Elastic Volume Service

# Service Overview

**Issue** 04

**Date** 2018-09-10

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.

# Contents

1 What Is Elastic Volume Service?..... 1
2 Disk Types and Disk Performance..... 5
3 Device Types and Usage Instructions..... 11
4 Shared EVS Disks and Usage Instructions..... 13
5 EVS Disk Encryption..... 17
6 EVS Disk Backup..... 21
7 EVS Snapshot (OBT)..... 22
8 Differences Between EVS Disk Backup and EVS Snapshot..... 25
9 EVS Three-Copies of Data Mechanism..... 27
10 Billing..... 31
10.1 Billing for Disks..... 31
11 Constraints..... 33
12 EVS and Other Services..... 37
13 Basic Concepts..... 39
13.1 Basic Concepts..... 39
13.2 Region and AZ..... 39
A Change History..... 42

  

# 1 What Is Elastic Volume Service?

## Overview

Elastic Volume Service (EVS) offers scalable block storage for cloud servers. With high reliability, high performance, and rich specifications, EVS disks can be used for distributed file systems, development and testing environments, data warehouse applications, and high-performance computing (HPC) scenarios to meet diverse service requirements. Servers that EVS supports include Elastic Cloud Servers (ECSs) and Bare Metal Servers (BMSs).

EVS disks are similar to hard disks in PCs. They must be attached to servers for use and cannot be used alone. You can initialize EVS disks, create file systems on them, and store data persistently on them.

EVS disks are sometimes just referred to as disks.

**Figure 1-1** EVS architecture

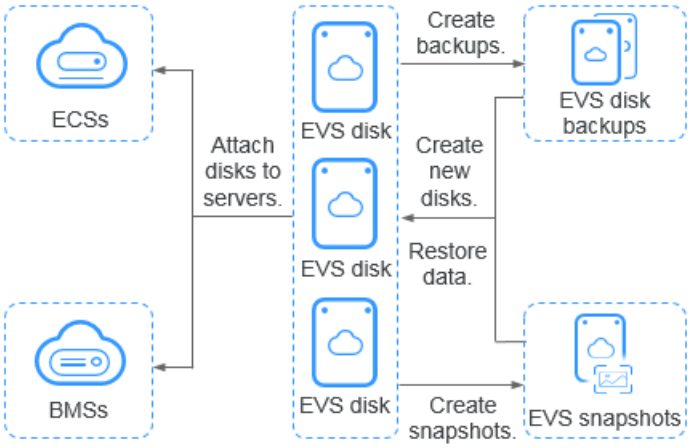


## EVS Advantages

EVS provides disk resources for servers, and its advantages are as follows:

  

**Table 1-1** EVS advantages

| Advantage | Description | Related Knowledge |
| --- | --- | --- |
| Various disk types | EVS provides various disk types for you to choose from, and EVS disks can be used as data disks and system disks for servers. You may select the disk type based on your budget and service requirements. | **Disk Types and Disk Performance** |
| Elastic scalability | The capacity of an EVS disk you can create ranges from 10 GB to 32 TB. Expand the disk capacity when it no longer meets your needs. The minimum expansion increment is 1 GB, and a disk can be expanded to up to 32 TB. EVS also supports smooth capacity expansion without interrupting services. | **Expansion Overview** |
| | Besides the disk capacity limit, the additional space you can add during an expansion is also affected by the capacity quota. The system will prompt you with the remaining quota, and the space added cannot exceed that. You may increase the quota if you want to expand your disk but the remaining quota is insufficient. | **Querying EVS Resource Quotas** |
| High security and reliability | Both system disks and data disks support data encryption to ensure data security. | **EVS Disk Encryption** |
| | Data protection functions, such as backups and snapshots, safeguard the disk data, preventing incorrect data caused by application exceptions or attacks. | **EVS Disk Backup**<br>**EVS Snapshot (OBT)** |
| Real-time monitoring | Working with Cloud Eye, EVS allows you to monitor the disk health and operating status at any time. | **Viewing EVS Monitoring Data** |

## Differences Among EVS, SFS, and OBS

Currently, there are three data storage services available for you to choose from: EVS, Scalable File Service (SFS), and Object Storage Service (OBS). The differences are described in the following table.

**Table 1-2** Comparison between SFS, OBS, and EVS

| Dimension | SFS | OBS | EVS |
| --- | --- | --- | --- |
| Concept | SFS provides on-demand high-performance file storage, which can be shared by multiple ECSs. SFS is similar to a remote directory for Windows or Linux OSs. | OBS provides massive, secure, reliable, and cost-effective data storage capabilities for users to store data of any type and size. | EVS provides scalable block storage that features high reliability, high performance, and rich specifications for ECSs to meet service requirements in different scenarios. An EVS disk is similar to a hard disk on a PC. |
| Data storage logic | Stores files and sorts and displays data in the hierarchy of files and folders. | Stores objects. Files can be directly stored. The files automatically generate corresponding system metadata. Users can also customize the metadata of files. | Stores binary data and cannot directly store files. To store files, you need to format the file system first. |
| Access method | Attach file systems to ECSs. You need to specify a network address for access or change the network address to a local directory for access. The NFS and CIFS protocols are used. | You can access OBS through the Internet or Direct Connect. You need to specify the bucket address for access. The transmission protocols such as HTTP and HTTPS are used. | An EVS disk can only be used by mounting to an ECS or BMS and cannot be directly accessed by OS applications. It must be formatted into a file system for access. |
| Application Scenario | High-performance computing (HPC), media processing, file sharing, content management, and web services<br>**NOTE**<br>HPC: High bandwidth is required for shared file storage, such as gene sequencing and image rendering. | Big data analysis, static website hosting, online video on demand (VoD), gene sequencing, and intelligent video surveillance | HPC, enterprise core cluster applications, enterprise application systems, and development and testing<br>**NOTE**<br>HPC: High-speed and high-IOPS storage is required, such as industrial design and energy exploration. |
| Capacity | PB-scale | EB-scale | TB-scale |

  

| Dimension | SFS | OBS | EVS |
| --- | --- | --- | --- |
| Latency | 3–10 ms | 10 ms | 1–2 ms |
| IOPS/TPS | 10,000 for a single file system | Quadrillion | 33,000 for a single disk |
| Bandwidth | GB/s | TB/s | MB/s |
| Data sharing supported | Yes | Yes | Yes |
| Remote access supported | Yes | Yes | No |
| Online editing supported | Yes | No | Yes |
| Used independently | Yes | Yes | No |

## Methods of Access

The public cloud system provides a web-based management console and HTTPS-based APIs for you 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. If you have registered with the public cloud, log in to the management console and choose **Elastic Volume Service** on the homepage. If not, see **Registering Yourself on the Public Cloud**.

# 2 Disk Types and Disk Performance

EVS disk types are classified into ultra-high I/O, general purpose SSD, high I/O, and common I/O based on the disk I/O performance. EVS disk types differ in performance and price. Choose the disk type based on your requirements. The details are described as follows:

## EVS Disk Performance

EVS disk performance metrics include:

- IOPS: Number of read/write operations performed by an EVS disk per second
- Throughput: Amount of data successfully transmitted by an EVS disk per second, that is, the amount of data read from and written into an EVS disk
- Read/write I/O latency: Minimum interval between two consecutive read/write operations of an EVS disk

**Table 2-1** EVS disk performance data

| Parameter | Ultra-high I/O | General Purpose SSD | High I/O | Common I/O (Previous Generation Product) |
| --- | --- | --- | --- | --- |
| Max. capacity | • System disk: 1024 GB<br>• Data disk: 32768 GB | • System disk: 1024 GB<br>• Data disk: 32768 GB | • System disk: 1024 GB<br>• Data disk: 32768 GB | • System disk: 1024 GB<br>• Data disk: 32768 GB |

  

| Parameter | Ultra-high I/O | General Purpose SSD | High I/O | Common I/O (Previous Generation Product) |
| --- | --- | --- | --- | --- |
| Short description | Highest performance disks excellent for enterprise mission-critical services as well as workloads demanding high throughput and low latency | Cost-effective disks designed for high-throughput, low-latency enterprise office applications | Disks suitable for commonly accessed workloads | Disks suitable for less commonly accessed workloads |
| Typical application scenarios | • Read/write-intensive applications that require ultra-large bandwidth<br>• Transcoding services<br>• I/O-intensive applications<br>– NoSQL<br>– Oracle<br>– SQL Server<br>– PostgreSQL<br>• Latency-sensitive applications<br>– Redis<br>– Memcache | Mainstream high-performance, low-latency interactive applications<br>• Enterprise office applications<br>• Large-scale development and testing<br>• Transcoding services<br>• Web server logs<br>• High-performance system disks, like container disks | Common workload applications<br>• Common development and testing | Applications demanding large capacity, medium read/write speed, and fewer transactions<br>• Common office applications<br>• Lightweight development and testing<br>• Not recommended to be used as system disks |
| Max. IOPS[a] | 33,000 | 20,000 | 5,000 | 2,200 |
| Max. throughput[a] | 350 MB/s | 250 MB/s | 150 MB/s | 50 MB/s |

| Parameter | Ultra-high I/O | General Purpose SSD | High I/O | Common I/O (Previous Generation Product) |
| --- | --- | --- | --- | --- |
| Formula used to calculate the disk throughput | Throughput = Min. (350, 120 + 0.5 × Capacity) MB/s | Throughput = Min. (250, 100 + 0.5 × Capacity) MB/s | Throughput = Min. (150, 100 + 0.15 × Capacity) MB/s | Throughput = 50 MB/s |
| Burst IOPS limit[a] | 16,000 | 8,000 | 5,000 | 2,200 |
| Formula used to calculate the disk IOPS | IOPS = Min. (33,000, 1,500 + 50 x Capacity) | IOPS = Min. (20,000, 1,500 + 8 x Capacity) | IOPS = Min. (5,000, 1,200 + 6 x Capacity) | IOPS = Min. (2,200, 500 + 2 x Capacity) |
| Single-queue access latency | 1 ms | 1 ms | 1 ms to 3 ms | 5 ms to 10 ms |
| API Name[b] | SSD | GPSSD | SAS | SATA |

NOTE

a: The maximum IOPS, maximum throughput, and burst IOPS limit are all calculated based on the sum of read and write operations.

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

EVS disk performance is closely related with the data block size. According to the formula, a large-capacity disk can achieve either the maximum IOPS or maximum throughput depending on which metric is reached first. When one has been reached, the other cannot increase any more.

- For data blocks of a small size, such as 4 KB or 8 KB, the disk performance can reach the maximum IOPS.
- For data blocks of a large size, greater than or equal to 16 KB, the disk performance can reach the maximum throughput.

The following uses an ultra-high I/O disk as an example. According to the formula, when the capacity of an ultra-high I/O disk is greater than or equal to 630 GB, the

  

disk can either reach the maximum IOPS 33,000 or the maximum throughput 350 MB/s. However, this is not the case in practice. For details, see **Table 2-2**.

**Table 2-2** Ultra-high I/O EVS disk maximum performance

| Data Block Size (KB) | Max. IOPS | Max. Throughput (MB/s) |
| --- | --- | --- |
| 4 | About 33,000 | About 130 |
| 8 | About 33,000 | About 260 |
| 16 | About 22,400 | About 350 |
| 32 | About 11,200 | About 350 |

## Description of the Disk IOPS Calculation Formula

Disk IOPS = Min. (Maximum IOPS, Baseline IOPS + IOPS per GB x Capacity)

The following example uses an ultra-high I/O EVS disk with a maximum IOPS of 33,000.

- If the disk capacity is 100 GB, the disk IOPS is calculated as follows:

  Disk IOPS = Min. (33,000, 1,500 + 50 x 100)

  Compare 33,000 and 6,500 and obtain the smaller value, which is 6,500. Therefore, the disk IOPS is 6,500.
- If the disk capacity is 1,000 GB, the disk IOPS is calculated as follows:

  Disk IOPS = Min. (33,000, 1,500 + 50 x 1,000)

  Compare 33,000 and 51,500 and obtain the smaller value, which is 33,000. Therefore, the disk IOPS is 33,000.

## EVS Disk Burst Capability and Principles

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. The IOPS limit indicates the performance of a single disk.

The burst capability is suitable for server startup scenarios. Normally, system disks have small capacities. For example, if a 50-GB ultra-high I/O disk does not have the burst capability, its IOPS limit can reach only 4,000 (1,500 + 50 x 50). However, if the disk has the burst capability, its IOPS limit can reach up to 16,000, thus improving the server startup speed.

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 GB, the disk IOPS limit is 6,500. In this case, the disk maximum IOPS can reach 16,000 in a certain duration.
- If the disk capacity is 1,000 GB, the disk IOPS limit is 33,000. In this case, the disk IOPS limit already exceeds its IOPS burst limit, and the disk does not need the burst capability.

The burst IOPS consumption and reservation principles are described as follows:

The burst capability is implemented based on a token bucket. 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, the fixed burst duration is 1800s, and a 100-GB ultra-high I/O EVS disk is used. Therefore, the number of initial tokens is 28,800,000 (1,800 x 16,000).

- Token production rate: This rate equals the disk IOPS limit, that is, 6,500 tokens/s.
- Token consumption rate: This rate is calculated 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 between the disk burst IOPS and IOPS limit.

Consumption principle

When the token consumption rate is greater than the production rate, the number of tokens decreases accordingly, and eventually the disk IOPS will be consistent with the token production rate, that is, the IOPS limit. In this example, the burst duration that the disk can sustain is approximately 3032s [28,800,000/(16,000 - 6,500)].

Reservation principle

When the token consumption rate is smaller than the production rate, the number of tokens increases accordingly, enabling the disk to regain the burst capability. In this example, if the disk is suspended for approximately 4431s (28,800,000/6,500), the token bucket will be filled up with tokens.

NOTE

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

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

- When the number of tokens is greater than zero, the disk IOPS can exceed 6,500 and has the capability to reach 16,000, the IOPS burst limit.
- When the number of tokens is zero, the disk does not have the burst capability, and the maximum IOPS is 6,500.
- When the actual IOPS is less than 6,500, the number of tokens starts to increase, and the disk can have the burst capability again.

  

**Figure 2-1** Principles of the burst capability

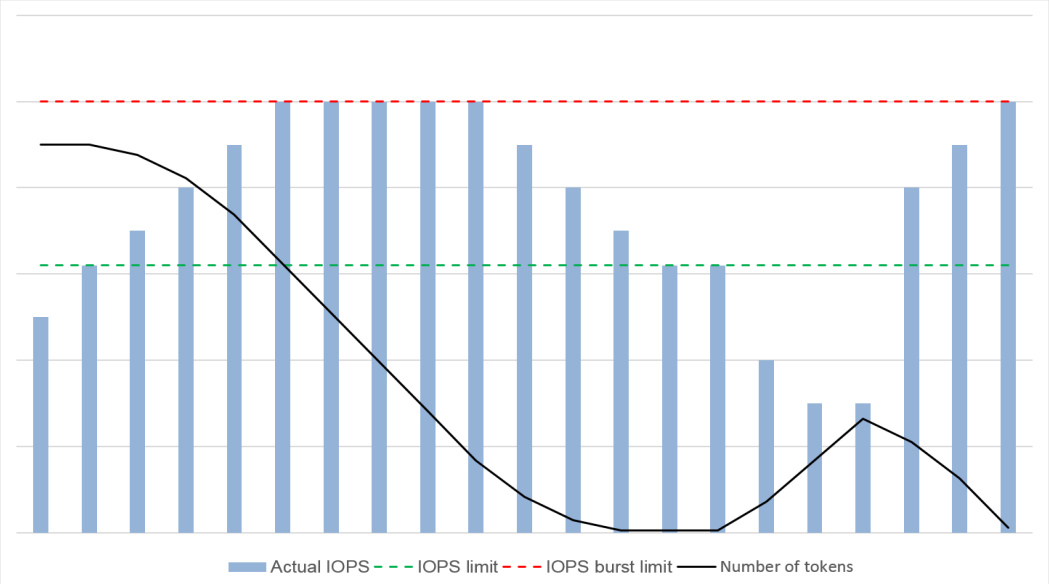


## Performance Test Method

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

# 3 Device Types and Usage Instructions

## What Are Device Types?

EVS device types are classified as Virtual Block Device (VBD) and Small Computer System Interface (SCSI) based on whether advanced SCSI commands are supported.

- VBD is the default EVS disk 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 also support advanced SCSI commands.

## Common Application Scenarios and Usage Instructions of SCSI EVS Disks

- SCSI EVS disks: BMSs support only SCSI EVS disks, which can be used as either system disks or data 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 the usage of 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. For more information about shared EVS disks, see **Shared EVS Disks and Usage Instructions**.

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

To use SCSI EVS disks, you need to install a driver for certain server OSs. The details are as follows:

- BMS

  Both the Windows and Linux images for BMSs are pre-installed with the required driver, that is, the SDI card driver. Therefore, no driver needs to be installed.
- KVM ECS

  When using SCSI EVS disks, you are advised to use them with KVM ECSs. Linux images for KVM ECSs already have the required driver built in the Linux kernel, and Windows images for KVM ECSs are also included with the 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 limitations in the driver support, you are advised not to use SCSI EVS disk with Xen ECSs.

  However, there are a few images available that support SCSI EVS disks on Xen ECSs. For the supported images, see **Table 3-1**.

  NOTE

  After confirming that the OS images of Xen ECSs support SCSI EVS disks, determine whether the driver needs to be installed based on the following conditions:

  - 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 so that 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**.

**Table 3-1** OSs supporting SCSI EVS disks

| Virtualization Type | OS | |
|---|---|---|
| XEN | Windows | See the Windows images listed on the **Public Images** page.<br>For details about how to view the information: Log in to the management console, choose **Image Mgmt Service**, click the **Public Images** tab, and select **ECS system disk 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.)<br>• SUSE Linux Enterprise Server 12 64bit (The kernel version is 3.12.51-52.31-default.)<br>• SUSE Linux Enterprise Server 12 SP1 64bit (The kernel version is 3.12.67-60.64.24-default.)<br>• SUSE Linux Enterprise Server 12 SP2 64bit (The kernel version is 4.4.74-92.35.1-default.) |

# 4 Shared EVS Disks and Usage Instructions

## What Are 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 feature multiple attachments, high-concurrency, high-performance, and high-reliability. They are usually used for enterprise key applications that require cluster deployment and high availability (HA) cluster. 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 Elastic Cloud Servers (ECSs) and Bare Metal Servers (BMSs). To implement file sharing, you need to deploy a shared file system or a cluster management system, such as Windows MSCS, Veritas VCS, or CFS.

**NOTICE**

To use shared EVS disks, you must set up a shared file system or similar cluster management system. If you directly attach EVS disks to multiple servers, the EVS disks cannot be shared and data may be overwritten.

  

**Figure 4-1** Application scenario of shared EVS disks

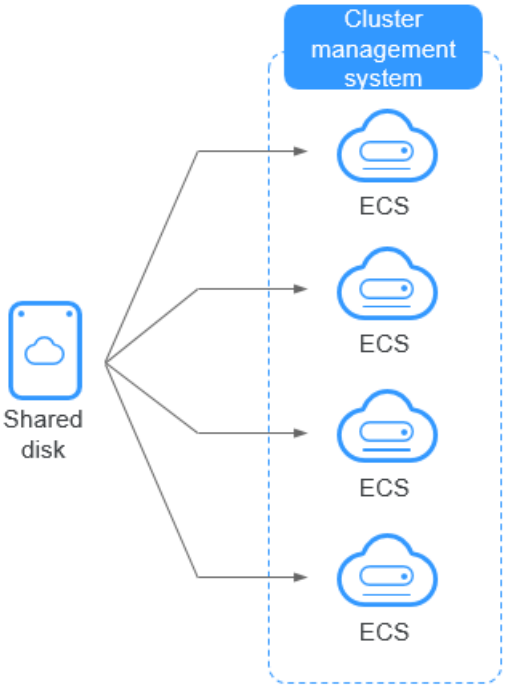


## Usage Precautions

Most common clusters, such as Windows MSCS and Veritas VCS and CFS, require SCSI reservations. Therefore, you are advised to use shared SCSI EVS disks for clusters. If a SCSI EVS disk is attached to a Xen ECS for use, you must install the driver. For details, see **Device Types and Usage Instructions**.

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

- Shared VBD EVS disks: The device type of a newly created shared EVS 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 EVS disks: These EVS 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 EVS 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 EVS 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 read/write operations 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. Therefore, 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: When multiple servers concurrently access a shared ultra-high I/O EVS disk, random read/write IOPS can reach up to 160,000.
- High-reliability: Shared EVS disks support both manual and automatic backup, delivering highly reliable data storage.
- Wide application scenarios: Shared EVS disks can be used for Linux RHCS clusters where only VBD EVS disks are needed. Whereas, they can also be used for Windows MSCS and Veritas VCS clusters that require SCSI reservations.

## Specifications and Performance

The specifications and performance of shared EVS disks are the same as those of non-shared EVS disks. For details, see **Disk Types and Disk Performance**.

## Data Sharing Principle 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 the data read/write rules, such as the read/write sequence and meaning, between these servers are not defined, data read/write interference between servers 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. Therefore, 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 the 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 the 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 EVS disk to multiple servers, files cannot be shared between the servers as shared EVS disks do not have the cluster capability. Therefore, build a shared file system or deploy a cluster management system if you need to share files between servers.

# 5 EVS Disk Encryption

## What Is EVS Disk Encryption

In case your services require encryption for the data stored on EVS disks, EVS provides you with the encryption function. You can encrypt newly created EVS disks.

EVS uses the industry-standard XTS-AES-256 encryption 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. Therefore, 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.

## Keys Used for EVS Disk Encryption

The keys provided by KMS include a Default Master Key and Customer Master Keys (CMKs).

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

  The Default Master Key cannot be disabled and does not support scheduled deletion.

- CMKs: Keys created by users. You may use existing CMKs or create new CMKs to encrypt disks. For details, see **Key Management Service** > **Creating a CMK** in the *Data Encryption Workshop User Guide*.

If disks are encrypted using CMKs and a CMK is then disabled or scheduled for deletion, the disks encrypted by this CMK can no longer be read from or written to and data on these disks may never be restored. See **Table 5-1** for more information.

  

**Table 5-1** Impact of CMK unavailability

| CMK Status | Impact | How to Restore |
| --- | --- | --- |
| Disabled | • If an encrypted disk is attached to a server, the disk cannot be accessed or data on the disk cannot be restored after a period of time or even permanently. If this disk is detached later, it cannot be attached again.<br>• If an encrypted disk is not attached to any server, it cannot be attached any more. | Enable the CMK. For details, see **Enabling One or More CMKs**. |
| Scheduled deletion | | Cancel the scheduled deletion for the CMK. For details, see **Canceling the Scheduled Deletion of One or More CMKs**. |
| Deleted | | Data on the disks can never be restored. |

**NOTICE**

You will be charged for the CMKs you use. If basic keys are used, ensure that your account balance is sufficient. If professional keys are used, renew your order timely. Otherwise, your services may be interrupted and your data may never be restored as the encrypted disks become unreadable and unwritable.

## Relationships Among Encrypted Disks, Backups, Images, and Snapshots

The encryption function can be used to encrypt system disks, data disks, images, and snapshots. The details are as follows:

- System disk encryption relates to the image that is used to create the server.
  - If an encrypted image is used to create the server, the encryption function is enabled for the system disk by default, and the system disk and image share the same encryption method. For details, see **Managing Private Images** > **Encrypting Images** in the *Image Management Service User Guide*.
- When creating an empty disk, you can determine whether to encrypt the disk or not. The disk encryption attribute cannot be changed after the disk has been created.
- If a disk is created from a snapshot, the encryption attribute of the disk will be the same as that of the snapshot's source disk.
- If a disk is created from a backup, the encryption attribute of the disk will be the same as that of the backup's source disk.
- If a disk is created from an image, the encryption attribute of the disk will be the same as that of the image's source disk.
- If a backup is created for a disk, the encryption attribute of the backup is the same as that of the disk.
- If a snapshot is created for a disk, the encryption attribute of the snapshot is the same as that of the disk.

  

## Who Can Use the Disk Encryption Function?

- The security administrator (having the Security Administrator rights) can grant the KMS access rights to EVS for using the disk encryption function.
- When a common user who does not have the Security Administrator rights needs to use the disk encryption function, the condition varies depending on whether the user is the first one ever in the current region or project to use this feature.
  - If the common user is the first one ever in the current region or project to use the feature, the user must contact a user having the Security Administrator rights to grant the KMS access rights to EVS. Then, the common user can use the disk encryption function.
  - If the common user is not the first one ever in the current region or project to use the feature, the common user can use the disk encryption function directly.

From the perspective of a tenant, as long as the KMS access rights have been granted to EVS in a region, all the users in the same region can directly use the disk encryption function.

If there are multiple projects in the current region, the KMS access rights need to be granted to each project in this region.

## Application Scenarios of EVS Disk Encryption

Figure 5-1 shows the user relationships under regions and projects from the perspective of a tenant. The following example uses region B to describe the two application scenarios of the disk encryption function.

**Figure 5-1** User relationships

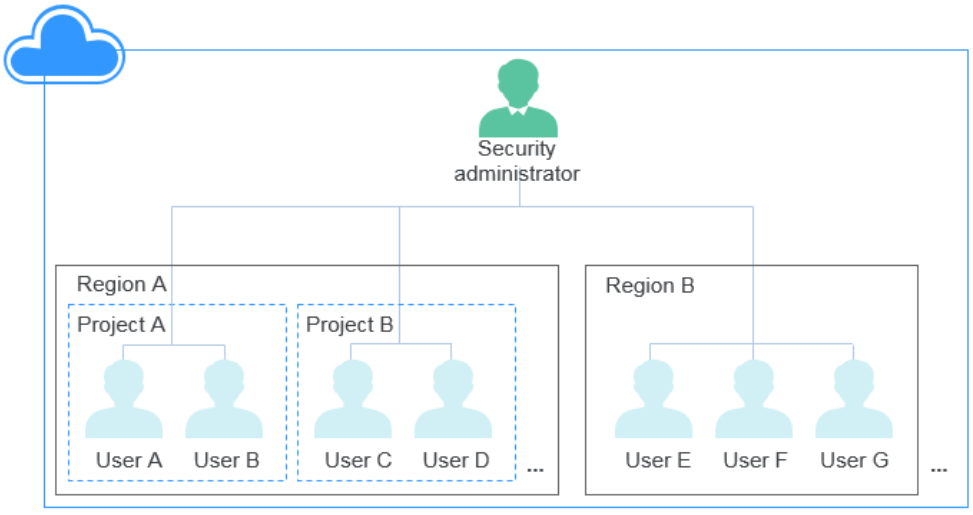


- If the security administrator uses the encryption function for the first time ever, the operation process is as follows:
  a. Grant the KMS access rights to EVS.

After the KMS access rights have been granted, the system automatically creates a Default Master Key and names it **evs/default**. DMK can be used for disk encryption.

NOTE

The EVS disk encryption relies on KMS. When the encryption function is used for the first time ever, the KMS access rights need to be granted to EVS. After the KMS access rights have been granted, all users in this region can use the encryption function, without requiring the KMS access rights to be granted again.

b. Select a key.

You can select one of the following keys:

- DMK: **evs/default**
- CMKs: Existing or newly created CMKs. For details, see **Creating a CMK**.

After the security administrator has used the disk encryption function, all users in Region B can directly use the encryption function.

- If User E (common user) uses the encryption function for the first time ever, the operation process is as follows:

  a. When user E uses the encryption function, and the system prompts a message indicating that the KMS access rights have not been granted to EVS.

  b. Contact the security administrator to grant the KMS access rights to EVS.

  After the KMS access rights have been granted to EVS, User E as well as all users in Region B can directly use the disk encryption function and do not need to contact the security administrator to grant the KMS access rights to EVS again.

# 6 EVS Disk Backup

## What Is EVS Disk Backup

The disk backup function provided by Cloud Backup and Recovery (CBR) allows you to create backups for your EVS disks. During the backup, you do not need to stop the server. When data loss or data damage occurred due to virus invasions, misoperations, or software and hardware faults, you can use backups to restore the data, guaranteeing your data correctness and security.

For details about EVS disk backup, see the *Cloud Backup and Recovery User Guide*.

## Backup Principle

For details about the disk backup principle, see **CBR Service Overview**.

NOTE

When a backup is created for a disk, the system automatically creates a snapshot, and the snapshot name starts with **autobk_snapshot_vbs_**. Only the snapshot automatically created during the latest backup is retained.

## Application Scenarios

After a backup policy has been set, the EVS disk data can be automatically backed up based on the policy. You can use the backups created on a timely basis as the baseline data to create new EVS disks or to restore the backup data to EVS disks.

## Usage Instructions

For details about the usage instructions of disk backups, see **Cloud Backup and Recovery User Guide**.

  

# 7 EVS Snapshot (OBT)

## What Is EVS Snapshot

EVS allows you to create snapshots for disks on the management console or by making API calls. An EVS snapshot is a complete copy or image of the disk data at a specific time point. As a major disaster recovery (DR) approach, you can use a snapshot to completely restore the data to the time point when the snapshot was created.

EVS 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.

## Differences Between Snapshots and Backups

- Both snapshots and backups are key approaches for data disaster recovery, but they use different storage plans.
  - The snapshot data is stored with the disk data so that you can rapidly back up and restore the disk data using snapshots.
  - The backup data is stored in the Object Storage Service (OBS). If the disk data is damaged, you can restore the data using backups.
- EVS snapshot does not support automatic creation, whereas backup does. You can set a backup policy, and the system will automatically back up the disk data according to this policy.

## Snapshot Principle

Snapshots and backups are different in that a backup saves the data as another copy in the storage system other than on the disk, whereas a snapshot establishes a relationship between the snapshot and disk data.

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

1. Create disk v1, which contains no data.
2. Write data d1 and d2 to disk v1. Data d1 and d2 are written to new spaces.

  

3. Create snapshot s1 for disk v1 that is modified in **2**. Data d1 and d2 are not saved as another copy elsewhere. Instead, the relationship between snapshot s1 and data d1 and d2 is established.
4. Write data d3 to disk v1 and change data d2 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. Therefore, snapshot s1 can be used to restore data if needed.
5. Create snapshot s2 for disk v1 that is modified in **4**. The relationship between s2 and data d1, d3, and d4 is established.

**Figure 7-1** Snapshot principle

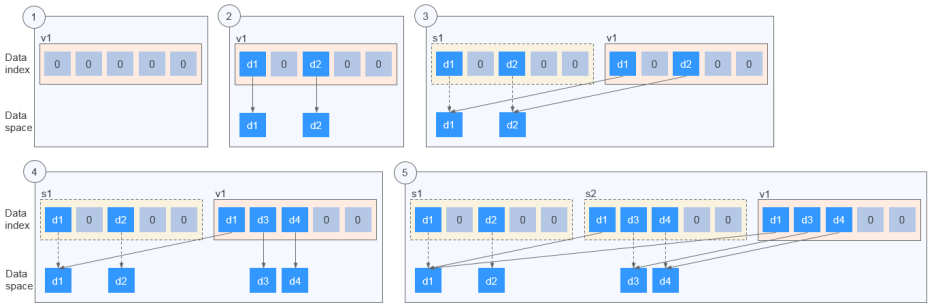


## Application Scenarios

The snapshot function helps address your following needs:

- Routine data backup

  You can create snapshots for disks on a timely basis and use snapshots to recover your data in case that data loss or data inconsistency occurred due to misoperations, viruses, or attacks.

- Rapid data restoration

  You can create a snapshot or multiple snapshots before an application software upgrade or a service data migration. If an exception occurs during the upgrade or migration, service data can be rapidly restored to the time point when the snapshot was created.

  For example, a fault occurred on system disk A of server A, and therefore server A cannot be started. Because system disk A is already faulty, the data on system disk A cannot be restored by rolling back snapshots. However, you can create disk B using an existing snapshot of system disk A and attach disk B to a properly running server, for example server B. In this case, server B can read the data of system disk A from disk B.

  **NOTE**

  Currently, when rolling back data from snapshots, the snapshot data can be rolled back only to its source EVS disk, and a rollback to another EVS disk is not possible.

- Multi-service quick deployment

  You can use a snapshot to create multiple disks containing the same initial data, and these disks can be used as data resources for various services, for example data mining, report query, and development and testing. This method protects the initial data and creates disks rapidly, meeting the diversified service data requirements.

  

## Charging Standards During OBT

The EVS snapshot function is currently in Open Beta Test (OBT), and you can use the function for free. The function will be charged after commercial use. The commercial use time and charging standards will be notified later.

During the OBT, the function adopts a limited free trial policy. That is, you can use the snapshot function for free, but the number of snapshots you can create is limited.

- Snapshot quota requirements
  - A maximum of 7 snapshots can be created for a disk.
  - The total number of snapshots that can be created by a user is calculated by the total number of disks multiplying seven. This total number includes both system disks and data disks.

  Once the snapshot quantity has exceeded the snapshot quota, new snapshots cannot be created. For example, a user who has five disks can create a maximum of 35 snapshots.
- Snapshot retention policy

  The system does not automatically delete user snapshots. A snapshot can be deleted in either of the following ways:
  - A user deletes the snapshot.
  - A user deletes a disk so that all snapshots created for this disk are also deleted.

NOTE

When a backup is created for a disk, the system automatically creates a snapshot, and the snapshot name starts with **autobk_snapshot_vbs_**. Only the snapshot automatically created during the latest backup is retained.

You can only view details of this snapshot but cannot perform any operations on it.

## Usage Instructions

For details about the snapshot usages, see **Creating a Snapshot (OBT)**.

  

# 8 Differences Between EVS Disk Backup and EVS Snapshot

Both EVS disk backup and EVS snapshot provide redundancies for the EVS disk data to improve reliability. **Table 8-1** lists the differences between them.

**Table 8-1** Differences between backups and snapshots

| Item | Storage Solution | Data Synchronization | DR Range | Service Recovery |
| --- | --- | --- | --- | --- |
| Backup | Backup data is stored in OBS, instead of disks. This ensures data restoration upon disk data loss or corruption. | A backup is the data copy of a disk at a given point in time. CBR supports automatic backup by configuring backup policies. Deleting a disk will not clear its backups. | A backup and its source disk reside in the same AZ. | Data can be recovered and services can be restored by restoring the backup data to original disks or creating new disks from backups, ensuring excellent data reliability. |

  

| Item | Storage Solution | Data Synchronization | DR Range | Service Recovery |
| --- | --- | --- | --- | --- |
| Snapshot | Snapshot data is stored with disk data.<br>**NOTE**<br>Creating a backup requires a certain amount of time because data needs to be transferred. Therefore, 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. When you delete the disk, the snapshots of the disk are also deleted. | A snapshot and its source disk reside in the same AZ. | You can use a snapshot to roll back its original disk or create a disk for data restoration and service recovery. |

# 9 EVS Three-Copies of Data Mechanism

## What Is the Three-Copies of Data Mechanism?

The backend storage system of EVS employs the three-copies of data mechanism to guarantee data reliability. In this mechanism, one piece of data is by default divided into multiple 1 MB 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.

The three-copies of data mechanism has the following characteristics:

- The storage system saves the data copies on different disks of different servers, ensuring that services are not interrupted in case that a physical device fails.
- The storage system guarantees tight 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 9-1** Three-copies of data mechanism

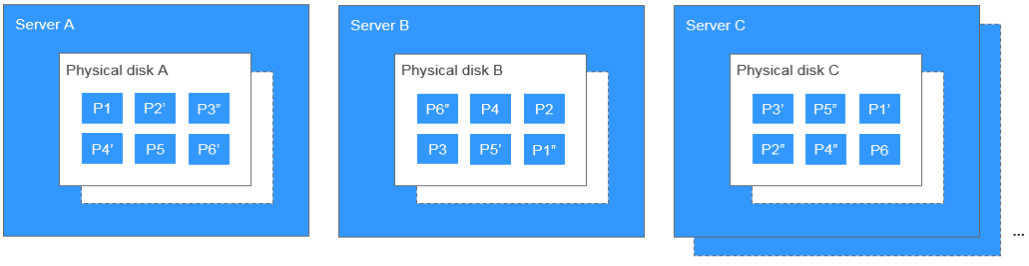


## How Does the Three-Copies of Data Mechanism Keep Data Consistency?

Data consistency includes the following two aspects: 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.

  

The three copies of data mechanism 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 its copies.
- Storage system automatically restores the damaged copy in case 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 mechanism ensures the correct number of data copies and data consistency among data copies.

## How Does the Three-Copies of Data Mechanism 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 9-2** shows the data rebuild process.

  

**Figure 9-2** Data rebuild process

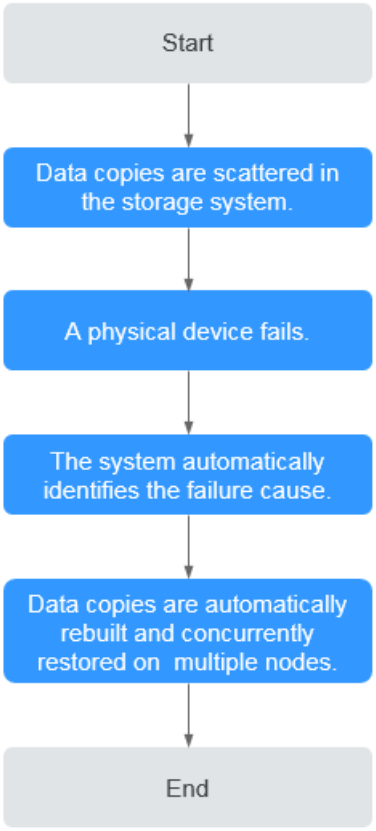


Figure 9-3 shows the data rebuild principle. For example, if physical disks on server F are faulty, the data blocks on these physical disks will be rebuilt on the physical disks of other servers.

  

**Figure 9-3** Data rebuild principle


## What Are the Differences Between Three-Copies of Data, EVS Disk Backup, and EVS Snapshot?

The three-copies of data mechanism improves the reliability of the data stored on EVS disks. It is used to tackle data loss or inconsistency caused by physical device faults.

Whereas, EVS disk backup and EVS snapshot are used to prevent data loss or inconsistency caused by misoperation, viruses, or hacker attacks. Therefore, you are advised to create backups and snapshots to back up the EVS disk data on a timely basis.

  

# 10 Billing

## 10.1 Billing for Disks

### Billing Items

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

- Billing start: You will be billed for the EVS disks right after you have purchased them, regardless of whether they are attached or not.
- Billing end:
  - For a yearly/monthly disk, the billing ends after the disk is successfully unsubscribed from, and the fee to be refunded is calculated as follows: Refund amount = Your actual payment - Amount due - Handling fees. For more information, see **How Do I View the Unsubscription Fee? How Is It Calculated?**
  - For a pay-per-use disk, the billing ends after the disk is successfully deleted.

### Billing Modes

EVS disks are billed by disk capacity. An EVS disk can be billed on a yearly/monthly or pay-per-use basis.

- Yearly/Monthly: prepaid.
- Pay-per-use: postpaid. EVS disks are billed by the second and settled by the hour. For a duration of less than an hour, the payment is based on the actual duration that the service was used for.

  

## Billing Involved in Configuration Modifications

| Item | Yearly/Monthly | Pay-per-Use |
| --- | --- | --- |
| Capacity change | • EVS does not support the reduction of disk capacities.<br>• EVS supports the expansion of disk capacities. Additional capacities need to be paid.<br>**NOTE**<br>The expiration time of the EVS disk remains unchanged after the capacity expansion. | • EVS does not support the reduction of disk capacities.<br>• EVS supports the expansion of disk capacities.<br>Multiple pieces of billing records will be generated within a billing cycle (an hour) when an expansion succeeded.<br>For example, if you expand the capacity of an EVS disk from 100 GB to 200 GB at 01:30:01, two pieces of billing records will be generated in billing cycle 01:00:00-02:00:00. One is the billing record generated for the 100 GB in 01:00:00-01:30:00, and the other is the billing record generated for the 200 GB in 01:30:01-02:00:00. |
| Billing mode change | EVS supports the billing mode change from pay-per-use to yearly/monthly.<br>For details, see **From Pay-per-Use to Yearly/Monthly**. | EVS supports the billing mode change from pay-per-use to yearly/monthly.<br>For details, see **From Yearly/Monthly to a Pay-per-Use**. |

# 11 Constraints

This topic describes the constraints on using EVS.

**Table 11-1** Constraints on using EVS

| Scenario | Item | Restrictions |
| --- | --- | --- |
| Disk creation | Maximum number of disks that can be created at a time | 100 |
| | Disk creation from snapshot | • The disk type of the new disk is the same as that of the snapshot's source disk.<br>• The device type of the new disk is the same as that of the snapshot's source disk.<br>• The encryption attribute of the new disk is the same as that of the snapshot's source disk.<br>• Batch creation is not supported. One can create only one disk from a snapshot at a time. |
| | Disk creation from backup | • Batch creation is not supported. One can create only one disk from a backup at a time.<br>• 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.<br>• If a disk is created from a backup of a system disk, the new disk can be used as a data disk only. |

  

| Scenario | Item | Restrictions |
| --- | --- | --- |
| | Disk creation from image | • The device type of the new disk is the same as that of the image's source disk.<br>• If a disk is created from an image, the encryption attribute of the disk will be 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 | Number of servers that a non-shared disk can be attached to | 1 |
| | Number of servers that a shared disk can be attached to | 16 |
| | Number of disks that can be attached to one server | 60 (system disk included) |
| Disk capacity expansion | Capacity expansion | Disk capacities can be expanded only, but cannot be reduced. |
| | Capacity expansion of non-shared disks | Some server OSs support the capacity expansion of non-shared, In-use disks.<br>For details, see **Expanding Capacity for an In-use EVS Disk**. |
| | 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 GB |
| Disk detachment | System disk detachment | A system disk can only be detached offline, that is, its 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 | Restrictions |
| --- | --- | --- |
| Disk deletion | Deletion of pay-per-use disks | • The disk status is **Available**, **Error**, **Expansion failed**, **Restoration failed**, or **Rollback failed**.<br>• The disk is not added to any replication pair in SDRS. If the disk has been added to a replication pair, **delete the replication pair** and then delete the disk.<br>• The disk is not locked by any service. |
| | Deletion of yearly/monthly disks | Yearly/monthly disks cannot be deleted right way. One can only unsubscribe from such disks.<br>For details about the unsubscription rules and operation methods, see **Billing Center User Guide**. |
| Disk capacity | Maximum capacity of a system disk | • Common I/O: 1024 GB<br>• High I/O: 1024 GB<br>• General Purpose SSD: 1024 GB<br>• Ultra-high I/O: 1024 GB<br>• Extreme SSD: 1024 GB |
| | Maximum capacity of a data disk | • Common I/O: 32768 GB<br>• High I/O: 32768 GB<br>• General Purpose SSD: 32768 GB<br>• Ultra-high I/O: 32768 GB<br>• Extreme SSD: 32,768 GB |
| | Maximum capacity supported by the MBR partition style | 2 TB |
| | Maximum capacity supported by the GPT partition style | 18 EB |
| Disk performance | Common I/O | • Maximum IOPS per disk: 2,200<br>• Maximum throughput per disk: 50 MB/s |
| | High I/O | • Maximum IOPS per disk: 5,000<br>• Maximum throughput per disk: 150 MB/s |
| | General Purpose SSD | • Maximum IOPS per disk: 2,000<br>• Maximum throughput per disk: 250 MB/s |

| Scenario | Item | Restrictions |
| --- | --- | --- |
| | Ultra-high I/O | • Maximum IOPS per disk: 33,000<br>• Maximum throughput per disk: 350 MB/s |
| | Extreme SSD | • Maximum IOPS per disk: 128,000<br>• Maximum throughput per disk: 1,000 MB/s |
| Snapshot | Maximum number of snapshots that can be created for a disk | 7 |
| | Maximum number of disks that can be created from a snapshot | 128 |
| | Snapshot data rollback to disk | • A snapshot can be rolled back only to its source EVS disk. A rollback to another disk is not possible.<br>• A snapshot can be rolled back only when the snapshot status is **Available** and the source disk status is **Available** (not attached to any server) or **Rollback failed**.<br>• When a backup is created for a disk, the system automatically creates a snapshot, and the snapshot name starts with **autobk_snapshot_vbs_**. Only the snapshot automatically created during the latest backup is retained. This snapshot can be viewed only, but cannot be used to roll back the disk data. |
| Tag | Maximum number of tags that can be added for a disk | 10 |

  

# 12 EVS and Other Services

Figure 12-1 shows the relationships between EVS and other services.

**Figure 12-1** Relationships between EVS and other services

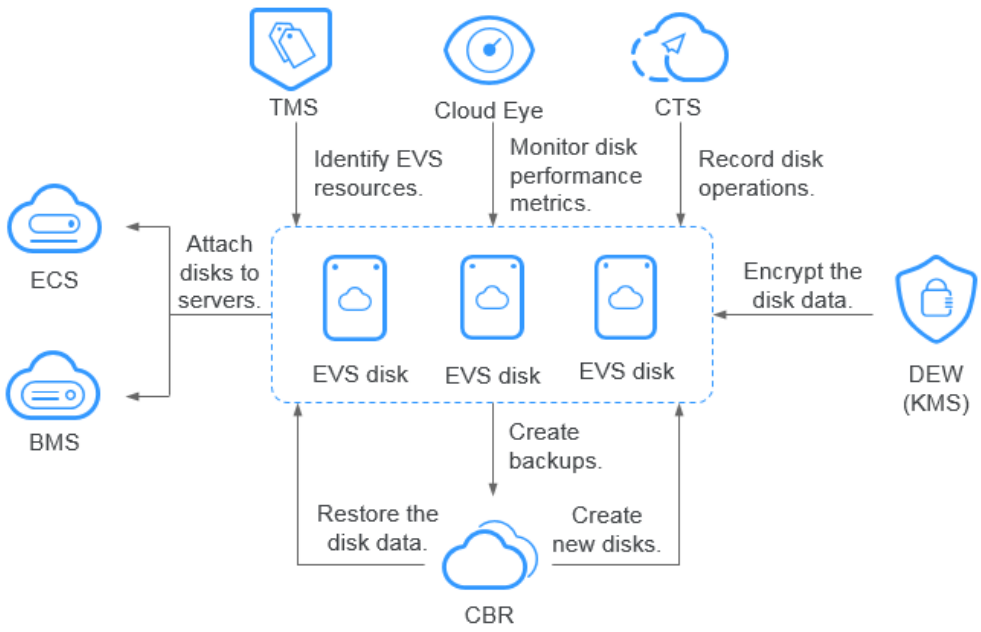


**Table 12-1** Related services

| Interactive Function | Related Service | Reference |
| --- | --- | --- |
| EVS disks can be attached to ECSs and used as scalable block storage devices. | Elastic Cloud Server (ECS) | • Attaching a Non-shared Disk<br>• Attaching a Shared Disk |

  

| Interactive Function | Related Service | Reference |
| --- | --- | --- |
| SCSI EVS disks can be attached to BMSs and used as scalable block storage devices. | Bare Metal Server (BMS) | • **Attaching a Non-shared Disk**<br>• **Attaching a Shared Disk** |
| Backups can be created for EVS disks to guarantee the reliability and security of the server data. | Cloud Backup and Recovery (CBR) | • **EVS Disk Backup**<br>• **Managing EVS Backup** |
| EVS disk encryption depends on the KMS service in DEW. Keys provided by KMS can be used to encrypt EVS disks (both system and data disks), thus improving EVS disk data security. | Data Encryption Workshop (DEW) | • **EVS Disk Encryption**<br>• **Managing an Encrypted EVS Disk** |
| After EVS is enabled, the performance metrics of monitored disks can be viewed through Cloud Eye without installing any additional plug-in. The performance metrics include Disk Read Rate, Disk Write Rate, Disk Read Requests, and Disk Write Requests. | Cloud Eye | **Viewing EVS Monitoring Data** |
| CTS records operations of EVS resources, facilitating user query, audit, and backtracking. | Cloud Trace Service (CTS) | **Querying EVS Traces** |
| Tags identify EVS resources for purposes of easy categorization and quick search. | Tag Management Service (TMS) | **Adding a Tag** |

# 13 Basic Concepts

## 13.1 Basic Concepts

**Table 13-1** Basic EVS concepts

| Concept | Description |
| --- | --- |
| IOPS | Number of read/write operations performed by an EVS disk per second |
| Throughput | Amount of data successfully transmitted by an EVS disk per second, that is, the amount of data read from and written into an EVS disk |
| 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. |

## 13.2 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 13-1** shows the relationship between regions and AZs.

**Figure 13-1** Regions and AZs

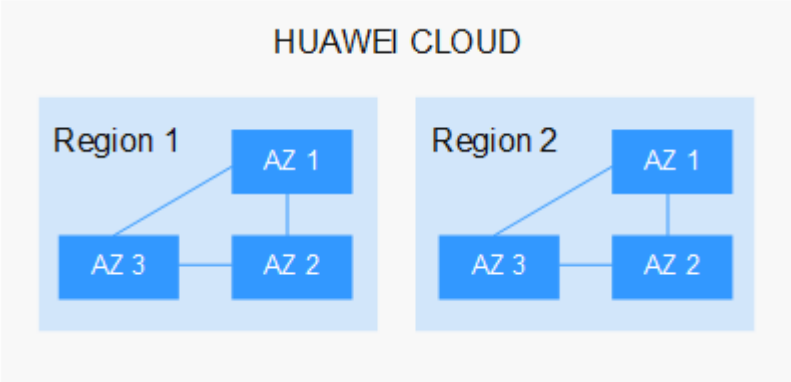


HUAWEI CLOUD provides services in many regions around the world. Select a region and AZ based on requirements. For more information, see **HUAWEI CLOUD Global Regions**.

## Selecting a Region

When selecting a region, consider the following factors:

- Location

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

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

    NOTE

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

  

- Resource price

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

## Selecting an AZ

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

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

## Regions and Endpoints

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

  

# A Change History

| Release Date | What's New |
| --- | --- |
| 2018-09-10 | This issue is the fourth official release, which incorporates the following change:<br>● Added section **EVS Three-Copies of Data Mechanism**. |
| 2018-07-30 | This issue is the third official release, which incorporates the following changes:<br>● Added content **Differences Between EVS, SFS, and OBS** in section **What Is Elastic Volume Service?**<br>● Added precautions for using shared EVS disks together with SCSI.<br>● Modified disk performance metrics. |
| 2018-06-30 | This issue is the second official release, which incorporates the following changes:<br>● Added section **Differences Between EVS Disk Backup and EVS Snapshot**.<br>● 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. |

  