frozen, Unfrozen, or Released

- Frozen resources: Resource access and usage are restricted, which will interrupt your services. For example, if a server is frozen, it will be automatically powered off or shut down. If a disk is frozen, disk I/Os will be restricted.
- Unfrozen resources: Resource restrictions are removed, but you need to check and restore your services. For example, after a server is unfrozen, you need to power it on.
- Released resources: Resources are released. Data stored on the resources will be deleted and cannot be retrieved.

Usage Suggestions on Yearly/Monthly Resources

If you no longer need to use a yearly/monthly disk after it expires, you can sign in to the console, detach the disk, and release the resource. For details, see section "Releasing Resources" in the *Billing Center User Guide*.

Table 4-1 lists the common usage scenarios and suggestions on yearly/monthly disks. You can refer to usage suggestions to enable auto renewal and set a renewal date, and pay attention to resource expiration and freezing notifications to keep up with the latest resource information, ensuring that your services and data are not affected.

Table 4-1 Common usage scenarios and suggestions

Common Usage Scenario	Suggestions
Resources are billed in yearly/monthly mode.	 Manually renew the resources. For details, see Manually Renewing a Resource. Enable auto renewal and keep sufficient balance in your account. For details, see Enabling Auto-Renewal. Pay attention to notifications about auto renewal failures and top up your account in time. Pay attention to notifications about to-be-expired resources and renew the resources in time. Pay attention to notifications about to-be-frozen resources and renew the resources in time. Pay attention to notifications about to-be-released resources and renew the resources in time.
The server is billed in yearly/monthly mode, and the attached disks are also billed in yearly/monthly mode. The server expiration date is inconsistent with the disk expiration date.	 Set a renewal date. Renew the server and disks in a batch before the expiration date, and set the renewal date for these resources to a same date. For details, see Setting a Renewal Date. For details, see Manually Renewing a Resource. NOTE You can only set the renewal date to a day (from the 1st day to the 28th day of a month, or the last day of a month) but not to a month. If you want to set the renewal date to a whole month, you need to set a unified expiration month when setting the renewal duration. Refer to suggestions for the scenario where resources are billed in yearly/monthly mode.
The server is billed in yearly/monthly mode, but the attached disks are billed in pay-per-use mode.	 Change the disk billing mode from pay-per-use to yearly/monthly. For details, see From Pay-per-Use to Yearly/Monthly. Refer to suggestions for the scenario where resources are billed in yearly/monthly mode. If the disk billing mode is not changed, refer to suggestions for the scenario where resources are billed in pay-per-use mode.

Common Usage Scenario	Suggestions
The server is billed in pay-per-use mode, but the attached disks are billed in yearly/monthly	 Change the server billing mode from pay-per-use to yearly/monthly. For details, see From Pay-per-Use to Yearly/ Monthly.
mode.	 Refer to suggestions for the scenario where resources are billed in yearly/monthly mode.
	 If the server billing mode is not changed, refer to suggestions for the scenario where resources are billed in pay-per-use mode.
Resources are billed in pay-per-use mode.	 Top up your account in time to keep sufficient account balance.
	 Pay attention to notifications about insufficient balance alert and top up your account in time.
	 Pay attention to notifications about account arrears and top up your account in time.

4.3 Impacts and Usage Suggestions on Pay-per-Use Disks Before and After Account Arrears

Introduction to Retention Period of Pay-per-Use Resources

Pay-per-use is a postpaid billing mode, of which resource charges are deducted from the account balance based on the resource usage duration. You can choose pay-per-use billing when purchasing disks.

If you do not top up your account after the account balance falls below zero, your disk will enter the retention period instead of being released directly.

- During the retention period, if you top up your account, the disk will be unfrozen.
- During the retention period, if you do not top up your account, the disk will be released after the retention period ends.

◯ NOTE

The duration of the retention period varies depending on user levels. For more information, see **Resource Suspension and Release**.

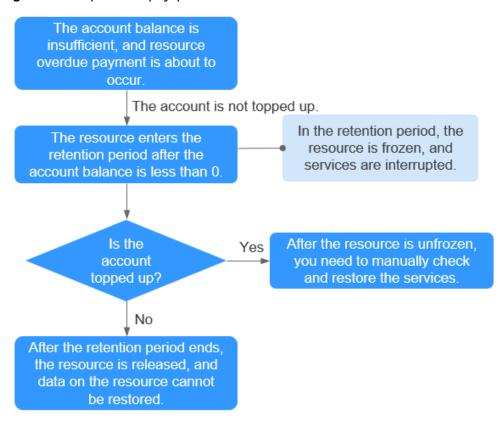


Figure 4-2 Impacts on pay-per-use resources before and after account arrears

Impact on Services When Resources Are Frozen, Unfrozen, or Released

- Frozen resources: Resource access and usage are restricted, which will interrupt your services. For example, if a server is frozen, it will be automatically powered off or shut down. If a disk is frozen, disk I/Os will be restricted.
- Unfrozen resources: Resource restrictions are removed, but you need to check and restore your services. For example, after a server is unfrozen, you need to power it on.
- Released resources: Resources are released. Data stored on the resources will be deleted and cannot be retrieved.

Usage Suggestions on Pay-per-Use Resources

If you no longer need to use a pay-per-use disk, you can sign in to the console, detach the disk, and then delete it. For how to delete a disk, see **Deleting an EVS Disk**.

5 Permissions

If you need to assign different permissions to personnel in your enterprise to access your EVS resources on Huawei Cloud, IAM is a good choice for fine-grained permissions management. IAM provides identity authentication, permissions management, and access control, helping you securely access your Huawei Cloud resources.

With IAM, you can control access to specific Huawei Cloud resources. For example, if you want some resource management personnel in your enterprise to view EVS resources but do not want them to delete EVS resources or perform any other high-risk operations, you can grant permission to view EVS resources but not permission to delete them.

If your Huawei Cloud account does not require IAM for permissions management, you can skip this section.

IAM is a free service. You only pay for the resources in your account.

For more information about IAM, see IAM Service Overview.

EVS Permissions

New IAM users do not have any permissions assigned by default. You need to first add them to one or more groups and attach policies or roles to these groups. The users then inherit permissions from the groups and can perform specified operations on cloud services based on the permissions they have been assigned.

EVS is a project-level service deployed for specific regions. To assign EVS permissions to a user group, specify the scope as region-specific projects and select a project for the permissions to take effect. If **All projects** is selected, the permissions will take effect for the user group in all region-specific projects. When accessing EVS, users need to switch to a region where they have been authorized to use EVS.

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 need to also assign other roles on which the
permissions depend to take effect. However, roles are not an ideal choice for
fine-grained authorization and secure access control.

Policies: A type of fine-grained authorization mechanism 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 EVS, see "Permissions and Supported Actions" > "Introduction"
in the Elastic Volume Service API Reference.

Table 5-1 lists all the system-defined roles and policies supported by EVS.

Table 5-1 System-defined roles and policies supported by EVS

Role/Policy Name	Description	Туре	Dependen cies
EVS FullAccess	Full permissions for EVS. Users granted these permissions can create, attach, detach, query, and delete EVS resources, and expand capacity of EVS disks.	System- defined policy	None
EVS ReadOnlyAcc ess	Read-only permissions for EVS. Users granted these permissions can view EVS resource data only.	System- defined policy	None
Server Administrato r	Full permissions for EVS	System- defined role	None

Table 5-2 lists the common operations supported by each system-defined policy of EVS. Select the policies as required.

Table 5-2 Common operations supported by each system-defined policy of EVS

Operation	EVS FullAccess	EVS ReadOnlyAccess
Creating disks	√	x
Viewing the disk list	√	√
Viewing disk details	√	√
Attaching disks	√	х
Detaching disks	√	х
Deleting disks	√	х
Expanding capacities of disks	√	х
Creating snapshots	√	х

Operation	EVS FullAccess	EVS ReadOnlyAccess
Deleting snapshots	√	x
Rolling back data from snapshots	√	x
Creating disks from snapshots	√	x
Adding tags for disks	√	х
Modifying disk tags	√	х
Deleting disk tags	√	х
Searching for disks by tag	√	√
Changing disk names	√	х

Related Links

- IAM Service Overview
- Creating a User and Granting EVS Permissions
- Permissions Policies and Supported Actions

6 Notes and Constraints

This section describes the constraints on using EVS.

Specifications

Table 6-1 Specifications

Resource Type	Item	Description
Disk capacity	Capacity of a system disk	40 GiB - 1024 GiB
Capacity of a data disk		10 GiB - 32768 GiB
	Maximum capacity supported by MBR	2 TiB
	Maximum capacity supported by GPT	18 EiB
Disk performance	Major disk performance metrics include IOPS, throughput, and latency.	Different types of EVS disks have different performance. For details, see Disk Types and Performance.

Security

Table 6-2 Security

Item	Description	
Disk encryption	The encryption attribute of a disk cannot be changed after the disk is purchased.	
	If you use an encrypted disk to create a backup, the backup generated will be an encrypted backup. You cannot modify the encryption attribute of the backup.	
	 If you use an encrypted disk to create an image, the image generated will be an encrypted image. You cannot modify the encryption attribute of the image. 	
	 If you use an encrypted disk to create a snapshot, the snapshot generated will be an encrypted snapshot. You cannot modify the encryption attribute of the snapshot. 	
	If you use an image that does not support lazyloading to create a disk, the disk created will be an encrypted disk. You cannot modify the encryption of the disk.	

Quotas

You can sign in to the console to view default quotas. You can **submit a service ticket** to apply for a larger quota if needed.

Table 6-3 Quotas

Resource Type	Default Quota
Tags	10
Disks	Default quotas vary depending on regions. See
Disk capacity (GB)	the quotas shown on the console.
Snapshots	

Operations

Table 6-4 Operations

Scenario	Item	Description	
Disk creation	Maximum number of disks that can be created at a time	100	
Disk creation	Disk creation from snapshot	For details, see the constraints in Creating a Disk from a Snapshot.	
Disk creation	Disk creation from backup	 Batch creation is not supported. One can create only one disk from a backup at a time. One backup cannot be used for concurrent disk creation operations at the same time. For example, if you are creating disk A from a backup, this backup can only be used to create another disk after disk A has been created. If a disk is created from a backup of a system disk, the new disk can be used as a data disk only. 	
Disk creation	Disk creation from image	 The device type of the new disk is the same as that of the image's source disk. The encryption attribute of the new disk is the same as that of the image's source disk. 	
Disk creation	Device type	The device type of a disk cannot be changed after the disk is created.	
Disk creation	Disk sharing	The sharing attribute of a disk cannot be changed after the disk is created.	
Disk creation	Disk encryption	The encryption attribute of a disk cannot be changed after the disk is created.	
Disk attachment	Constraints on region and AZ	The disk and the server must be in the same region and AZ.	
	Maximum number of servers that a non- shared disk can be attached to	1	

Scenario	Item	Description	
	Maximum number of servers that a shared disk can be attached to	16	
	Maximum number of disks that can be attached to an ECS	This number varies depending on the ECS type. For details, see Can I Attach Multiple Disks to an ECS? For a BMS, a maximum of 60 EVS disks (1 system disk and 59 data disks) are supported. Only SCSI disks can be attached.	
	Device name	 System disk: /dev/vda, /dev/sda, and /dev/xvda Data disk: /dev/vd[b-z], /dev/sd[b-z], and /dev/xvd[b-z] 	
Disk capacity expansion	Capacity expansion	Disk capacity can be expanded, but cannot be reduced.	
	Capacity expansion of non-shared disks	Some server OSs support the capacity expansion of non-shared, In-use disks. For details, see Expand Disk Capacity.	
	Capacity expansion of shared disks	A shared disk must be detached from all its servers before expansion. This means that the shared disk status must be Available .	
	Expansion increment	1 GiB	
Disk detachment	System disk detachment	A system disk can only be detached offline, which means that the server must be in the Stopped state.	
	Data disk detachment	A data disk can be detached online or offline, meaning that its server can either be in the Running or Stopped state.	
Disk deletion	Deletion of pay-per- use disks	 The disk status is Available, Error, Expansion failed, Restoration failed, or Rollback failed. The disk is not locked by any service. The shared disk has been detached from all its servers. 	
Snapshot creation	/	For details, see the constraints in Creating an EVS Snapshot.	

Scenario	Item	Description
Snapshot data rollback to disks		 Snapshot data can only be rolled back to source EVS disks. Rollback to a different disk is not possible. If the snapshot status is Creating, it cannot be used to roll back disk data. A snapshot whose name starts with autobk_snapshot_vbs_, manualbk_snapshot_vbs_, autobk_snapshot_csbs_, or manualbk_snapshot_csbs_ is automatically generated during backup. Such a snapshot can only be viewed. It cannot be used to roll back the disk data.
Snapshot deletion	/	For details, see the constraints in Deleting an EVS Snapshot .
Disk type change	Before the change	 You can only change the disk type when the disk status is Available or In-use. The disk type cannot be changed when any snapshot of the disk is being deleted. A disk having more than 128 snapshots cannot have its disk type changed. You can delete some snapshots and then perform the change. In rare cases, the disk type may fail to be changed due to a background resource issue. If this happens, submit a service ticket. A disk protected by SDRS cannot have its disk type changed.

Scenario	Item	Description
	During the change	You can only change the disk type when the disk status is Available or In-use .
		 The disk type cannot be changed when any snapshot of the disk is being deleted.
		Some operations cannot be performed on the disk. Such operations include creating snapshots, creating backups, expanding the disk capacity, rolling back data from a snapshot, restoring data from a backup, attaching or detaching the disk, deleting the disk, transferring the disk, and creating an image from the ECS.
		 Changing the disk type may take several hours or even longer, and cannot be stopped. The time depends on the throughput, storage space, and original disk type at the time of the change.
		 You can have a maximum of 10 disks with their types being changed at the same time.
		The OS cannot be changed if you are changing the disk type of a system disk.
	After the change	In rare cases, the disk type may fail to be changed due to a background resource issue. If this happens, try again later.
		 In rare cases, the disk type may fail to be changed after a data rollback from a snapshot. If this happens, submit a service ticket.

Scenario	Item	Description
Recycle bin management	/	 When you delete a disk, regardless of whether the disk will be moved to the recycle bin or not, snapshots of the disk will always be deleted permanently.
		 There are no limits on the capacity and quantity of disks in the recycle bin.
		 You can recover or permanently delete the disks in the recycle bin. After the disks expire, they are permanently deleted and cannot be recovered.
		• If the ECS recycle bin is enabled, EVS disks deleted or unsubscribed from together with ECSs will be retained in the recycle bin the same period of time as the ECSs and can only be recovered or permanently deleted at the ECS side. After the disks expire, they are permanently deleted and cannot be recovered.

EVS and Other Services

Figure 7-1 shows the relationships between EVS and other services.

Figure 7-1 Relationships between EVS and other services

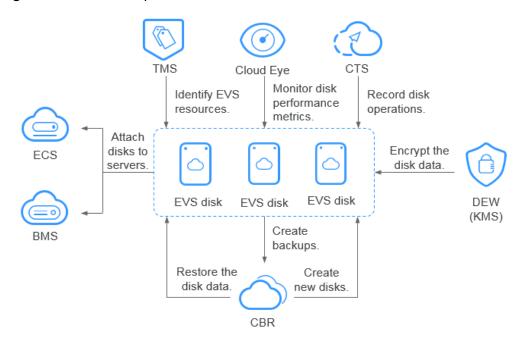


Table 7-1 EVS and other services

Interactive Function	Related Service	Reference
EVS disks can be attached to ECSs and used as scalable block storage devices.	ECS	Attaching a Non- Shared Disk Attaching a Shared Disk

Interactive Function	Related Service	Reference
SCSI EVS disks can be attached to BMSs and used as scalable block storage devices.	BMS	Attaching a Non- Shared Disk
		Attaching a Shared Disk
Backups can be created for EVS disks to guarantee the reliability and security of the server data.	CBR	Managing EVS Disk Backups
EVS disk encryption depends on the KMS service in DEW. You can use keys provided by KMS to encrypt EVS disks (both system and data disks), thus improving EVS disk data security.	DEW	Managing Encrypted EVS Disks
After EVS is enabled, the performance metrics of monitored disks can be viewed through Cloud Eye without any additional plug-in installed. The monitored metrics include Disk Read Rate, Disk Write Rate, Disk Read Requests, and Disk Write Requests.	Cloud Eye	Viewing EVS Monitoring Data
Cloud Trace Service (CTS) records operations of EVS resources, facilitating user query, audit, and backtracking.	CTS	Auditing
Tag Management Service (TMS) tags are used to identify EVS resources for purposes of easy categorization and quick search.	TMS	Adding a Tag

8 Basic Concepts

8.1 EVS Concepts

Table 8-1 EVS concepts

Concept	Description	Reference
IOPS	Number of read/write operations performed by an EVS disk per second.	Disk Types and Performance
Through put	Amount of data read from and written into an EVS disk per second.	
Read/ write I/O latency	Minimum interval between two consecutive read/write operations of an EVS disk.	
Burst capabilit y	A small-capacity disk can surpass its official maximum IOPS for a short period of time.	
VBD	A device type of EVS disks. VBD EVS disks only support basic SCSI read/write commands.	Device Types
SCSI	A device type of EVS disks. SCSI EVS disks support transparent SCSI command transmission and allow the server OS to directly access the underlying storage media.	

8.2 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.

8.3 Three-Copy Redundancy

What Is the Three-Copy Redundancy?

The backend storage system of EVS employs three-copy redundancy to guarantee data reliability. With this mechanism, one piece of data is by default divided into multiple 1 MiB data blocks. Each data block is saved in three copies, and these copies are stored on different nodes in the system according to the distributed algorithms.

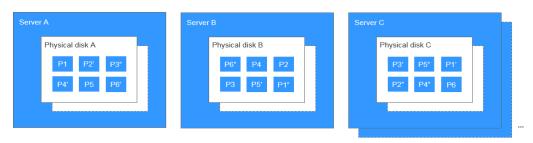
Three-copy redundancy has the following characteristics:

- The storage system saves the data copies on different disks of different servers across cabinets, ensuring that services are not interrupted if a physical device fails.
- The storage system guarantees strong consistency between the data copies.

For example, for data block P1 on physical disk A of server A, the storage system backs up its data to P1" on physical disk B of server B and to P1' on physical disk

C of server C. Data blocks P1, P1', and P1" are the three copies of the same data block. If physical disk A where P1 resides is faulty, P1' and P1" can continue providing storage services, ensuring service continuity.

Figure 8-1 Three-copy redundancy



How Does the Three-Copy Redundancy Keep Data Consistency?

When an application writes a piece of data to the system, the three copies of the data in the storage system must be consistent. When any of the three copies is read by the application later, the data on this copy is consistent with the data previously written to it.

Three-copy redundancy keeps data consistency in the following ways:

- Data is simultaneously written to the three copies of the data.
 When an application writes data, the storage system writes it to the three copies of the data simultaneously. In addition, the system returns the write success response to the application only after the data has been written to all of the three copies.
- Storage system automatically restores the damaged copy in the event of a data read failure.

When an application fails to read data, the system automatically identifies the failure cause. If the data cannot be read from a physical disk sector, the system reads the data from another copy of the data on another node and writes it back to the original disk sector. This ensures the correct number of data copies and data consistency among data copies.

How Does Three-Copy Redundancy Rapidly Rebuild Data?

Each physical disk in the storage system stores multiple data blocks, whose copies are scattered on the nodes in the system according to certain distribution rules. When a physical server or disk fault is detected, the storage system automatically rebuilds the data. Since the copies of data blocks are scattered on different nodes, the storage system will start the data rebuild on multiple nodes simultaneously during a data restore, with only a small amount of data on each node. In this way, the system eliminates the potential performance bottlenecks that may occur when a large amount of data needs to be rebuilt on a single node, and therefore minimizes the adverse impacts exerted on upper-layer applications.

Figure 8-2 shows the data rebuild process.

Figure 8-2 Data rebuild process

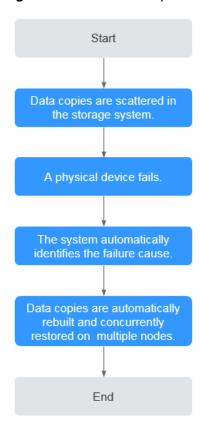


Figure 8-3 shows the data rebuild principle. For example, if physical disks on server F are faulty, the data blocks on these physical disks will be rebuilt on the physical disks of other servers.

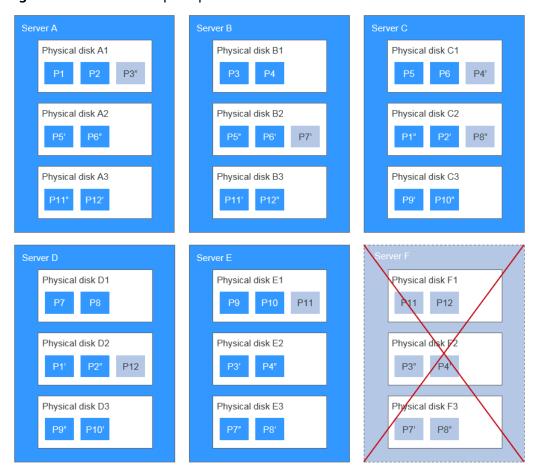


Figure 8-3 Data rebuild principle

What Are the Differences Between Three-Copy Redundancy, EVS Snapshots, and EVS Backups?

Three-copy redundancy improves the reliability of the data stored on EVS disks. It is used to tackle data loss or inconsistency caused by physical device faults.

EVS backups and EVS snapshots are used to prevent data loss or data inconsistency caused by incorrect operations, viruses, or hacker attacks. So you are advised to create backups or snapshots to back up the disk data on a timely basis.

8.4 Disk Sharing

Disk sharing allows you to create shared EVS disks. Shared EVS disks are block storage devices that support concurrent read/write operations and can be attached to multiple servers. Shared EVS disks provide multiple attachments, high concurrency, high performance, and high reliability. They are usually used for enterprise business-critical applications that require cluster deployment and high availability (HA). Multiple servers can access the same shared EVS disk at the same time.

A shared EVS disk can be attached to a maximum of 16 servers, including ECSs or BMSs. To share files, you need to deploy a shared file system or a cluster management system, such as Windows MSCS, Veritas VCS, or CFS.

NOTICE

A shared file system or cluster management system must be set up before you can properly use a shared disk. If you simply attach a shared disk to multiple servers, data cannot be shared between those servers and may be overwritten.

Cluster management system

ECS

ECS

Shared disk

ECS

ECS

Figure 8-4 Application scenario of shared EVS disks

8.5 Disk Encryption

EVS enables you to encrypt data on newly created disks as required.

It uses the industry-standard XTS-AES-256 cryptographic algorithm and keys to encrypt EVS disks. Keys used to encrypt EVS disks are provided by the Key Management Service (KMS) of Data Encryption Workshop (DEW), which is secure and convenient. You do not need to establish and maintain the key management infrastructure. KMS uses the Hardware Security Module (HSM) that complies with FIPS 140-2 level 3 requirements to protect keys. All user keys are protected by the root key in HSM to prevent key exposure.

NOTICE

The encryption attribute of a disk cannot be changed after the disk is purchased.