Elastic Volume Service(EVS)

Service Overview

Issue	01
Date	2023-11-01

HUAWEI TECHNOLOGIES CO., LTD.

Trademarks and Permissions


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?

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 specifications. 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

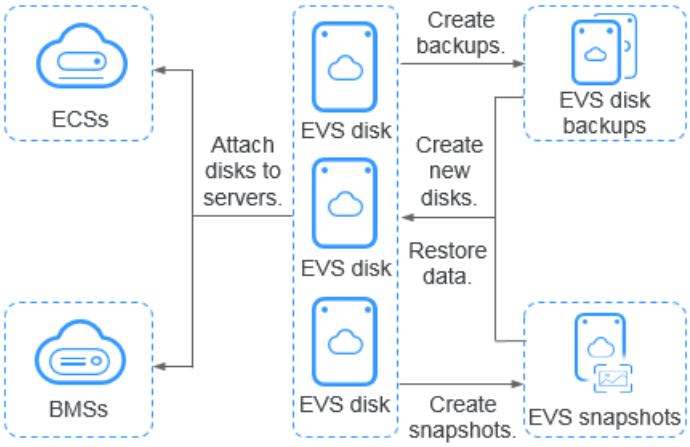

EVS Advantages

EVS has the following advantages.

Table 1-1 EVS advantages

Advantage	Description	Related Knowledge
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 SFS, OBS, and EVS

Dimension	SFS	OBS	EVS
Concept	SFS provides on-demand high-performance file storage, which 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 users 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 can be 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 ECSs or BMSs 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 ECSs or BMSs and initialized.

Dimension	SFS	OBS	EVS
Application 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	Sub-millisecond
IOPS/TPS	10,000 per file system	Tens of millions	128,000 per disk
Bandwidth	GiB/s	TiB/s	MiB/s
Data sharing	Supported	Supported	Supported
Remote access	Supported	Supported	Not supported
Used independently	Supported	Supported	Not supported

Access Methods

The public cloud system provides a web-based management 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**.

- Management console

  Use the management console if you do not need to integrate EVS with a third-party system. Log in to the management console with your account and

choose **Elastic Volume Service** from the service list. If you do not have an account, **register yourself on the public cloud**.

2 Disk Types and Performance

There are the following types of EVS disks, each with different levels of I/O performance: Extreme SSD, General Purpose SSD V2, Ultra-high I/O, General Purpose SSD, High I/O, and Common I/O. EVS disks differ in performance and price. You can choose whichever disk type that is the best fit for your applications.

Extreme SSD EVS disks use the congestion control algorithms for Remote Direct Memory Access (RDMA) deployments. An extreme SSD disk can reach up to 1,000 MiB/s of throughput and with extremely low single-channel latency.

EVS Performance

EVS performance metrics include:

- IOPS: number of read/write operations performed by an EVS disk per second
- Throughput: 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 on an EVS disk

Table 2-1 EVS performance data

Parameter	Extreme SSD	General Purpose SSD V2	Ultra-high I/O	General Purpose SSD	High I/O	Common I/O (Previous Generation Product)
Max. capacity (GiB)	• System disk: 1,024 • Data disk: 32,768	• System disk: 1,024 • Data disk: 32,768	• System disk: 1,024 • Data disk: 32,768	• System disk: 1,024 • Data disk: 32,768	• System disk: 1,024 • Data disk: 32,768	• System disk: 1,024 • Data disk: 32,768

Parameter	Extreme SSD	General Purpose SSD V2	Ultra-high I/O	General Purpose SSD	High I/O	Common I/O (Previous Generation Product)
Short description	Superfast disks for workloads demanding ultra-high bandwidth and ultra-low latency	SSD-backed disks allowing for tailored IOPS and throughput and targeting for transactional workloads that demand high performance and 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 commonly accessed workloads[f]	Disks suitable for less commonly accessed workloads

Parameter	Extreme SSD	General Purpose SSD V2	Ultra-high I/O	General Purpose SSD	High I/O	Common I/O (Previous Generation Product)
Typical workloads	• Database workloads – Oracle – SQL Server – ClickHouse • AI workloads	• Enterprise OA and virtual desktops • Large-scale development and test environments • Transcoding services • System disks • Medium- and large-sized databases (SQL Server, Oracle, NoSQL, and PostgreSQL)	• Transcoding services • I/O-intensive workloads – NoSQL – Oracle – SQL Server – PostgreSQL • Latency-sensitive applications – Redis – Memcache	• Enterprise OA • Medium-scale development and test environments • Small- and medium-sized databases • Web applications • System disks	Common development and test environments	Applications demanding large capacity, medium read/write speed, but having fewer transactions • Common office applications • Lightweight development and testing • Not recommended to be used as system disks
Max. IOPS[a]	128,000	128,000	50,000	20,000	5,000	2,200

Parameter	Extreme SSD	General Purpose SSD V2	Ultra-high I/O	General Purpose SSD	High I/O	Common I/O (Previous Generation Product)
Max. throughput[a] (MiB/s)	1,000	1,000	350	250	150	50
Burst IOPS limit[a]	64,000	N/A	16,000	8,000	5,000	2,200
Disk IOPS[c]	Min. [128,000, 1,800 + 50 x Capacity (GiB)]	You preconfigure an IOPS ranging from 3,000 to 128,000. This IOPS must also be less than or equal to 500 times the capacity (GiB).	Min. [50,000, 1,800 + 50 x Capacity (GiB)]	Min. [20,000, 1,800 + 12 x Capacity (GiB)]	Min. [5,000, 1,800 + 8 x Capacity (GiB)]	Min. [2,200, 500 + 2 x Capacity (GiB)]

Parameter	Extreme SSD	General Purpose SSD V2	Ultra-high I/O	General Purpose SSD	High I/O	Common I/O (Previous Generation Product)
Disk throughput[b] (MiB/s)	Min. [1,000, 120 + 0.5 × Capacity (GiB)]	You preconfigure a throughput ranging from 125 to 1,000. This throughput must also be less than or equal to the IOPS divided by 4.	Min. [350, 120 + 0.5 × Capacity (GiB)]	Min. [250, 100 + 0.5 × Capacity (GiB)]	Min. [150, 100 + 0.15 × Capacity (GiB)]	50
Single-queue access latency[d] (ms)	Sub-millisecond	1	1	1	1–3	5–10
API Name[e]	ESSD	GPSSD2	SSD	GPSSD	SAS	SATA

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.

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.

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 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, 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 greater than or equal to 16 KiB, 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 = Min. (Maximum IOPS, Baseline IOPS + IOPS per GiB x Capacity)

> **NOTE**
>
> This formula does not apply to General Purpose SSD V2 disks.
>
> For a General Purpose SSD V2 disk: You preconfigure an IOPS ranging from 3,000 to 128,000. This IOPS must also be less than or equal to 500 times the capacity (GiB).

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 is calculated as follows: Disk IOPS = Min. (50,000, 1,800 + 50 × 100)

  The disk IOPS 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 is calculated as follows: Disk IOPS = Min. (50,000, 1,800 + 50 × 1,000)

  The disk IOPS 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 x 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 burst to 16,000 IOPS in a certain duration.
- 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 x IOPS burst limit

In the following example, a 100-GiB ultra-high I/O EVS disk is used, and the fixed burst duration is 1800s. Therefore, the number of initial tokens is 28,800,000 (1,800 x 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.
- Any time 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

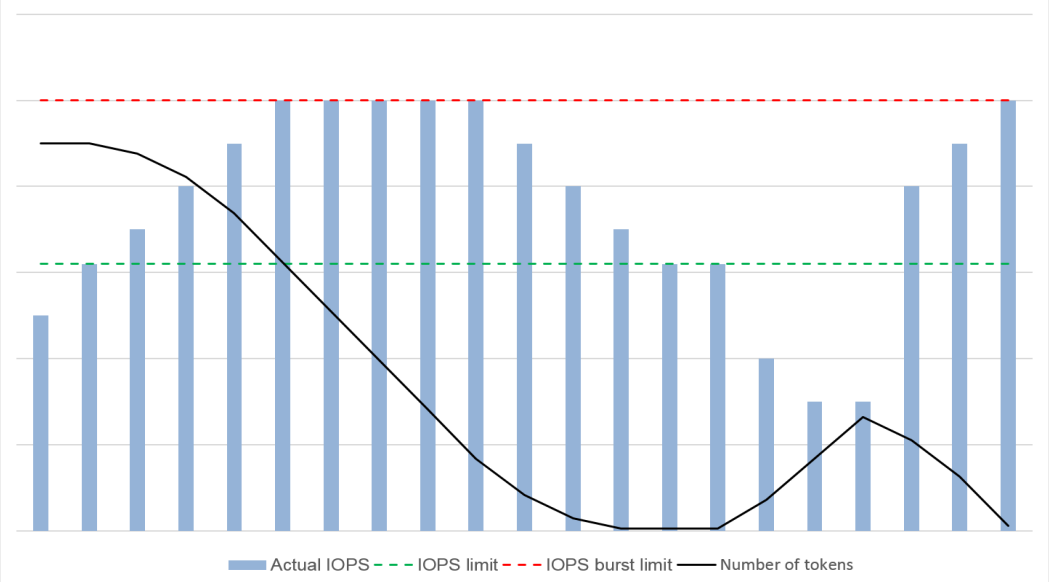

Performance Test Method

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

3 General Purpose SSD V2 Disks

General Purpose SSD V2 disks provide a baseline IOPS of 3,000 and a baseline throughput of 125 MiB/s regardless of the disk capacity.

With the General Purpose SSD V2 type, you can buy disks with the IOPS and throughput tailored to your workloads. The disk performance no longer changes with the disk capacity.

Performance

Table 3-1 EVS performance data

Parameter	General Purpose SSD V2
Max. capacity	• System disk: 1,024 GiB • Data disk: 32,768 GiB
Short description	General-purpose SSD-backed disks targeting for transactional workloads with optimized performance and cost
Typical scenarios	Mainstream high-performance, low-latency interactive applications • Enterprise OA and virtual desktops • Large-scale development and test environments • Transcoding services • System disks • Medium- and large-sized databases (SQL Server, Oracle, NoSQL, and PostgreSQL)
Max. IOPS	128,000
Max. throughput	1,000 MiB/s
Burst IOPS limit	N/A

Parameter	General Purpose SSD V2
Disk IOPS	You preconfigure an IOPS ranging from 3,000 to 128,000. This IOPS must also be less than or equal to 500 multiplying the capacity.
Disk throughput	You preconfigure a throughput ranging from 125 to 1,000 MiB/s. This throughput must also be less than or equal to the IOPS divided by 4.
Single-queue access latency	1 ms
API name **NOTE** 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.	GPSSD2

Configuration

Step 1 Go to the **Buy Disk** page.

Step 2 Configure the disk parameters.

- Choose the **General Purpose SSD V2** type and enter a desired disk size.
- Configure a desired IOPS.
- Configure a desired throughput.
- Configure other parameters by referring to **Purchasing an EVS Disk**.

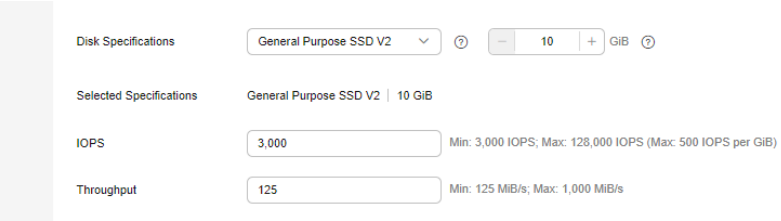

Step 3 Click **Next**.

- If you select **Yearly/Monthly** for **Billing Mode**:
  a. Check the disk details on the **Confirm** page.
  b. Confirm the information and click **Submit**.
  c. On the **Pay** page, select a desired payment method and confirm the payment. The system displays a message indicating payment processed successfully.
  d. Click **Back to Elastic Volume Service** to return to the **Elastic Volume Service** page.
- If you select **Pay-per-use** for **Billing Mode**:
  a. Check the disk details on the **Confirm** page.

b. Confirm the information and click **Submit**. The system displays a message indicating request submitted successfully.

c. Click **Back to Disk List** to return to the **Elastic Volume Service** page.

----End

NOTE

If you do not have a clear IOPS: throughput ratio in mind, you are advised to use the ratio of 50:1. For example, if your planned throughput is 600 MiB/s, configure 30,000 for the IOPS; if your planned throughput is 1,000 MiB/s, configure 50,000 for the IOPS.

If the preconfigured IOPS or throughput cannot meet your service requirement or is way more than what your need, you can adjust them at any time.

Billing

Table 3-2 Billing items

Billing Item	Billing Mode	Description
Capacity	Pay-per-use and yearly/ monthly	For more information, see **EVS Pricing Details**.
IOPS	Pay-per-use	
Throughput	Pay-per-use	

Billing Examples

Example 1: Purchasing a yearly/monthly General Purpose SSD V2 disk

A customer purchases a 100-GiB General Purpose SSD V2 disk preconfigured with an IOPS of 5,000 and throughput of 325 MiB/s.

If the yearly/monthly capacity unit price is $0.5 per GiB per month, the pay-per-use IOPS unit price is $0.0000153 per IOPS per hour, and the pay-per-use throughput unit price is $0.00194 per MiB/s per hour:

One month (30 days) after purchasing the disk, the customer is billed for $351.392 (Capacity charge + IOPS charge + throughput charge = 100 x 0.5 x 1+ (5,000 – 3,000) x 0.0000153 x 24 x 30 + (325 – 125) x 0.00194 x 24 x 30 = 50 + 22.032 + 279.36 = 351.392).

NOTE

- Capacity charge = Yearly/Monthly capacity charge
- IOPS charge = (Preconfigured IOPS – Baseline IOPS) x IOPS unit price x Duration
- Throughput charge = (Preconfigured throughput – Baseline throughput) x Throughput unit price x Duration

Example 2: Purchasing a pay-per-use General Purpose SSD V2 disk

A customer purchases a 100-GiB General Purpose SSD V2 disk preconfigured with an IOPS of 5,000 and throughput of 325 MiB/s.

If the pay-per-use capacity unit price is $0.000695 per GiB per hour, the pay-per-use IOPS unit price is $0.0000153 per IOPS per hour, and the pay-per-use throughput unit price is $0.00194 per MiB/s per hour:

24 hours after purchasing the disk, the customer is billed for $11.7344 (Capacity charge + IOPS charge + throughput charge = 100 x 0.000695 x 24 + (5,000 – 3,000) x 0.0000153 x 24 + (325 – 125) x 0.00194 x 24 = 1.668 + 0.7344 + 9.312 = 11.7344).

NOTE

- Capacity charge = Capacity x Storage unit price x Duration
- IOPS charge = (Preconfigured IOPS – Baseline IOPS) x IOPS unit price x Duration
- Throughput charge = (Preconfigured throughput – Baseline throughput) x Throughput unit price x Duration

4 EVS Features

4.1 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.

  NOTE

  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 4-1**.

  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 **(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 4-1 OSs supporting SCSI EVS disks

Virtualization Type	OS	
Xen	Windows	See the Windows images listed on the **Public Images** page. Log in to the management 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.)

4.2 Disk Sharing

What Is 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. Servers that EVS supports include ECSs and 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, the sharing function will not work and data may be overwritten.

Figure 4-1 Application scenario of shared EVS disks

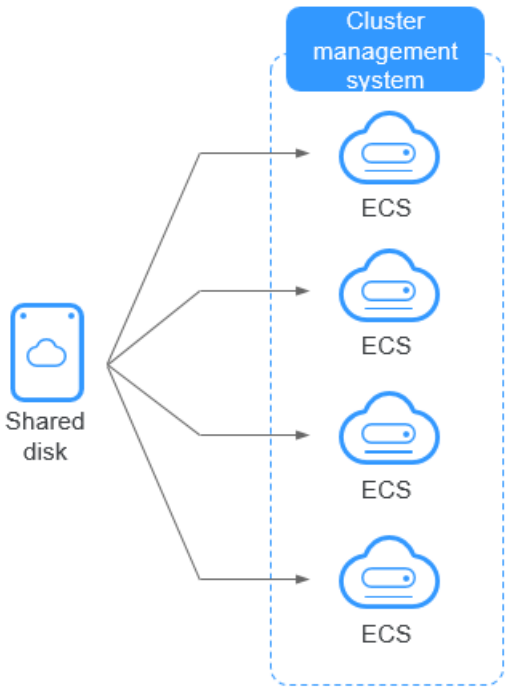

Usage Precautions

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. If a SCSI EVS disk is attached to a Xen ECS, you must install the driver. For details, see **Device Types**.

You can create shared VBD disks or shared SCSI disks. It is recommended that you attach a shared disk to the ECSs in the same ECS group to improve service reliability.

- Shared VBD disks: The device type of a newly created shared disk is VBD by default. Such disks can be used as virtual block storage devices, but do not support SCSI reservations. If SCSI reservations are required for your applications, create shared SCSI EVS disks.
- Shared SCSI disks: Such disks support SCSI reservations.

NOTICE

- To improve data security, you are advised to use SCSI reservations together with the anti-affinity policy of an ECS group. That said, ensure that shared SCSI disks are only attached to ECSs in the same anti-affinity ECS group.
- If an ECS does not belong to any anti-affinity ECS group, you are advised not to attach shared SCSI disks to this ECS. Otherwise, SCSI reservations may not work properly, which may put your data at risk.

Concepts of the anti-affinity ECS group and SCSI reservations:

- The anti-affinity policy of an ECS group allows ECSs to be created on different physical servers to improve service reliability.

  For details about ECS groups, see **Managing ECS Groups**.
- The SCSI reservation mechanism uses a SCSI reservation command to perform SCSI reservation operations. If an ECS sends such a command to an EVS disk, the disk is displayed as locked to other ECSs, preventing the data damage that may be caused by simultaneous reads/writes to the disk from multiple ECSs.
- ECS groups and SCSI reservations have the following relationship: A SCSI reservation on a single EVS disk cannot differentiate multiple ECSs on the same physical host. For that reason, if multiple ECSs that use the same shared EVS disk are running on the same physical host, SCSI reservations will not work properly. So you are advised to use SCSI reservations only on ECSs that are in the same ECS group, thus having a working anti-affinity policy.

Advantages

- Multiple attachments: A shared EVS disk can be attached to a maximum of 16 servers.
- High-performance: The random read/write IOPS of a shared ultra-high I/O disk can reach up to 160,000.
- High-reliability: Shared EVS disks support both manual and automatic backup, delivering highly reliable data storage.

- Wide range of use: Shared EVS disks can be used for Linux RHCS clusters where only shared VBD disks are needed. They can also be used for Windows MSCS and Veritas VCS clusters that require SCSI reservations.

Specifications and Performance

Shared EVS disks have the same specifications and performance as non-shared EVS disks.

Data Sharing Principles and Common Usage Mistakes

A shared EVS disk is essentially the disk that can be attached to multiple servers for use, which is similar to a physical disk in that the disk can be attached to multiple physical servers, and each server can read data from and write data into any space on the disk. If no data read/write rules, such as the read/write sequence and meaning, between these servers are defined, data reads and writes between these servers may conflict, or other unpredictable errors may occur.

Though shared EVS disks are block storage devices that provide shared access for servers, shared EVS disks do not have the cluster management capability. You need to deploy a cluster system to manage shared EVS disks. Common cluster management systems include Windows MSCS, Linux RHCS, Veritas VCS, and Veritas CFS.

If shared EVS disks are not managed by a cluster system, the following issues may occur:

- Data inconsistency caused by read/write conflicts

  When a shared EVS disk is attached to two servers (server A and server B), server A cannot recognize the disk spaces allocated to server B, vice versa. That said, a disk space allocated to server A may be already used by server B. In this case, repeated disk space allocation occurs, which leads to data errors.

  For example, a shared EVS disk has been formatted into an ext3 file system and attached to server A and server B. Server A has written metadata into the file system in space R and space G. Then server B has written metadata into space E and space G. In this case, the data written into space G by server A will be replaced. When the metadata in space G is read, an error will occur.

- Data inconsistency caused by data caching

  When a shared EVS disk is attached to two servers (server A and server B), the application on server A has read the data in space R and space G, then cached the data. At that time, other processes and threads on server A would then read this data directly from the cache. At the same time, if the application on server B has modified the data in space R and space G, the application on server A cannot detect this data change and still reads this data from the cache. As a result, the user cannot view the modified data on server A.

  For example, a shared EVS disk has been formatted into an ext3 file system and attached to server A and server B. Both servers have cached the metadata in the file system. Then server A has created a new file (file F) on the shared disk, but server B cannot detect this modification and still reads data from its cached data. As a result, the user cannot view file F on server B.

Before you attach a shared EVS disk to multiple servers, the disk device type needs to be determined. The device type can be either VBD or SCSI. Shared SCSI EVS

disks support SCSI reservations. Before using SCSI reservations, you need to install a driver in the server OS and ensure that the OS image is included in the compatibility list.

NOTICE

If you simply attach a shared disk to multiple servers, files cannot be shared between the servers, because the shared disk does not have the cluster capability. To share files between servers, build a shared file system or deploy a cluster management system.

4.3 Disk Encryption

What Is EVS 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 by encrypted EVS disks are provided by the Key Management Service (KMS) of Data Encryption Workshop (DEW), which is secure and convenient. So 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.

Keys Used for EVS Encryption

Keys provided by KMS include a Default Key and Custom Keys.

- Default Key: A key that is automatically created by EVS through KMS and named **evs/default**.

  It cannot be disabled and does not support scheduled deletion.
- Custom keys: Keys created by users. You can use existing keys or create new ones to encrypt disks. For details, see "Key Management Service" > "Creating a CMK" in the *Data Encryption Workshop User Guide*.

When an encrypted disk is attached, EVS accesses KMS, and KMS sends the data key (DK) to the host memory for use. The disk uses the DK plaintext to encrypt and decrypt disk I/Os. The DK plaintext is only stored in the memory of the host housing the ECS and is not stored persistently on the media. If a custom key is disabled or deleted in KMS, the disk encrypted using this custom key can still use the DK plaintext stored in the host memory. If this disk is later detached, the DK plaintext will be deleted from the memory, and data can no longer be read from or written to the disk. Before you re-attach this encrypted disk, ensure that the custom key is enabled.

If you use a custom key to encrypt disks and this custom key is then disabled or scheduled for deletion, data cannot be read from or written to these disks or may never be restored. See **Table 4-2** for more information.

Table 4-2 Impact of custom key unavailability

Custom Key Status	Impact	How to Restore
Disabled	• For an encrypted disk already attached: Reads and writes to the disk are normal. If the disk is detached, it cannot be attached again. • For an encrypted disk not attached: The disk cannot be attached anymore.	Enable the custom key. For details, see **Enabling One or More Custom Keys**.
Scheduled deletion		Cancel the scheduled deletion for the custom key. For details, see **Canceling the Scheduled Deletion of One or More Custom Keys**.
Deleted		Data on the disks can never be restored.

NOTICE

You will be billed for the custom keys you use. If pay-per-use keys are used, ensure that you have sufficient account balance. If yearly/monthly keys are used, renew your order timely. Or, your services may be interrupted and data may never be restored if encrypted disks become inaccessible.

Encryption Scenarios

- **System disk encryption**

  System disks are purchased along with servers and cannot be purchased separately. So whether a system disk is encrypted or not depends on the image selected during the server creation. See the following table for details.

  Table 4-3 Encryption relationship between images and system disks

Creating Server Using Encrypted Image	Whether System Disk Will Be Encrypted	Description
Yes	Yes	For details, see **Creating Encrypted Images**.

Creating Server Using Encrypted Image	Whether System Disk Will Be Encrypted	Description
No	No	If you want to use a non-encrypted image to create an encrypted system disk, replicate the image as an encrypted image and then use it to create a server. For details, see **Replicating Images Within a Region**.

- **Data disk encryption**

  Data disks can be purchased along with servers or separately. Whether data disks are encrypted depends on their data sources. See the following table for details.

  Table 4-4 Encryption relationship between backups, snapshots, images, and data disks

Purchased On	Method of Purchase	Whether Data Disk Will Be Encrypted	Description
The ECS console	Purchased together with the server	Yes/No	When a data disk is purchased together with a server, you can choose to encrypt the disk or not. For details, see **Getting Started** > **Creating an ECS** > **Step 1: Configure Basic Settings** in the *Elastic Cloud Server User Guide*.
The EVS console	No data source selected	Yes/No	When an empty disk is created, you can choose whether to encrypt the disk or not. The encryption attribute of the disk cannot be changed after the disk has been created.

Purchased On	Method of Purchase	Whether Data Disk Will Be Encrypted	Description
	Creating from a backup	Yes/No	• When a disk is created from a backup, you can choose whether to encrypt the disk or not. The encryption attributes of the disk and backup do not need to be the same. • When you create a backup for a system or data disk, the encryption attribute of the backup will be the same as that of the disk.
	Creating from a snapshot (The snapshot's source disk is encrypted.)	Yes	A snapshot created from an encrypted disk is also encrypted.
	Creating from a snapshot (The snapshot's source disk is not encrypted.)	No	A snapshot created from a non-encrypted disk is not encrypted.
	Creating from an image (The image's source disk is encrypted.)	Yes	-
	Creating from an image (The image's source disk is not encrypted.)	No	-

Who Can Use the Encryption Function?

When a user uses the encryption function, the condition varies depending on whether the user is the first one ever in the current region or project to use this function.

- If the user is the first user, the user needs to follow the prompt to create an agency, which grants KMS Administrator permissions to EVS. Then the user can create and obtain keys to encrypt and decrypt disks.

  NOTE

  The first user must have the KMS Administrator permissions to create the agency. If the user does not have the KMS Administrator permissions, contact the account administrator to grant the permissions first.

- If the user is not the first user, the user can use encryption directly.

4.4 Disk Backup

What Is EVS Disk Backup?

Cloud Disk Backup provided by Cloud Backup and Recovery (CBR) allows you to create backups for your EVS disks while servers are running. If data loss or damage occurred due to virus invasions, accidental deletions, or software/hardware faults, you can use backups to restore data, guaranteeing your data integrity and security.

CBR Architecture

CBR involves backups, vaults, and policies.

Backup

A backup is a copy of a particular chunk of data and is usually stored elsewhere so that it may be used to restore the original data in the event of data loss.

There are the following types of backups:

- Cloud disk backup: provides snapshot-based backups for EVS disks.
- Cloud server backup: uses the consistency snapshot technology to protect data for ECSs and BMSs. Backups of non-database servers are non-database server backups, and those of database servers are application-consistent backups.
- SFS Turbo backup: backs up data of SFS Turbo file systems.
- Hybrid cloud backup: protects data of VMware VMs by storing their backups to the cloud. You can manage the backups on the CBR console.
- File backup: backs up data of a single or multiple files, instead of the entire cloud servers or on-premises hosts.
- Desktop backup: backs up data of Workspace desktops.

Vault

CBR stores backups in vaults. Before creating a backup, you need to create at least one vault and associate the resources you want to back up with the vaults. Then the resources can be backed up to the associated vaults.

Different types of resources must be backed up to different types of vaults. For example, cloud servers must be backed up to server backup vaults, not disk backup vaults or any other types of vaults.

Policy

There are backup policies and replication policies.

- A backup policy defines when you want to take a backup and for how long you would retain each backup.
- A replication policy defines when you want to replicate from backup vaults and for how long you would retain each replica. Backup replicas are stored in replication vaults.

Organizational policies

You can manage backup and replication policies for a given organization. The organization administrator or delegated CBR administrator can centrally create and configure organizational backup policies and replication policies for member accounts in the organization.

- Organizational backup policies: An enterprise can use an organization's management account to configure organizational backup policies for all the member accounts in the organization. All member accounts in the organization can use the created organizational backup policies.
- Organizational replication policies: An enterprise can use an organization's management account to configure organizational replication policies for all the member accounts in the organization. All member accounts in the organization can use the created organizational replication policies.

Figure 4-2 CBR architecture

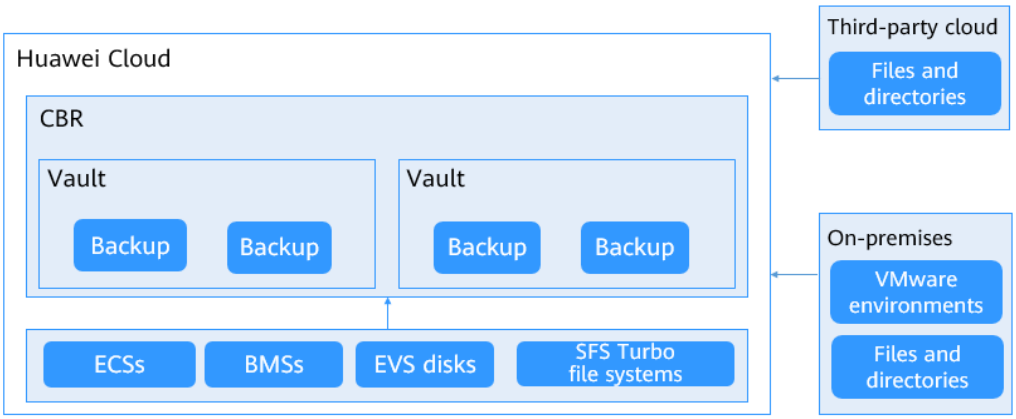

Who Can Use the Backup Function?

Only users with the CBR FullAccess permissions can use the cloud disk backup function. If the user does not have the permissions, contact the account administrator to grant the permissions first.

Application Scenarios

EVS backup can help address your following needs:

- Create and apply backup policies to schedule periodic backups for your EVS disks. You can use the backup data to create new EVS disks or restore to source disks.

- Share backups with other users. You can use the backups shared by other users to create new EVS disks.

4.5 Disk Snapshot

What Is EVS Disk Snapshot?

A snapshot is a complete copy or image of the disk data taken at a specific time. Snapshot is a major DR approach, and you can use a snapshot to restore disk data to the time when the snapshot was created. You can create snapshots for disks on the console or via the API.

EVS disk snapshots are sometimes referred to as snapshots in this document.

You can create snapshots to rapidly save the disk data at specified time points. In addition, you can use snapshots to create new disks so that the created disks will contain the snapshot data in the beginning.

Legacy Snapshot Principles

The following example describes the snapshot principles by creating snapshots s1 and s2 for disk v1 at different points in time:

1. Disk v1 is created, which contains no data.
2. Data d1 and d2 are written to disk v1. Data d1 and d2 are written to new spaces.
3. Snapshot s1 is created for disk v1 modified in step 2. Data d1 and d2 are not saved as another copy elsewhere. Instead, a relationship between snapshot s1 and data d1 and d2 is established.
4. Data d3 is written to disk v1, and data d2 is changed to d4. Data d3 and d4 are written to new spaces, and data d2 is not overwritten. The relationship between snapshot s1 and data d1 and d2 is still valid. Snapshot s1 can be used to restore data if needed.
5. Snapshot s2 is created for disk v1 modified in step 4, and a relationship between snapshot s2 and data d1, d3, and d4 is established.

Figure 4-3 Snapshot Principles

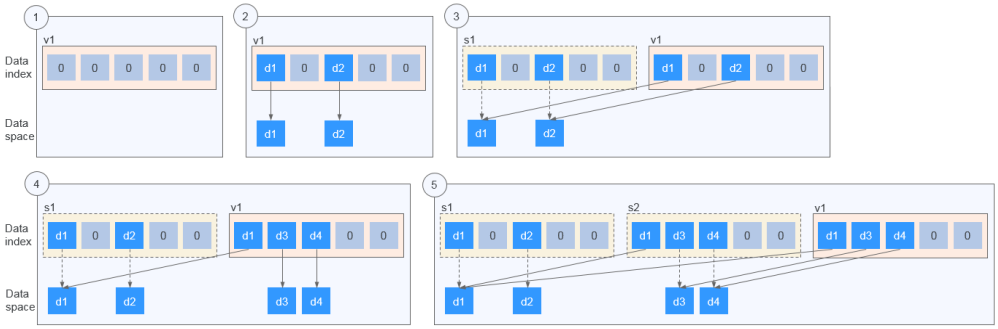

Standard Snapshot Principles

Standard snapshots back up data by data block. They include **full snapshots** and **incremental snapshots**. The first snapshot created for an EVS disk is a full

snapshot, which backs up all data blocks on the disk at the time of the snapshot. Subsequent snapshots are incremental snapshots, which back up only changed data blocks since the last snapshot.

Metadata files of full and incremental snapshots record information about all data blocks when the snapshots were created. So, you can use any snapshot to restore your disk data to the state when the snapshot was created.

Figure 4-4 Standard snapshot principles

Based on the source of data blocks, a snapshot's metadata file contains information about three types of data blocks: **inherited data blocks** (inherited from the last snapshot), **modified data blocks** (have modifications compared with the last snapshot), and **new data blocks** (new compared with the last snapshot).

A snapshot's data file stores only the changed data blocks (modified and new data blocks) compared with the last snapshot.

Let's use the preceding figure for illustration. Assume that data is written to an EVS disk at 09:30 and 10:30. Snapshot 1 is created at 09:00, snapshot 2 is created at 10:00, and snapshot 3 is created at 11:00.

- At 09:00, snapshot 1 is created for the disk. This is the first time that a snapshot is created for this disk, so snapshot 1 is a full snapshot and it contains all the data on the disk, including data blocks A, B, and C. The metadata file of snapshot 1 records information about the disk's full data blocks: A, B, and C.
- After snapshot 1 is created, data block A is changed to A1, data block B is changed to B1, and data block D is added. Then, snapshot 2 is created at 10:00. It is an incremental snapshot. Compared with snapshot 1, data blocks A1, B1, and D are changed data blocks. The metadata file of snapshot 2 records information about the disk's full data blocks: A1, B1, C, and D, among which data block C is inherited from snapshot 1.
- After snapshot 2 is created, data block A1 is changed to A2, data block C is changed to C1, and data block E is added. Then, snapshot 3 is created at 11:00. It is an incremental snapshot. Compared with snapshot 2, data blocks

A2, C1, and E are changed data blocks. The metadata file of snapshot 3 records information about the disk's full data blocks: A2, B1, C1, D, and E, among which data blocks B1 and D are inherited from snapshot 2.

Calculating the Standard Snapshot Storage Usage

The total snapshot storage usage of an EVS disk is calculated by snapshot chain. A snapshot chain collects the storage space used by data blocks of all the snapshots of a disk.

- **Snapshot chain's storage usage calculation after snapshots are added**

  Figure 4-5 Snapshot chain with snapshots added

  Take the scenario in **Figure 4-5** as an example. Assume that the size of a snapshot's data block is fixed at 2 MiB. The snapshot chain's storage usage is calculated as follows:

  - After snapshot 1 is created, the snapshot chain of the disk contains only one snapshot. Snapshot chain's storage usage = Snapshot 1's storage usage = Size of data block A + Size of data block B + Size of data block C = 6 MiB
  - After snapshot 2 is created, the snapshot chain of the disk contains two snapshots: snapshot 1 and snapshot 2. Snapshot chain's storage usage = Snapshot 1's storage usage + Snapshot 2' storage usage = 6 MiB + (Size of data block A1 + Size of data block B1 + Size of data block D) = 12 MiB
  - After snapshot 3 is created, the snapshot chain of the disk contains three snapshots: snapshot 1, snapshot 2, and snapshot 3. Snapshot chain's storage usage = Snapshot 1's storage usage + Snapshot 2' storage usage + Snapshot 3's storage usage = 6 MiB + 6 MiB + (Size of data block A2 + Size of data block C1 + Size of data block E) = 18 MiB

- **Snapshot chain's storage usage calculation after snapshots are deleted**

  When a snapshot is deleted, all data block information in this snapshot's metadata file is traversed, and the following deletion rules are applied:

  - If a data block is inherited by the next snapshot, it will not be deleted.
  - If a data block is not inherited by the next snapshot:

- For an inherited data block, if the previous snapshot that the data block is inherited from is not deleted, the data block will not be deleted. Otherwise, it will be deleted.
- For a modified data block, it will be deleted.
- For a new data block, it will be deleted.

Figure 4-6 Snapshot data block deletion rules

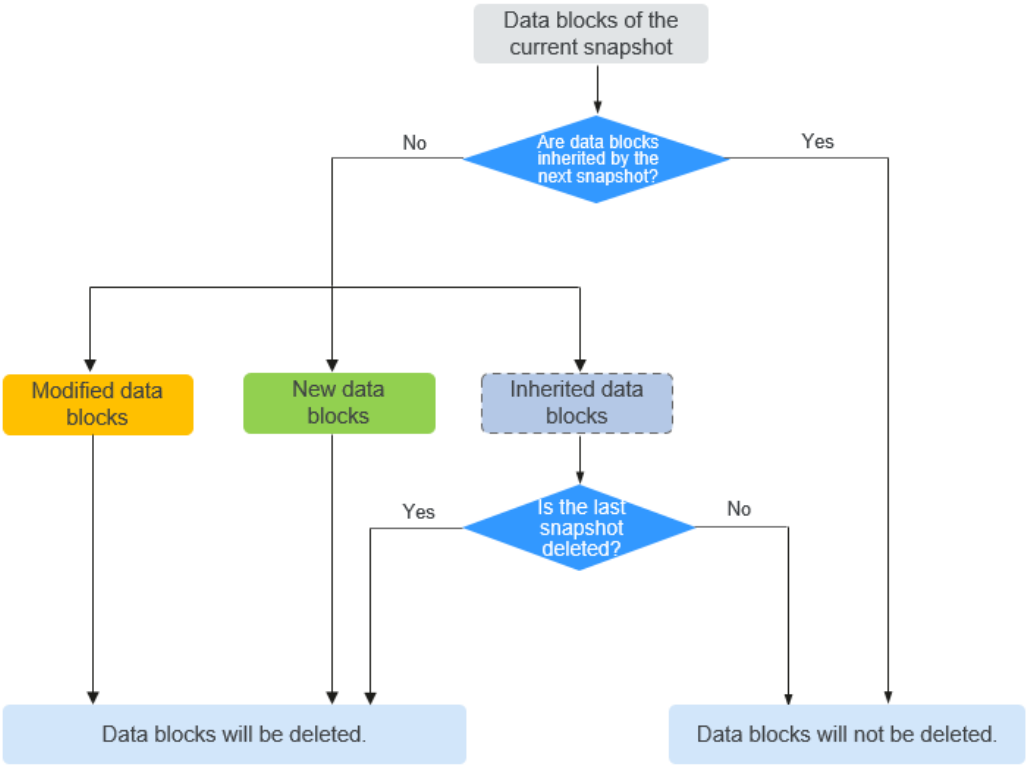

The following example describes how to calculate a snapshot chain's storage usage after snapshots are deleted.

Figure 4-7 Snapshot chain with snapshots deleted

Take the scenario in **Figure 4-7** as an example. Assume that snapshot 2 is deleted at 14:00 and snapshot 3 is deleted at 15:00. The snapshot chain's storage usage is calculated as follows:

- Before any snapshot is deleted, the snapshot chain's storage usage is 18 MiB (Snapshot 1's storage usage + Snapshot 2's storage usage + Snapshot 3's storage usage).
- When snapshot 2 is deleted at 14:00, information about all data blocks in the metadata file of snapshot 2 is traversed.
  - Data block A1: It is not inherited by snapshot 3 and is modified from data block A of snapshot 1. So, data block A1 will be deleted.
  - Data block B1: It is inherited by snapshot 3, so it will not be deleted.
  - Data block C: It is not inherited by snapshot 3, but is inherited from snapshot 1 and snapshot 1 is not deleted. So, data block C will not be deleted.
  - Data block D: It is inherited by snapshot 3. So, it will not be deleted.

  After snapshot 2 is deleted, the snapshot chain's storage usage is 16 MiB (18 MiB – Size of data block A1).
- When snapshot 3 is deleted at 15:00, information about all data blocks in the metadata file of snapshot 3 is traversed.
  - Data block A2: It is not inherited by the next snapshot and is modified from data block A1 of snapshot 2. So, data block A2 will be deleted.
  - Data block B1: It is not inherited by the next snapshot, but is inherited from snapshot 2 and snapshot 2 has been deleted. So, data block B1 will be deleted.
  - Data block C1: It is not inherited by the next snapshot and is modified from data block C of snapshot 2. So, data block C1 will be deleted.

- Data block D: It is not inherited by the next snapshot, but is inherited from snapshot 2 and snapshot 2 has been deleted. So, data block D will be deleted.
- Data block E: It is not inherited by the next snapshot and is newly added in snapshot 3. So, data block E will be deleted.

After snapshot 3 is deleted, the snapshot chain's storage usage is 6 MiB (16 MiB – Size of data block A2 – Size of data block B1 – Size of data block C1 – Size of data block D – Size of data block E).

4.6 Differences Between Disk Backups and Disk Snapshots

Both disk backups and disk snapshots provide redundancies for improved disk data reliability. **Table 4-5** lists the differences between them.

Table 4-5 Differences between backups and snapshots

Metric	Storage Solution	Data Synchronization	DR Range	Service Recovery
Backup	Backups are stored in OBS, instead of disks. This ensures data restoration upon disk damage or corruption.	A backup is a copy of a disk taken at a given time and is stored in a different location. Automatic backup can be performed based on backup policies. Deleting a disk will not delete its backups.	A backup and its source disk reside in different AZs.	To restore data and recover services, you can restore the backups to their original disks or create new disks from the backups.

Metric	Storage Solution	Data Synchronization	DR Range	Service Recovery
Legacy Snapshot	Snapshots are stored on the same disk as the original data. **NOTE** Creating a backup requires a certain amount of time because data needs to be transferred to OBS. Creating or rolling back a snapshot consumes less time than creating a backup.	A snapshot is the state of a disk at a specific point in time and is stored on the same disk. If the disk is deleted, all its snapshots will also be deleted. For example, if you reinstalled or changed the server OS, snapshots of the system disk were also automatically deleted. Snapshots of the data disks can be used as usual.	A snapshot and its source disk reside in the same AZ.	You can use a snapshot to roll back data to its source disk or create a new disk.
Standard Snapshot	Snapshots are stored in OBS, instead of disks. This ensures data restoration upon disk damage or corruption.	A snapshot is the state of a disk at a specific point in time and is stored in OBS. If a disk is deleted, all its snapshots will not be deleted.	A snapshot and its source disk reside in different AZs.	You can use a snapshot to roll back data to its source disk or create a new disk. The data reliability is high.

4.7 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 4-8 Three-copy redundancy

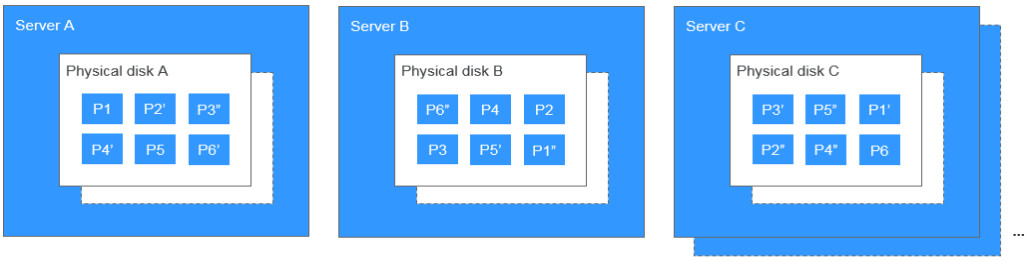

How Does the Three-Copy Redundancy Keep Data Consistency?

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 4-9 shows the data rebuild process.

Figure 4-9 Data rebuild process

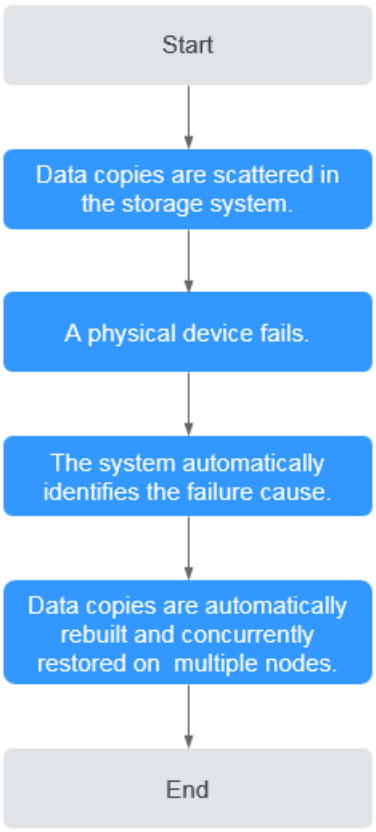

Figure 4-10 shows the data rebuild principle. For example, if physical disks on server F are faulty, the data blocks on these physical disks will be rebuilt on the physical disks of other servers.

Figure 4-10 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.

5 Billing

5.1 Billing for EVS Disks

Billing Items

EVS disks are billed based on the disk type, size, and usage duration. For details, see **EVS Pricing Details**.

- 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.

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 the 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 01:00:00-01:30:00, and the other is the billing record generated for the 200 GiB from 01:30:01-02:00:00.
Performance change	Throughput and IOPS cannot be billed on a yearly/monthly basis. For yearly/monthly disks that allow you to configure performance, their capacities are billed on a yearly/monthly basis, and their performance is billed pay per use. Pay-per-use pricing applies after a performance change.	Throughput and IOPS can be billed on a pay-per-use basis. Pay-per-use pricing applies after a performance change.
Disk type change	You need to pay for the price difference incurred by a disk type change. **NOTE** The expiration time of the EVS disk remains unchanged after a disk type change.	Pay-per-use pricing of the new disk type applies.

Item	Yearly/Monthly	Pay-per-Use
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 disks cannot be changed to yearly/monthly billing separately. They must be changed together with servers. After the change, they have the same expiration times as the 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.

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.

5.2 Billing for EVS Recycle Bin

Billing Items

EVS disks in the recycle bin are billed based on the disk type, size, and storage period. For details, see **EVS Pricing Details**.

- The billing starts after the disks have been moved to the recycle bin upon deletion.
- The billing ends after the disks have been deleted from the recycle bin.

Billing Modes

Disks in the recycle bin are billed on a pay-per-use basis.

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.

Viewing EVS Recycle Bin Bills

Step 1 Log in to the **management console**.

Step 2 Click **Billing** from the top menu bar to go to the Billing Center.

Step 3 Choose **Billing** > **Transactions and Detailed Bills** and click the **Bill Details** tab.

Step 4 Select a billing cycle, select **Usage** for **Settings**, and select **By billing cycle** for **Data Period**.

Step 5 In the list, select **Elastic Volume Service (EVS)** for **Service Type**.

Step 6 In the **Resource Tag** column, find the **_sys_recycle_bin_volume:true** tag, which identifies the EVS recycle bin bill.

Figure 5-1 Viewing the EVS recycle bin bill

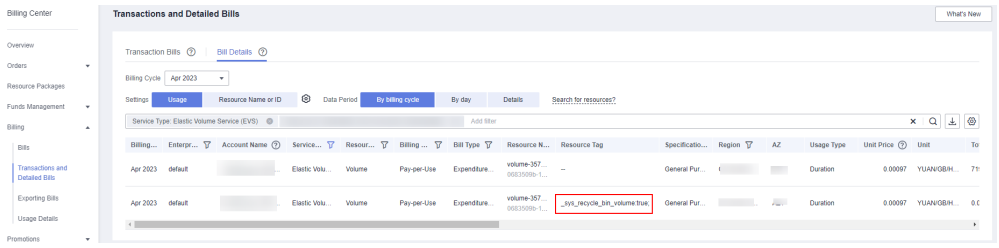

NOTE

You can also click the **Export** button next to the search box to export all bills and find out the recycle bin bill by filtering the **_sys_recycle_bin_volume:true** tag.

----End

5.3 Billing for EVS Snapshots

Legacy snapshots are in OBT. You can use them free of charge.

Standard snapshots are in commercial use. This section describes the billing of standard snapshots.

There are differences between legacy snapshots and standard snapshots. For details, see the snapshot constraints in **Notes and Constraints**.

Billing Items

EVS snapshots are billed based on the storage usage of snapshot chains and the usage periods. A snapshot chain collects the storage space used by data blocks of all the snapshots of a disk.

NOTE

Instant Snapshot Restore is currently not billed. Snapshots with Instant Snapshot Restored enabled are only billed based on the storage usage.

Billing Modes

EVS snapshots support pay-per-use billing.

Pay-per-use is a postpaid payment model. Your snapshot 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 Formula

Standard snapshot storage price = Snapshot unit price x Snapshot chain storage usage x Billed usage period

- Snapshot unit price: The storage price per GiB per hour. The unit is $/GiB/hour. For details, see **EVS Pricing Details**.
- Snapshot chain storage usage: The total storage space used by data blocks of all the snapshots of a disk. For details, see **Calculating the Snapshot Chain Storage Usage**.
- Billed usage period: The billing starts after snapshots are created and ends after snapshots are deleted.

Billing example:

Assume that you created snapshots for a disk, the snapshots took up 100 GiB of storage, and you kept the snapshots for three hours. If the snapshot unit price was $0.000065 USD/GiB/hour, then the price of snapshots would be calculated as follows:

Snapshot storage price = $0.000065/GiB/hour x 100 GiB x 3 hours = $0.0195

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 snapshot will enter the grace period and retention period before being deleted.

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

How Do I View the Bill of My Snapshots?

NOTE

You can only view the snapshot bill by snapshot chain. You cannot view the bill of a single snapshot.

Step 1 Log in to the **management console**.

Step 2 (Optional) View the snapshot chain ID of the desired disk.

If you already got the snapshot chain ID, skip this step.

1. Choose **Storage** > **Elastic Volume Service**.
2. In the navigation pane on the left, choose **Elastic Volume Service** > **Snapshots** to go to the **Snapshots** page.
3. Click the **Snapshot Chains** tab. In the search box above the snapshot chain list, search for the snapshot chain using the disk ID and take note of the snapshot chain ID.

Step 3 Click **Billing** from the top menu bar to go to the Billing Center.

Step 4 Choose **Billing** > **Transactions and Detailed Bills** and click the **Bill Details** tab.

Step 5 Select a billing cycle, select **Usage** for **Settings**, and select **By billing cycle** for **Data Period**.

Step 6 In the search box, select **Elastic Volume Service EVS** for **Service Type**, select **Resource ID**, and enter the snapshot chain ID.

Step 7 Click , and the displayed bill is the snapshot bill.

----End

5.4 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.

Figure 5-2 Impacts on yearly/monthly resources before and after expiration

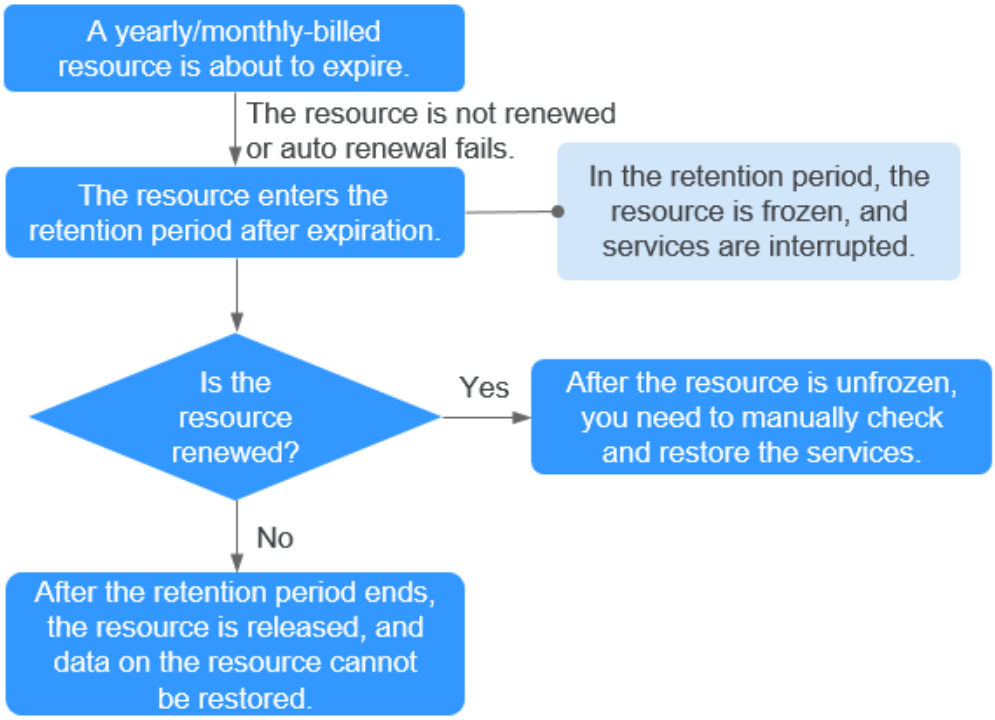

Impact on Services When Resources Are Frozen, Unfrozen, or Released

- 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 log in to the management console, detach the disk, and release the resource. For details, see section "Releasing Resources" in the *Billing Center User Guide*.

Table 5-1 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 5-1 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 mode.	• Change the server 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 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.

5.5 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 5-3 Impacts on pay-per-use resources before and after account arrears

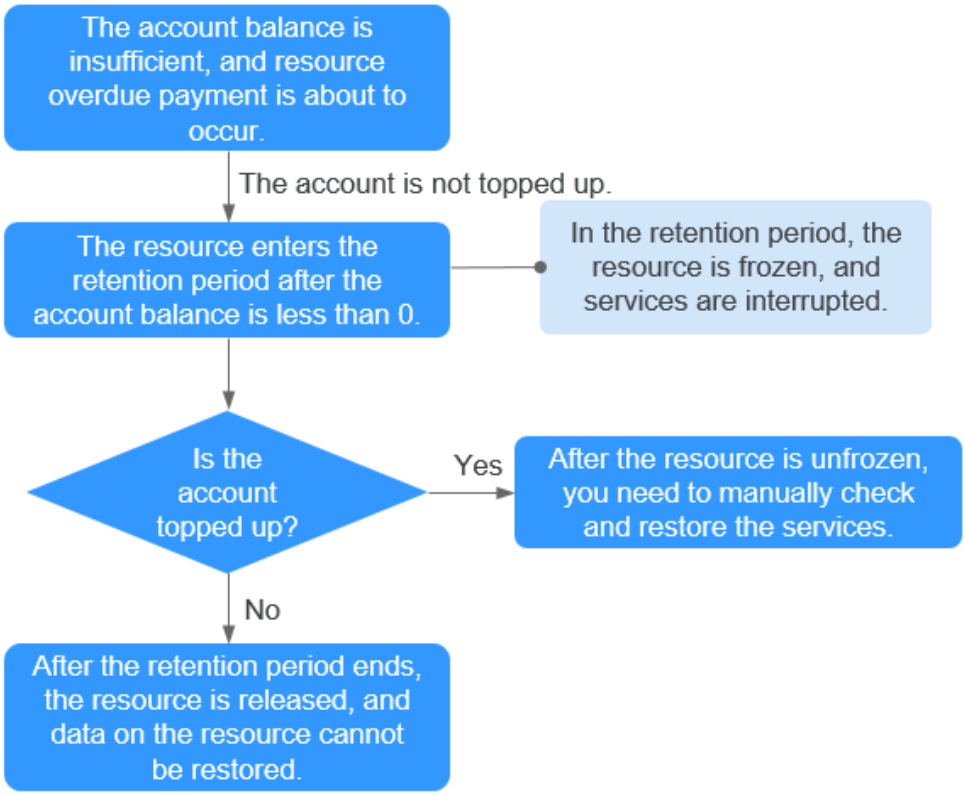

Impact on Services When Resources Are Frozen, Unfrozen, or Released

- 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 log in to the management console, detach the disk, and then delete it. For how to delete a disk, see **Deleting EVS Disks**.

Table 5-2 lists the common usage scenarios and suggestions on pay-per-use disks. You can enable account balance alert, change disk billing mode from pay-per-use to yearly/monthly, and pay attention to account balance and resource freezing notifications to keep up with the latest resource information, ensuring that your services and data are not affected.

Table 5-2 Common usage scenarios and suggestions

Common Usage Scenario	Suggestions
Resources are billed in pay-per-use mode.	• Top up your account in time to keep sufficient account balance. • Pay attention to notifications about account arrears and top up your account in time.
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.
The server is billed in pay-per-use mode, but the attached disks are billed in yearly/monthly mode.	• Change the server 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 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 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.

Common Usage Scenario	Suggestions
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.

6 Security

6.1 Shared Responsibilities

Huawei guarantees that its commitment to cyber security will never be outweighed by the consideration of commercial interests. To cope with emerging cloud security challenges and pervasive cloud security threats and attacks, Huawei Cloud builds a comprehensive cloud service security assurance system for different regions and industries based on Huawei's unique software and hardware advantages, laws, regulations, industry standards, and security ecosystem.

Figure 6-1 illustrates the responsibilities shared by Huawei Cloud and users.

- **Huawei Cloud**: Ensure the security of cloud services and provide secure clouds. Huawei Cloud's security responsibilities include ensuring the security of our IaaS, PaaS, and SaaS services, as well as the physical environments of the Huawei Cloud data centers where our IaaS, PaaS, and SaaS services operate. Huawei Cloud is responsible for not only the security functions and performance of our infrastructure, cloud services, and technologies, but also for the overall cloud O&M security and, in the broader sense, the security compliance of our infrastructure and services.
- **Tenant**: Use the cloud securely. Tenants of Huawei Cloud are responsible for the secure and effective management of the tenant-customized configurations of cloud services including IaaS, PaaS, and SaaS. This includes but is not limited to virtual networks, the OS of virtual machine hosts and guests, virtual firewalls, API Gateway, advanced security services, all types of cloud services, tenant data, identity accounts, and key management.

Huawei Cloud Security White Paper elaborates on the ideas and measures for building Huawei Cloud security, including cloud security strategies, the shared responsibility model, compliance and privacy, security organizations and personnel, infrastructure security, tenant service and security, engineering security, O&M security, and ecosystem security.

Figure 6-1 Huawei Cloud shared security responsibility model

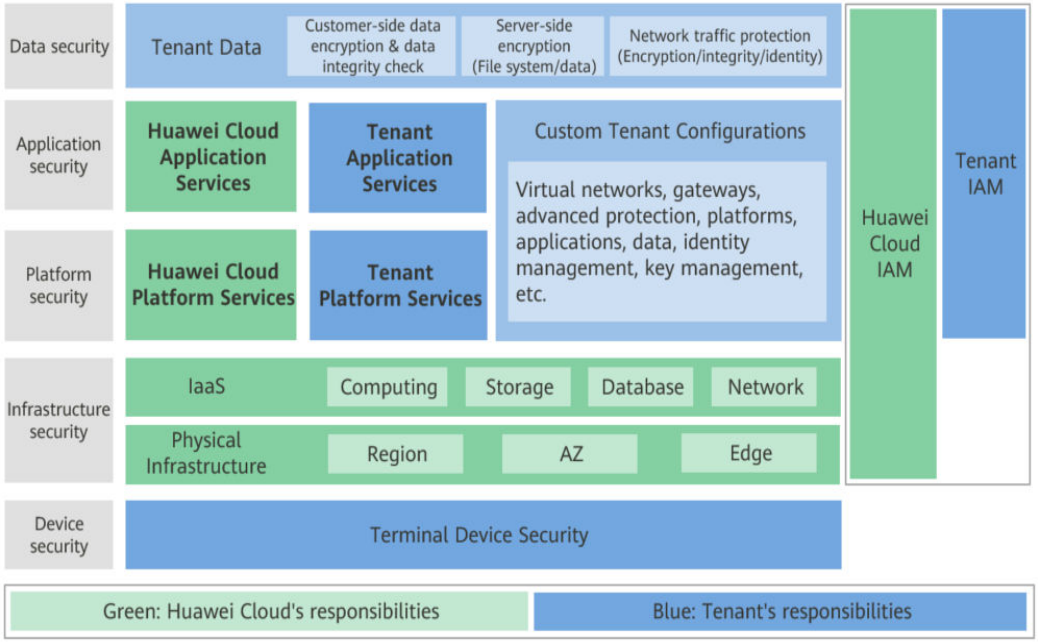

6.2 Identity Authentication and Access Control

You can access EVS through the EVS console, APIs, and SDKs. No matter which method you choose, you actually use REST APIs to access EVS.

EVS APIs support only authenticated requests. You must obtain the authentication information from Huawei Cloud Identity and Access Management (IAM) before you can access EVS. For details about IAM authentication, see **Authentication**.

Access Control

You can use IAM to securely control access to your EVS resources.

Table 6-1 EVS access control

Method		Description	Reference
Permissions management	IAM permissions	IAM permissions define which actions on your cloud resources are allowed or denied. After creating an IAM user, the administrator needs to add it to a user group and grant the permissions required by EVS to the user group. Then, all users in this group automatically inherit the granted permissions.	**Permissions**

6.3 Data Protection

EVS uses the encryption function to protect the confidentiality of static data stored on EVS disks.

Table 6-2 EVS data protection

Measure	Description	Reference
Disk encryption	1. Empty encrypted disks can be created. 2. Encrypted disks can be created from snapshots, backups, and images. 3. AES-256 is used to encrypt the server-side static data by default. 4. KMS keys can be used to encrypt static data. 5. Both data disks and system disks can be encrypted. 6. Snapshots, backups, and images created from encrypted disks are encrypted by default.	**Managing Encrypted EVS Disks**

6.4 Auditing

Cloud Trace Service (CTS) records operations on the cloud resources in your account. You can use the logs generated by CTS to perform security analysis, track resource changes, audit compliance, and locate faults. After you enable CTS and configure a tracker, CTS can record management and data traces of EVS for auditing.

For details about how to enable and configure CTS, see **CTS Getting Started**.

For the EVS management and data traces supported by CTS, see **Auditing**.

6.5 Risk Monitoring

EVS uses Cloud Eye to perform monitoring over resources and operations, helping you monitor your disk usages and receive alarms and notifications in real time. Alternatively, you can monitor the IOPS, throughput, and latency of your disks in real time.

For details about supported EVS metrics and how to create alarm rules, see **Monitoring**.

6.6 Fault Recovery

EVS offers a variety of fault recovery methods. For details, **Table 6-3**.

Table 6-3 Fault recovery

Method	Description	Reference
EVS backup	CBR provides the cloud disk backup function, which allows you to back up EVS disks while the servers are running. In case of a virus, accidental deletion, or software/hardware fault, you can restore data to any point in the past when backups were created to guarantee data integrity and security.	**Restoring Data Using a Cloud Disk Backup**
Disk creation from snapshots	You can use EVS snapshots to create new disks so that the new disks will contain the snapshot data once being created.	**Creating an EVS Disk from a Snapshot (OBT)**
Snapshot data rollback to disks	If data on an EVS disk is incorrect or damaged, you can quickly restore data by rolling back the disk to the state when the snapshot was created.	**Rolling Back Data from a Snapshot (OBT)**

7 Permissions

If you need to assign different permissions to employees in your enterprise to access your EVS resources, 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 (such as **na-mexico-1** in the **LA-Mexico City1** region) 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 Policies and Supported Actions**.

Table 7-1 lists all the system-defined roles and policies supported by EVS.

Table 7-1 System-defined roles and policies supported by EVS

Role/Policy Name	Description	Type	Dependency
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 ReadOnlyAccess	Read-only permissions for EVS. Users granted these permissions can view EVS resource data only.	System-defined policy	None
Server Administrator	Full permissions for EVS	System-defined role	None

Table 7-2 lists the common operations supported by each system-defined policy of EVS. Select the policies as required.

Table 7-2 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	√	x
Detaching disks	√	x
Deleting disks	√	x
Expanding disk capacities	√	x
Creating snapshots	√	x
Deleting snapshots	√	x

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

Helpful Links

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

8 Notes and Constraints

This section describes the constraints on using EVS.

Specifications

Table 8-1 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 Disk Performance**.

Security

Table 8-2 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. • If you use a standard snapshot with Instant Snapshot Restore enabled to create a disk, the disk created will be an encrypted disk. You cannot modify the encryption of the disk.

Quotas

You can log in to the console to view default quotas. For details, see **How Do I View My Quotas?** You can **submit a service ticket** to apply for a larger quota if needed.

Table 8-3 Quotas

Resource Type	Default Quota
Tags	20
Disks	Default quotas vary depending on regions. See the quotas shown on the console.
Disk capacity (GB)	
Snapshots	

Operations

Table 8-4 Operations

Scenario	Item	Description
Disk creation	Maximum number of disks that can be created at a time	100
	Disk creation from snapshot	• Batch creation is not supported when disks are created from snapshots. Only one disk can be created from a snapshot at a time. **Snapshot function in OBT** • A disk created from a snapshot has the same device type (SCSI or VBD), encryption attribute, AZ, and disk type as the snapshot's source disk. • 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 create new disks. **Snapshot function in commercial use** • If Instant Snapshot Restore is not enabled for a standard snapshot, you can only use it to create disks when its upload progress is complete. • If Instant Snapshot Restore is enabled for a standard snapshot, when its upload is in progress, you can use it to create disks but cannot change the device type (SCSI or VBD), encryption attribute, AZ, and type of the new disk. They are kept the same as those of the snapshot's source disk. • After a standard snapshot is uploaded, you can change the device type (SCSI or VBD), encryption attribute, AZ, and type of the disks as required.

Scenario	Item	Description
	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 be used to create another disk only 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 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.
	Device type	The device type of a disk cannot be changed after the disk has been created.
	Disk sharing	The sharing attribute of a disk cannot be changed after the disk has been created.
	Disk encryption	The encryption attribute of a disk cannot be changed after the disk has been created.
Disk attachment	Constraints on region and AZ	The disk and server must be in the same region and AZ.
	Maximum number of servers that a non-shared disk can be attached to	1
	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 value varies with ECS types. For details, see **Can I Attach Multiple Disks to an ECS?**
	Max. number of disks that can be attached to a BMS	60 (1 system disk and 59 data disks) Only SCSI disks can be attached to BMSs.

Scenario	Item	Description
	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. That is, 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, that is, its server can either be in the **Running** or **Stopped** state.

Scenario	Item	Description
Disk deletion	Deletion of pay-per-use disks Unsubscription of yearly/monthly 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. • The disk is not added to any replication pair in the Storage Disaster Recovery Service (SDRS). For any disk already added to a replication pair, you need to first **delete the replication pair** and then delete the disk. • Yearly/Monthly system disks cannot be unsubscribed from separately. They must be unsubscribed from together with their servers. • Non-shared, yearly/monthly data disks purchased together with or later added to a yearly/monthly server have the same expiration time as the server. They can be unsubscribed from together with the server or separately when their statuses are **In-use**, **Available**, or **Error**. • Yearly/Monthly data disks purchased on the EVS console have different expiration times as the server. They can be unsubscribed from separately.

Scenario	Item	Description
Snapshot creation	/	• Snapshots can be created for both system disks and data disks. • Snapshots of encrypted disks are stored encrypted, and those of non-encrypted disks are stored non-encrypted. **Snapshot function in OBT** • You can manually create a maximum of seven snapshots for a disk. • The enterprise project of a snapshot is the same as that of the snapshot's source disk. **Snapshot function in commercial use** • You can manually create a maximum of 256 standard snapshots for a disk, of which up to seven can have Instant Snapshot Restore enabled. • You can create one standard snapshot for a disk at a time. You can only create the next standard snapshot for the same disk after the previous snapshot has been created. • Standard snapshots cannot be created for the disks in edge AZs. For details about the differences between edge AZs and general AZs, see the **CloudPond User Guide**. • When standard snapshots are created for Common I/O and High I/O disks, Instant Snapshot Restore cannot be enabled. • It usually takes several minutes to create a standard snapshot. The time required varies depending on the amounts of data written to the disk. The larger the data volume, the longer the time required. The initial standard snapshot usually takes more time because data of the entire disk is backed up. Subsequent standard snapshots are quicker, but the time required is still determined by the amounts of changed data compared with each last snapshot. The more the changed data, the longer the time required.

Scenario	Item	Description
		• If the data on a disk is rolled back from a snapshot, the next standard snapshot created for this disk will be a full snapshot. • During the creation of a standard snapshot, any incremental data written to the disk will not be backed up to the snapshot created. • During the creation of a standard snapshot, deleting the snapshot's source disk does not affect the creation of the snapshot.
Use of Instant Snapshot Restore	/	• Instant Snapshot Restore is supported for the following types of disks: Extreme SSD V2, Extreme SSD, General Purpose SSD V2, General Purpose SSD, and High I/O. • You can only enable Instant Snapshot Restore when creating standard snapshots. It cannot be enabled later. • You can enable Instant Snapshot Restore for up to seven snapshots for a disk. • When Instant Snapshot Restore is enabled and snapshots are being created, you cannot disable Instant Snapshot Restore. • When you delete a disk whose standard snapshots have Instant Snapshot Restore enabled, the snapshots will be not deleted, but Instant Snapshot Restore will be disabled automatically.

Scenario	Item	Description
Snapshot data rollback to disk	/	• Snapshot data can only be rolled back to source EVS disks. Rollback to a different disk is not possible. • You can only roll back disk data from a snapshot when the source disk status is **Available** (not attached to any server) or **Rollback failed**. If the source disk is attached, detach the disk first. • If a snapshot is being created, 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.

Scenario	Item	Description
Snapshot deletion	/	• If a snapshot is deleted, disks rolled back or created from this snapshot are not affected. **Snapshot function in commercial use** • Standard snapshots are not deleted even if their source disks are deleted. • When you delete a disk whose standard snapshots have Instant Snapshot Restore enabled, the standard snapshots will be not deleted, but Instant Snapshot Restore will be disabled automatically. • If you reinstall or change the server OS, standard snapshots will be not deleted, but Instant Snapshot Restore will be disabled automatically if it has been enabled for the standard snapshots of the system disk. **Snapshot function in OBT** • If a snapshot's source disk is deleted, all legacy snapshots of this disk are also deleted. • If you reinstall or change the server OS, snapshots of the system disk are automatically deleted. Those of the data disks can be used as usual. • A snapshot whose name starts with **autobk_snapshot_vbs_**, **manualbk_snapshot_vbs_**, **autobk_snapshot_csbs_**, or **manualbk_snapshot_csbs_** is automatically generated during backup. You can only check details of such snapshots and cannot delete them.
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. • Changing the disk type may affect the disk performance, so change the type during off-peak hours.

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 backend issue. If this happens, try again later.
Recycle bin management	/	• When you delete a disk, regardless of whether the disk will be moved to the recycle bin or not, the disk snapshots will always be deleted permanently. • There are no limits on the capacity and quantity of disks in the recycle bin. • In the recycle bin, disks can be stored for 7 days. After 7 days, the system will permanently delete the disks, and the disks cannot be recovered.

9 EVS and Other Services

Figure 9-1 shows the relationships between EVS and other services.

Figure 9-1 Relationships between EVS and other services

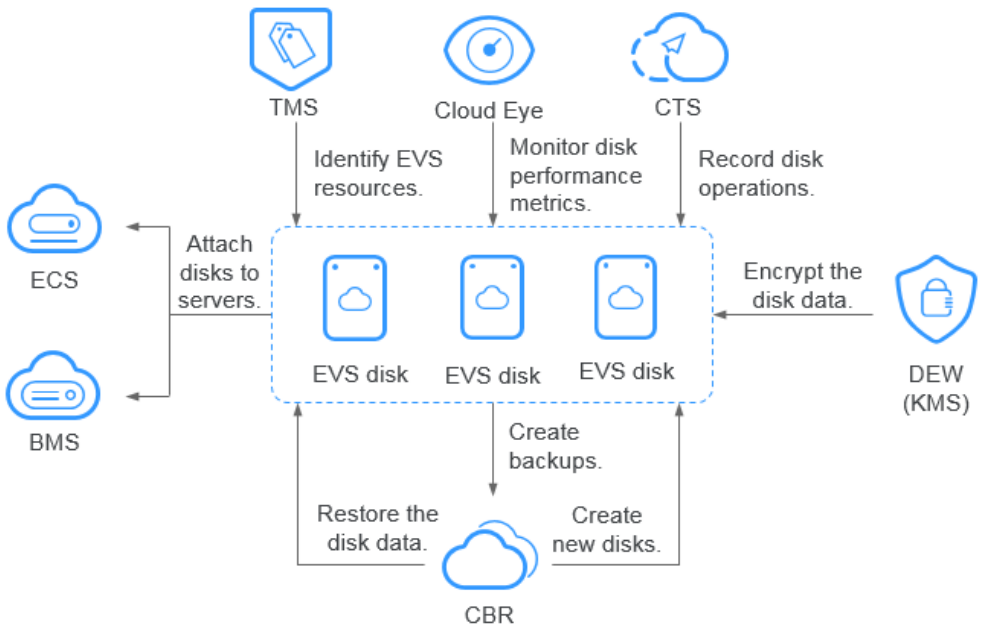

Table 9-1 EVS and other services

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	• **Disk Backup** • **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	• **Disk Encryption** • **Managing Encrypted EVS Disks**
After EVS is enabled, the performance metrics of monitored disks can be viewed through Cloud Eye without installing any additional plug-in. 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**

10 Basic Concepts

10.1 EVS Concepts

Table 10-1 EVS concepts

Concept	Description
IOPS	Number of read/write operations performed by an EVS disk per second
Throughput	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 capability	The burst capability allows the IOPS of a small-capacity disk to reach the disk IOPS burst limit, which can surpass the disk IOPS limit within a certain period of time.
VBD	A device type of EVS disks. VBD EVS disks only support basic SCSI read/write commands.
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.

10.2 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. AZs within a region are interconnected using high-speed optical fibers, to support cross-AZ high-availability systems.

Figure 10-1 shows the relationship between regions and AZs.

Figure 10-1 Regions and AZs

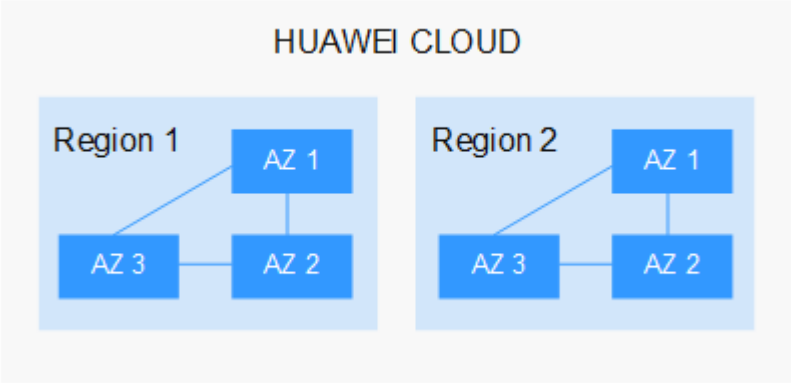

Huawei Cloud provides services in many regions around the world. You can select a region and an AZ based on requirements. For more information, see **Huawei Cloud Global Regions**.

Selecting a Region

When selecting a region, consider the following factors:

- Location

  It is recommended that you select the closest region for lower network latency and quick access.

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

    NOTE

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

- Resource price

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

Selecting an AZ

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

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

Regions and Endpoints

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

A Change History

Released On	Description
2023-11-01	This issue is the tenth official release, which incorporates the following change: Updated and added permissions constraints in sections "EVS Encryption" and "EVS Backup."
2023-07-20	This issue is the ninth official release, which incorporates the following change: Updated and added constraints.
2023-06-15	This issue is the eighth official release, which incorporates the following changes: Updated: Added descriptions about General Purpose SSD V2 disks in sections "Disk Types and Performance" and "Constraints." Added: Added support for the General Purpose SSD V2 disk type.
2023-02-14	This issue is the seventh official release, which incorporates the following change Changed the capacity unit to GiB in section "Disk Types and Performance."
2022-11-14	This issue is the sixth official release, which incorporates the following change: ● Added section "Security."
2022-03-26	This issue is the fifth official release, which incorporates the following change: ● Added the Extreme SSD disk type in section "Disk Types and Performance."

Released On	Description
2018-09-10	This issue is the fourth official release, which incorporates the following change: ● Added section "EVS Three-Copy Redundancy."
2018-07-30	This issue is the third official release, which incorporates the following changes: ● Added content "Differences Between EVS, SFS, and OBS" in section "What Is EVS?" ● Added precautions for using shared EVS disks together with SCSI. ● Modified disk performance metrics.
2018-06-30	This issue is the second official release, which incorporates the following changes: ● Added section "Differences Between EVS Backups and EVS Snapshots." ● Optimized the content under "Do I Need to Install a Driver for SCSI EVS Disks?" from the perspective of KVM and Xen ECSs in section "Device Types and Usage Instructions."
2018-06-15	This issue is the first official release.