

Elastic Volume Service 8.1.0

User Guide

Issue	01
Date	2021-09-30



Copyright © Huawei Technologies Co., Ltd. 2021. All rights reserved.

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

Trademarks and Permissions



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

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

Notice

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

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

Contents

1 Quick Start.....	1
1.1 Overview.....	1
1.2 Preparations.....	1
1.3 Applying for an EVS Disk.....	2
1.4 Attaching an EVS Disk.....	2
1.5 Initializing EVS Disks.....	3
1.5.1 Initializing a Windows Data Disk.....	3
1.5.2 Initializing a Linux Data Disk.....	8
2 Infographics.....	14
2.1 Introduction.....	15
2.2 Shared Disk.....	17
2.3 Changing the Disk Type.....	18
3 User Guide (for ECS).....	19
3.1 Introduction.....	19
3.1.1 What Is Elastic Volume Service?.....	19
3.1.2 Advantages.....	22
3.1.3 Application Scenarios.....	22
3.1.4 Implementation Principles.....	26
3.1.5 Related Services.....	29
3.1.6 Key Metrics.....	30
3.1.7 Restrictions.....	31
3.1.8 Accessing and Using the Cloud Service.....	45
3.2 Related Concepts.....	45
3.2.1 Device Type.....	45
3.2.2 Disk Type.....	47
3.2.3 EVS Disk Data Redundancy Protection.....	48
3.2.4 Shared Disk.....	50
3.2.5 EVS Disk Snapshot.....	54
3.2.6 EVS Disk Quota.....	60
3.2.7 Relationship Between the Disk Type, Backend Storage, and Storage Array.....	61
3.2.8 Mapping Between Mount Points and Device Names.....	62
3.2.9 Disk Migration.....	63

3.2.10 EVS Disk Performance.....	67
3.3 Operation Process.....	68
3.4 Applying for an EVS Disk.....	72
3.4.1 Applying for a Data Disk.....	72
3.4.2 Applying for a System Disk.....	76
3.4.3 Applying for an EVS Disk Using an Existing Disk.....	81
3.4.4 Applying for an EVS Disk Using a Snapshot.....	84
3.4.5 Applying for an EVS Disk Using a Backup.....	87
3.4.6 Applying for an EVS Disk Using an Image.....	91
3.5 Attaching an EVS Disk.....	94
3.6 Initializing a Data Disk.....	98
3.6.1 Initialization Overview.....	98
3.6.2 Initializing a Windows Data Disk.....	99
3.6.3 Initializing a Linux Data Disk (fdisk).....	104
3.6.4 Initializing a Linux Data Disk (parted).....	109
3.7 Expanding EVS Disk Capacity.....	113
3.7.1 Overview.....	113
3.7.2 Expanding Disk Capacity Online.....	115
3.7.3 Expanding Disk Capacity Offline.....	118
3.7.4 Operations After Expanding Disk Capacity in Windows.....	119
3.7.5 Operations After Expanding Disk Capacity in Linux (Adding Partitions Using fdisk).....	122
3.7.6 Operations After Expanding Disk Capacity in Linux (Adding Partitions Using parted).....	126
3.7.7 Operations After Expanding Disk Capacity in Linux (Replacing Original Partitions Using fdisk).....	130
3.7.8 Operations After Expanding Disk Capacity in Linux (Replacing Original Partitions Using parted)...	134
3.7.9 Adding a Data Disk.....	139
3.8 Releasing an EVS Disk.....	139
3.8.1 Detaching an EVS Disk.....	139
3.8.2 Deleting an EVS Disk.....	141
3.9 Managing Disk Types.....	143
3.9.1 Changing a Disk Type.....	143
3.10 Managing Snapshots.....	145
3.10.1 Applying for a Snapshot.....	145
3.10.2 Rolling Back an EVS Disk Using a Snapshot.....	147
3.10.3 Deleting a Snapshot.....	148
3.11 Managing Backups.....	149
3.11.1 Creating a Backup.....	149
3.11.2 Using a Backup to Restore Data.....	150
3.12 Managing EVS Disks.....	150
3.12.1 Extending the EVS Disk Validity Period.....	151
3.12.2 Viewing Monitoring Data.....	152
3.12.3 Changing the VDC Quota.....	152
3.13 FAQs.....	152

3.13.1 Can I Attach an EVS Disk to Multiple Instances?.....	152
3.13.2 How Can I Attach an EVS Disk to an ECS?.....	152
3.13.3 How Many States Does an EVS Disk Have?.....	152
3.13.4 Does an EVS Disk or Snapshot Generate Metering Information in All States?.....	155
3.13.5 Can I Change the EVS Disk Capacity?.....	155
3.13.6 What Can I Do If an EVS Disk Whose Disk Type Is Changed Remains in the Retyping State for a Long Time?.....	156
3.13.7 How Can I Test the Performance of a Windows EVS Disk?.....	157
3.13.8 How Can I Test the Performance of a Linux EVS Disk?.....	158
3.13.9 Logging In to the EVS Console as a VDC Administrator or VDC Operator.....	162
4 User Guide (for BMS).....	163
4.1 Introduction.....	163
4.1.1 What Is Elastic Volume Service?.....	163
4.1.2 Advantages.....	166
4.1.3 Application Scenarios.....	166
4.1.4 Implementation Principles.....	170
4.1.5 Related Services.....	173
4.1.6 Key Metrics.....	174
4.1.7 Restrictions.....	174
4.1.8 Accessing and Using the Cloud Service.....	181
4.2 Related Concepts.....	181
4.2.1 Device Type.....	181
4.2.2 Disk Type.....	181
4.2.3 EVS Disk Data Redundancy Protection.....	183
4.2.4 Shared Disk.....	185
4.2.5 EVS Disk Snapshot.....	188
4.2.6 EVS Disk Quota.....	194
4.2.7 Relationship Between the Disk Type, Backend Storage, and Storage Array.....	195
4.2.8 Mapping Between Mount Points and Device Names.....	196
4.2.9 EVS Disk Performance.....	197
4.3 Operation Process.....	198
4.4 Applying for an EVS Disk.....	202
4.4.1 Applying for a Data Disk.....	202
4.4.2 Applying for an EVS Disk Using an Existing Disk.....	205
4.4.3 Applying for an EVS Disk Using a Snapshot.....	208
4.4.4 Applying for an EVS Disk Using a Backup.....	210
4.5 Attaching an EVS Disk.....	214
4.6 Initializing a Data Disk.....	217
4.6.1 Initialization Overview.....	217
4.6.2 Initializing a Windows Data Disk.....	217
4.6.3 Initializing a Linux Data Disk (fdisk).....	222
4.6.4 Initializing a Linux Data Disk (parted).....	227

4.7 Expanding EVS Disk Capacity.....	232
4.7.1 Overview.....	232
4.7.2 Online Disk Capacity Expansion.....	233
4.7.3 Offline Disk Capacity Expansion.....	236
4.7.4 Post-Expansion Operations for a Windows EVS Disk.....	237
4.7.5 Operations After Expanding Disk Capacity in Linux (Adding Partitions Using fdisk).....	240
4.7.6 Operations After Expanding Disk Capacity in Linux (Adding Partitions Using parted).....	244
4.7.7 Operations After Expanding Disk Capacity in Linux (Replacing Original Partitions Using fdisk).....	248
4.7.8 Operations After Expanding Disk Capacity in Linux (Replacing Original Partitions Using parted)...	252
4.7.9 Adding a Data Disk.....	257
4.8 Releasing an EVS Disk.....	257
4.8.1 Detaching a Data Disk.....	257
4.8.2 Deleting Data Disks.....	259
4.9 Managing Disk Types.....	261
4.9.1 Changing the Disk Type.....	261
4.10 Managing Snapshots.....	262
4.10.1 Applying for a Snapshot.....	262
4.10.2 Rolling Back an EVS Disk Using a Snapshot.....	264
4.10.3 Deleting Snapshots.....	265
4.11 Managing Backups.....	266
4.11.1 Creating a Backup.....	266
4.11.2 Using a Backup to Restore Data.....	267
4.12 Managing EVS Disks.....	267
4.12.1 Extending the EVS Disk Validity Period.....	267
4.12.2 Viewing Monitoring Data.....	268
4.12.3 Changing the VDC Quota.....	269
4.13 FAQs.....	269
4.13.1 Can I Attach an EVS Disk to Multiple Instances?.....	269
4.13.2 How Can I Attach an EVS Disk to a BMS?.....	269
4.13.3 How Many States Does an EVS Disk Have?.....	269
4.13.4 Does an EVS Disk or Snapshot Generate Metering Information in All States?.....	272
4.13.5 Can I Change the EVS Disk Capacity?.....	272
4.13.6 Will Data in an EVS Disk Be Lost When the EVS Disk Is Detached?.....	273
4.13.7 How Can I Test the Performance of a Windows EVS Disk?.....	273
4.13.8 How Can I Test the Performance of a Linux EVS Disk?.....	274
4.13.9 Logging In to the EVS Console as a VDC Administrator or VDC Operator.....	277

1 Quick Start

- [1.1 Overview](#)
- [1.2 Preparations](#)
- [1.3 Applying for an EVS Disk](#)
- [1.4 Attaching an EVS Disk](#)
- [1.5 Initializing EVS Disks](#)

1.1 Overview

Elastic Volume Service (EVS) is a virtual block storage service, which provides block storage space for Elastic Cloud Servers (ECSs) and Bare Metal Servers (BMSs).

The procedure for applying for and using an EVS disk is as follows:

1. [Apply for an EVS disk.](#)
2. [Attach the EVS disk.](#)
3. [Initialize the EVS disk.](#)

This document describes how to apply for, attach, and initialize EVS disks. The following describes how to attach an EVS disk to an ECS instance.

For details about EVS, see **Operation Help Center > Storage > Elastic Volume Service**.

1.2 Preparations

Disk types are types and tags that can be selected during disk creation. A disk type corresponds to a backend storage device for a group of disks.


When applying for an EVS disk, you need to select a disk type. Therefore, you need to log in to ManageOne Operation Portal to create a disk type. For details, see "Storage Services" > "Elastic Volume Service" > "Configuration Before Applying for an EVS Disk" > "(Optional) Creating a Disk Type" in *HUAWEI CLOUD Stack 8.1.0 Resource Provisioning Guide*.

1.3 Applying for an EVS Disk

Step 1 Log in to ManageOne as a VDC administrator or operator using a browser.

URL in non-B2B scenarios: **https://Address for accessing ManageOne Operation Portal**, for example, **https://console.demo.com**

URL in B2B scenarios: **https://Address for accessing ManageOne Tenant Portal**, for example, **https://tenant.demo.com**

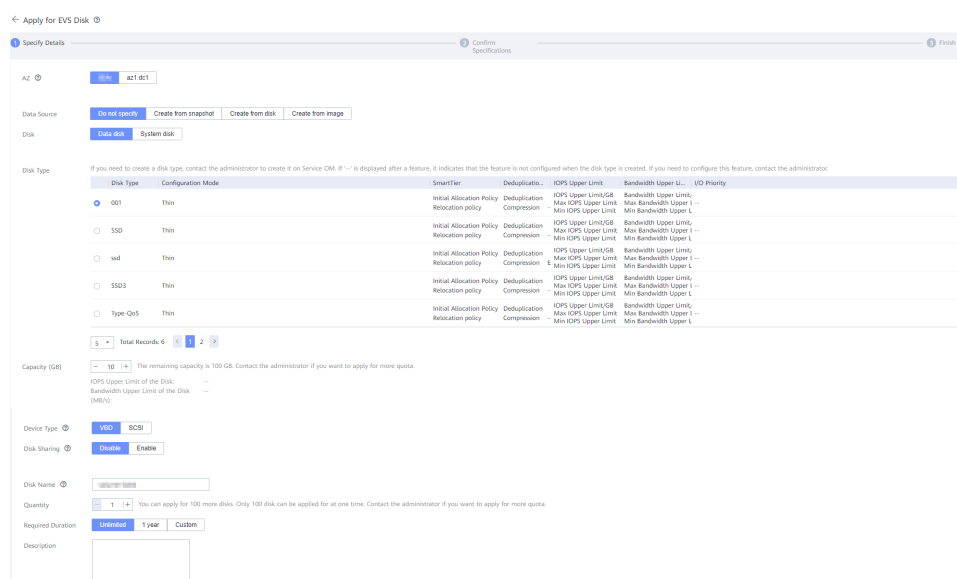
Step 2 Click  in the upper left corner and select a region. Choose **Storage > Elastic Volume Service**.

Step 3 Select a resource set from the menu bar.

Step 4 Click **Apply for EVS Disk**.

Step 5 In the **Select Service** dialog box, select the target service and then click **Apply Now**.

Step 6 In the **Apply for EVS Disk** dialog box, configure EVS disk information as prompted.



Step 7 Click **Next**.

Step 8 Confirm the order information and click **Apply Now**.

Step 9 After the state of the EVS disk changes to **Available**, the EVS disk is successfully created.

----End

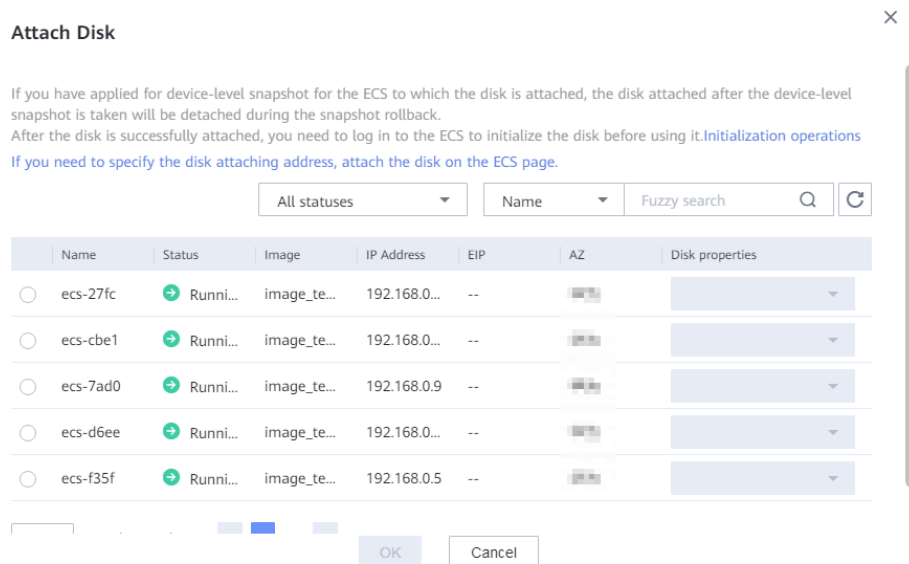
1.4 Attaching an EVS Disk

Step 1 Refer to [Step 1](#) to [Step 2](#) in section **Applying for an EVS Disk**.

Step 2 Click **Attach** in the **Operation** column where the EVS disk resides.

Step 3 In the displayed **Attach Disk** dialog box, select the instance to which the disk is to be attached. A data disk can only be attached to an instance as a data disk. A system disk can be attached to an instance as a data disk or a system disk.

To attach a shared disk, you can select multiple ECSs. A shared disk can be attached to a maximum of 16 ECSs by default.



Step 4 When the disk status changes to **In-use**, the disk is successfully attached to the instance.

----End

1.5 Initializing EVS Disks

1.5.1 Initializing a Windows Data Disk

Step 1 Log in to an ECS.

For details, see **Operation Help Center > Compute > Elastic Cloud Server > User Guide > Logging In to an ECS**.

Step 2 On your desktop, click **Start**, right-click **Computer** and choose **Manage** from the shortcut menu.

The **Computer Management** window is displayed.

Step 3 In the navigation pane, choose **Storage > Disk Management**.

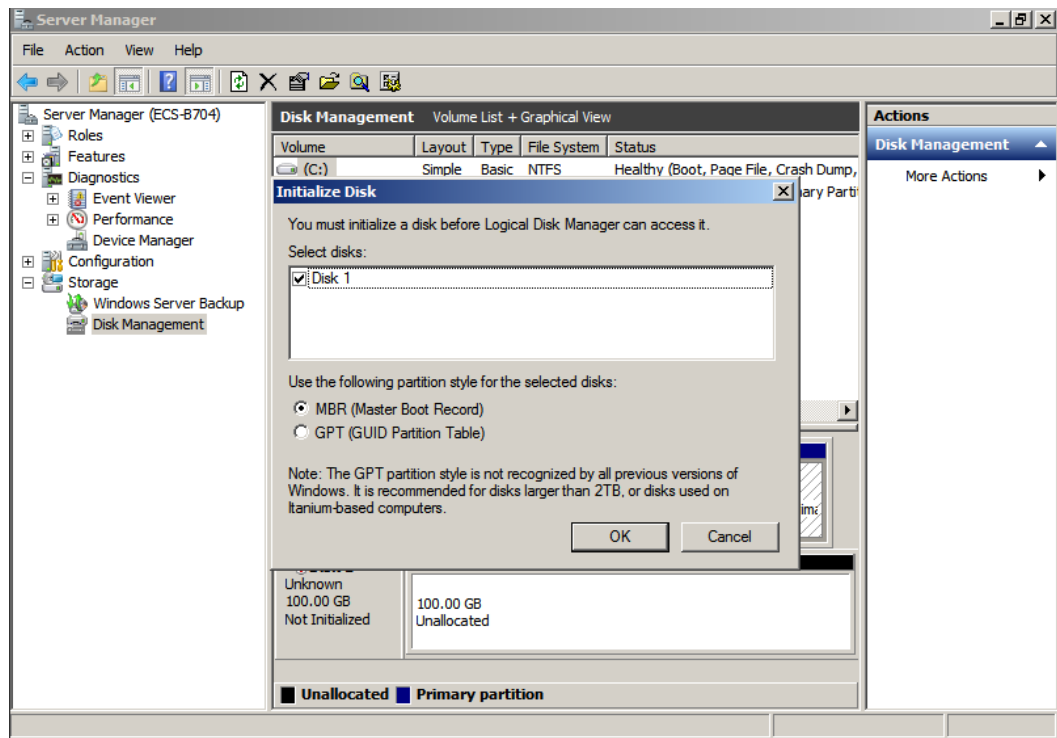
Step 4 If the disk to be initialized in the disk list is in **Offline** state, right-click on it and choose **Online** from the shortcut menu.

Then, the disk status should change from **Offline** to **Uninitialized**.

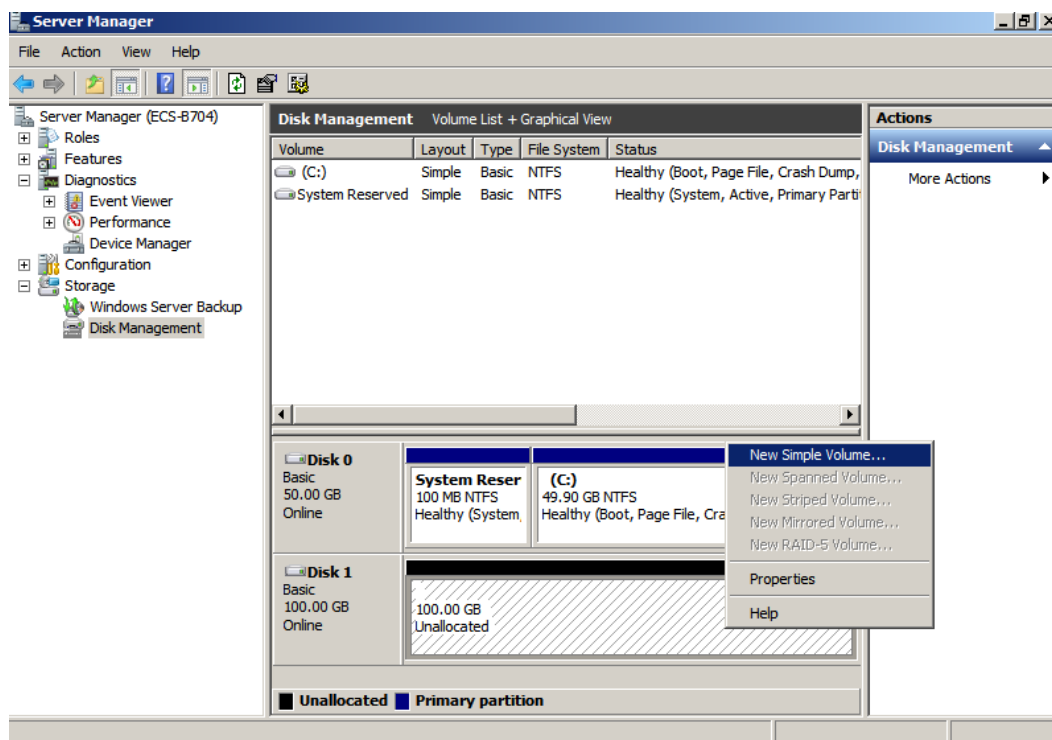
- Step 5** Right-click in the disk area and choose **Initialize Disk** from the shortcut menu. In the displayed **Initialize Disk** dialog box, select **MBR (Master Boot Record)** and click **OK**.

 **NOTE**

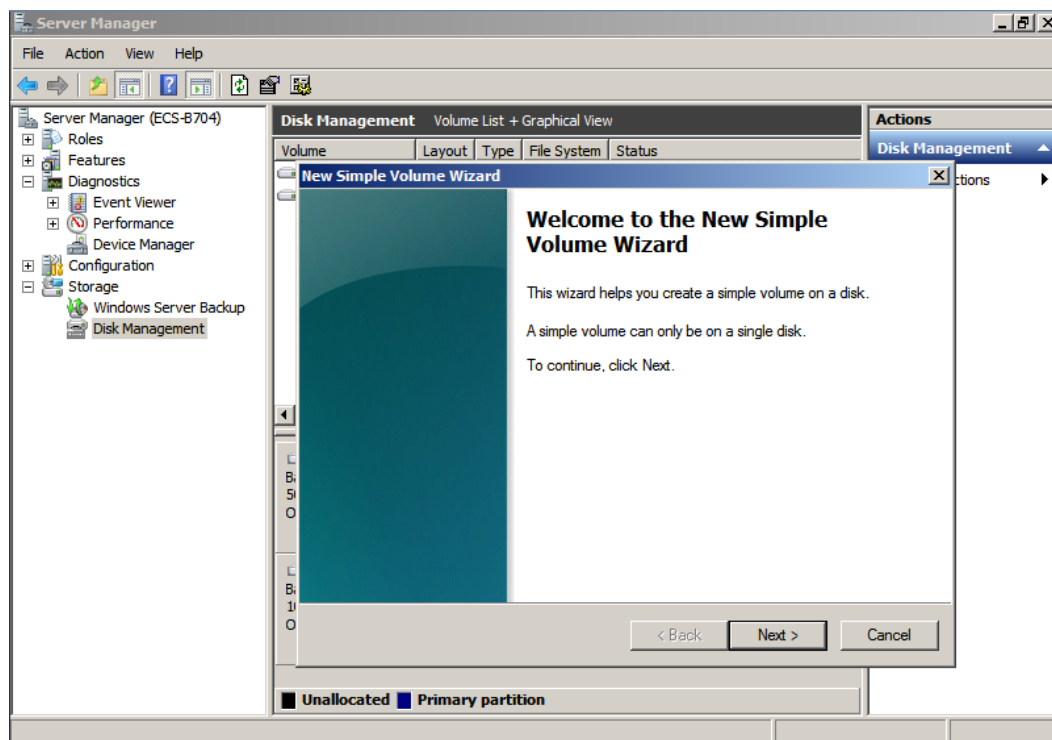
If the data disk to be initialized is larger than 2 TB, select **GPT (GUID Partition Table)** in the dialog box.



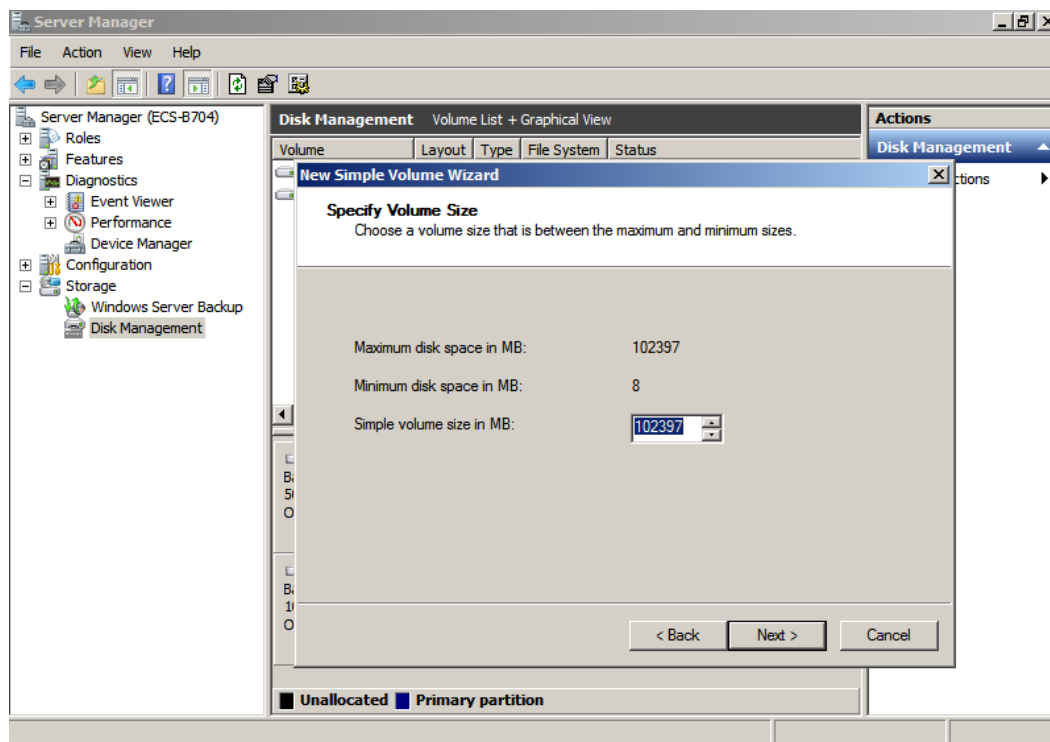
- Step 6** Right-click at the unallocated disk space and choose **New Simple Volume** from the shortcut menu.



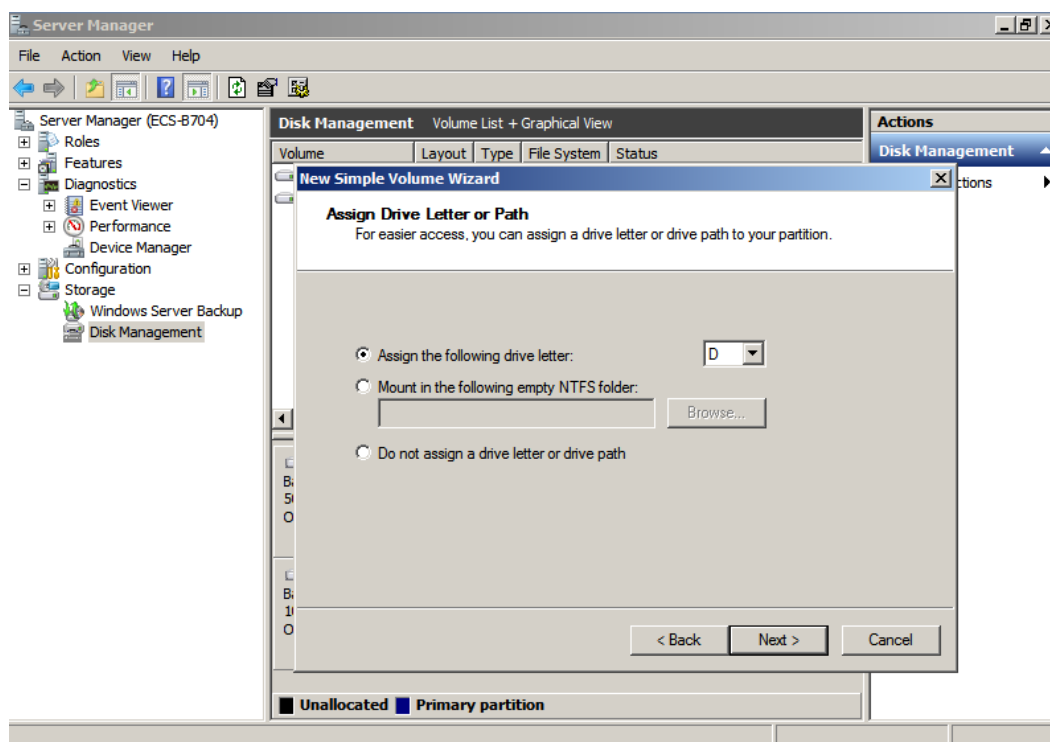
Step 7 On the displayed **New Simple Volume Wizard** page, click **Next**.



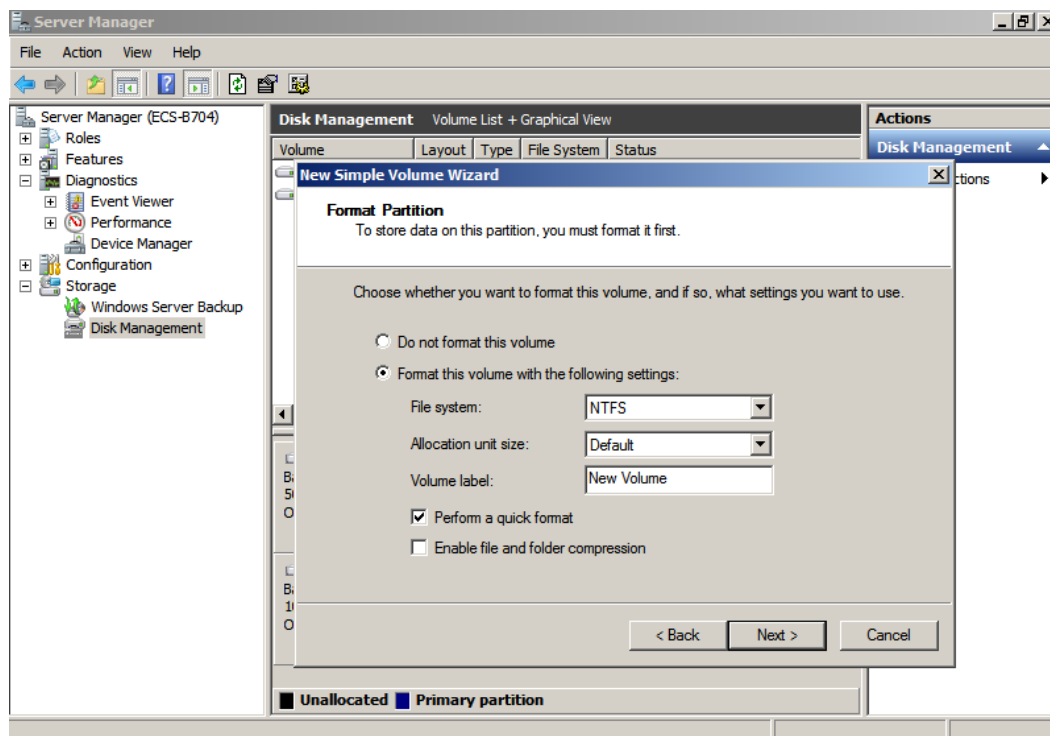
Step 8 Specify the simple volume size as required (the default value is the maximum) and click **Next**.



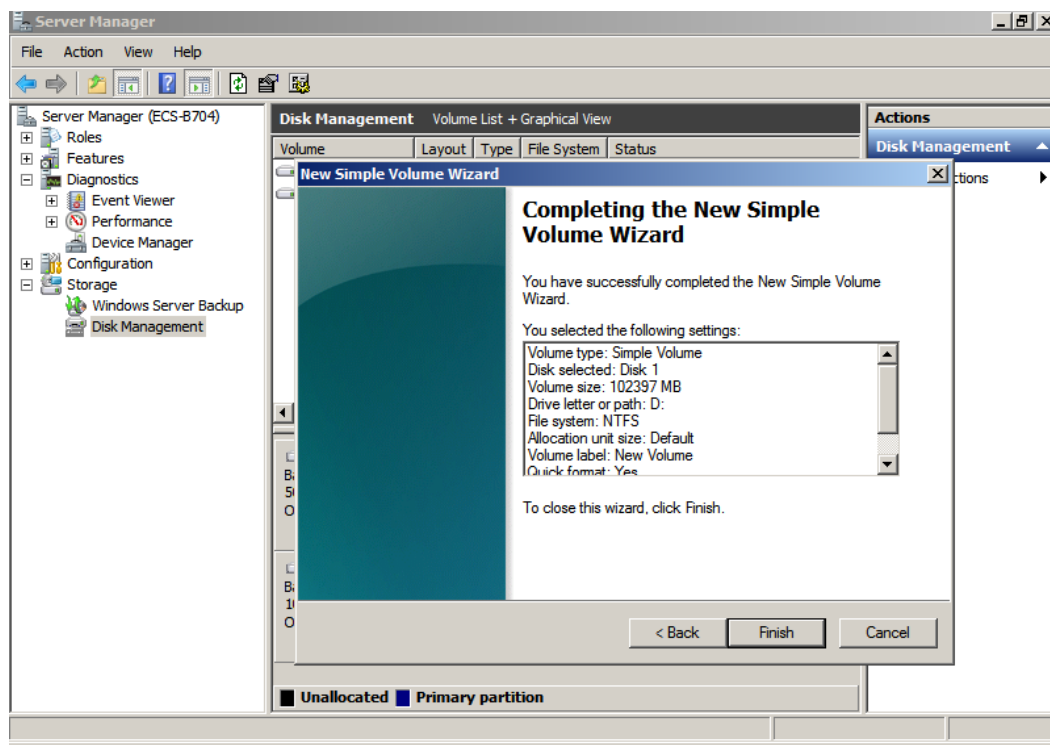
Step 9 Assign the driver letter and click **Next**.



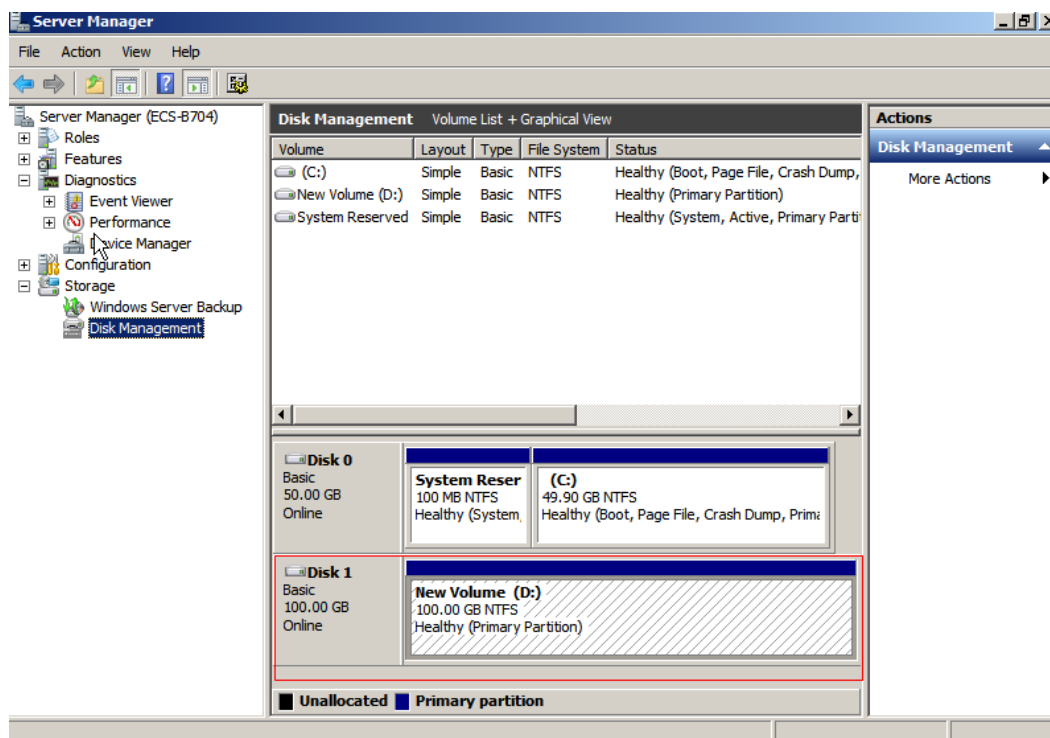
Step 10 Select **Format this volume with the following settings**, set parameters based on the actual requirements, and select **Perform a quick format**. Then click **Next**.



Step 11 Click **Finish**.



Wait for the initialization to complete. When the volume status changes to **Healthy**, the initialization has finished.



----End

1.5.2 Initializing a Linux Data Disk

Creating Partitions and Mounting a Disk

The following example shows how to create a new primary partition on a new data disk that has been attached to an instance. The primary partition will be created using `fdisk`, and MBR is the default partition style. Furthermore, the partition will be formatted using the `ext4` file system, mounted on the `/mnt/sdc` directory, and set to be automatically mounted upon a system start.

Step 1 Log in to an ECS.

For details, see [Operation Help Center > Compute > Elastic Cloud Server > User Guide > Logging In to an ECS](#).

Step 2 Run the following command to view information about the added data disk:

`fdisk -l`

Information similar to the following is displayed: (In the command output, the server contains two disks. `/dev/xvda` is the system disk, and `/dev/xvdb` is the added data disk.)

 **NOTE**

If you do not log in to the ECS and run the **umount** command but directly detach the **/dev/xvdb** or **/dev/vdb** EVS disk on the management console, the disk name in the ECS may encounter a release delay. When you attach the disk to the server again, the mount point displayed on the management console may be inconsistent with that in the server. For example, device name **/dev/sdb** or **/dev/vdb** is selected for attachment, but **/dev/xvdc** or **/dev/vdc** may be displayed as the disk name in the OS. This issue does not adversely affect services.

```
[root@ecs-b656 test]# fdisk -l
```

```
Disk /dev/xvda: 42.9 GB, 42949672960 bytes, 83886080 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x000cc4ad
```

Device	Boot	Start	End	Blocks	Id	System
/dev/xvda1	*	2048	2050047	1024000	83	Linux
/dev/xvda2		2050048	22530047	10240000	83	Linux
/dev/xvda3		22530048	24578047	1024000	83	Linux
/dev/xvda4		24578048	83886079	29654016	5	Extended
/dev/xvda5		24580096	26628095	1024000	82	Linux swap / Solaris

```
Disk /dev/xvdb: 10.7 GB, 10737418240 bytes, 20971520 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
```

 **NOTE**

The displayed capacity is inconsistent with the EVS disk capacity applied for on ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios). The reason is as follows: The capacity of EVS disks is calculated using the unit of GiB (Gibibyte), while the capacity unit in Linux OS is GB (Gigabyte). The GiB is calculated in binary mode, and the GB is calculated in decimal format. 1 GiB = 1,073,741,824 Bytes, 1 GB = 1,000,000,000 Bytes.

Step 3 Run the following command to allocate partitions for the added data disk using **fdisk**:

fdisk *Newly added data disk*

In this example, **/dev/xvdb** is the newly added data disk.

fdisk /dev/xvdb

Information similar to the following is displayed:

```
[root@ecs-b656 test]# fdisk /dev/xvdb
Welcome to fdisk (util-linux 2.23.2).
Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.
Device does not contain a recognized partition table
Building a new DOS disklabel with disk identifier 0xb00005bd.
Command (m for help):
```

Step 4 Enter **n** and press **Enter**.

Entering **n** creates a partition.

There are two types of disk partitions:

- Choosing **p** creates a primary partition.
- Choosing **e** creates an extended partition.

```
Command (m for help): n
Partition type:
  p primary (0 primary, 0 extended, 4 free)
  e extended
```

Step 5 Enter **p** and press **Enter**.

The following describes how to create a primary partition.

Information similar to the following is displayed. **Partition number** indicates the serial number of the primary partition. The value can be **1** to **4**.

```
Select (default p): p
Partition number (1-4, default 1):
```

Step 6 Enter the primary partition number **1** and press **Enter**.

For example, select **1** as the partition number.

Information similar to the following is displayed. **First sector** indicates the start sector number. The value can be **2048** to **20971519**, and the default value is **2048**.

```
Partition number (1-4, default 1): 1
First sector (2048-20971519, default 2048):
```

Step 7 Press **Enter**.

The default start sector number 2048 is used as an example.

Information similar to the following is displayed: (**Last sector** indicates the last sector number. The value can be from **2048** to **20971519**, and the default value is **20971519**.)

```
First sector (2048-20971519, default 2048):
Using default value 2048
Last sector, +sectors or +size{K,M,G} (2048-20971519, default 20971519):
```

Step 8 Press **Enter**.

The default last sector number 20971519 is used as an example.

Information similar to the following is displayed. A primary partition is created for a 10 GB data disk.

```
Last sector, +sectors or +size{K,M,G} (2048-20971519, default 20971519):
Using default value 20971519
Partition 1 of type Linux and of size 10 GiB is set
Command (m for help):
```

Step 9 Enter **p** and press **Enter** to view the details about the created partition.

Information similar to the following is displayed. Details about the **/dev/xvdb1** partition are displayed.

```
Command (m for help): p

Disk /dev/xvdb: 10.7 GB, 10737418240 bytes, 20971520 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0xb00005bd

   Device Boot      Start         End      Blocks   Id  System
/dev/xvdb1        2048     20971519     10484736    83   Linux

Command (m for help):
```

Step 10 Enter **w** and press **Enter** to write the changes into the partition table.

Information similar to the following is displayed. The partition is successfully created.

```
Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.
```

Step 11 Run the following command to synchronize the new partition table to the data disk:

partprobe

Step 12 Run the following command to set the format for the file system of the newly created partition:

mkfs -t *File system format* /dev/xvdb1

In this example, run the following command to set the **ext4** file system for the new partition:

mkfs -t ext4 /dev/xvdb1

Information similar to the following is displayed:

```
[root@ecs-b656 test]# mkfs -t ext4 /dev/xvdb1
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
655360 inodes, 2621184 blocks
131059 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2151677952
80 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632

Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

NOTE

The formatting takes a period of time. Observe the system running status and do not exit.

Step 13 Run the following command to create a mount directory:

mkdir *Mount directory*

/mnt/sdc is used in this example.

mkdir /mnt/sdc

Step 14 Run the following command to mount the new partition to the mount directory created in [Step 13](#):

mount /dev/xvdb1 *Mount directory*

For example, run the following command to mount the newly created partition on **/mnt/sdc**:

```
mount /dev/xvdb1 /mnt/sdc
```

Step 15 Run the following command to view the mount result:

```
df -TH
```

Information similar to the following is displayed. The newly created **/dev/xvdb1** partition has been mounted on **/mnt/sdc**.

```
[root@ecs-b656 test]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/xvda2      xfs       11G   7.4G  3.2G   71% /
devtmpfs        devtmpfs  4.1G    0  4.1G    0% /dev
tmpfs           tmpfs     4.1G   82k  4.1G    1% /dev/shm
tmpfs           tmpfs     4.1G   9.2M  4.1G    1% /run
tmpfs           tmpfs     4.1G    0  4.1G    0% /sys/fs/cgroup
/dev/xvda3      xfs       1.1G   39M   1.1G    4% /home
/dev/xvda1      xfs       1.1G  131M  915M   13% /boot
/dev/xvdb1      ext4      11G   38M   9.9G    1% /mnt/sdc
```

----End

Setting Automatic Disk Attachment at a System Start

If you require a disk to be automatically attached to an instance when the instance is started, enable automatic disk attachment upon an instance start by referring to operations provided in this section. When enabling automatic disk attachment, you cannot directly specify **/dev/xvdb1** in **/etc/fstab**. This is because the sequence codes of the instance may change during an instance stop or start process. You are advised to use the universally unique identifier (UUID) in **/etc/fstab** to automatically attach the disk at a system start.

NOTE

The UUID of a disk is a character string that uniquely identifies a storage device in a Linux system.

Step 1 Run the following command to query the partition UUID:

```
blkid Disk partition
```

For example, run the following command to query the UUID of **/dev/xvdb1**:

```
blkid /dev/xvdb1
```

Information similar to the following is displayed: (The UUID of **/dev/xvdb1** is displayed.)

```
[root@ecs-b656 test]# blkid /dev/xvdb1
/dev/xvdb1: UUID="1851e23f-1c57-40ab-86bb-5fc5fc606ffa" TYPE="ext4"
```

Step 2 Run the following command to open the **fstab** file using the vi editor:

```
vi /etc/fstab
```

Step 3 Press **i** to enter the editing mode.

Step 4 Move the cursor to the end of the file and press **Enter**. Then add the following information:

UUID=xxx attachment directory file system defaults 0 2

Assuming that the file system is **ext4** and the attachment directory is **/mnt/sdc**.

```
UUID=1851e23f-1c57-40ab-86bb-5fc5fc606ffa /mnt/sdc ext4 defaults 0 2
```

NOTICE

After automatic attachment upon instance start is configured, comment out or delete the line in the **fstab** file before detaching the disk. Otherwise, you may fail to access the operating system after the disk is detached.

Step 5 Press **Esc**, enter **:wq**, and press **Enter**.

The system saves the configuration and exits the editor.

----End

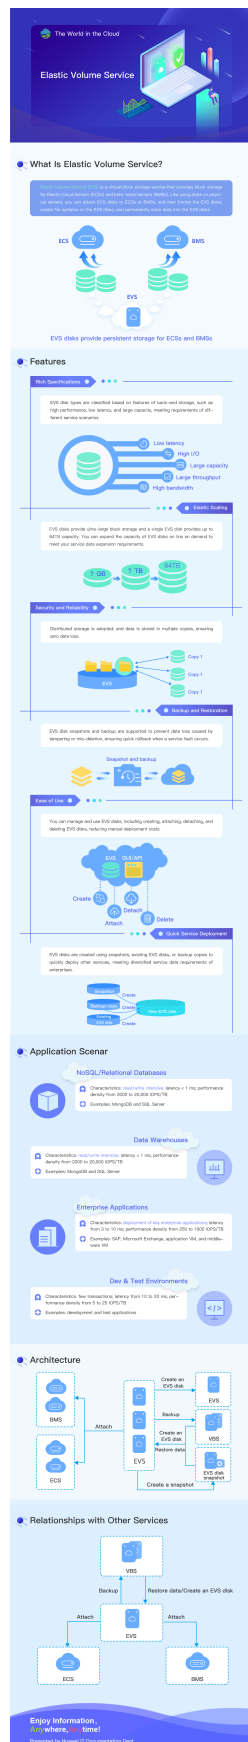
2 Infographics

[2.1 Introduction](#)

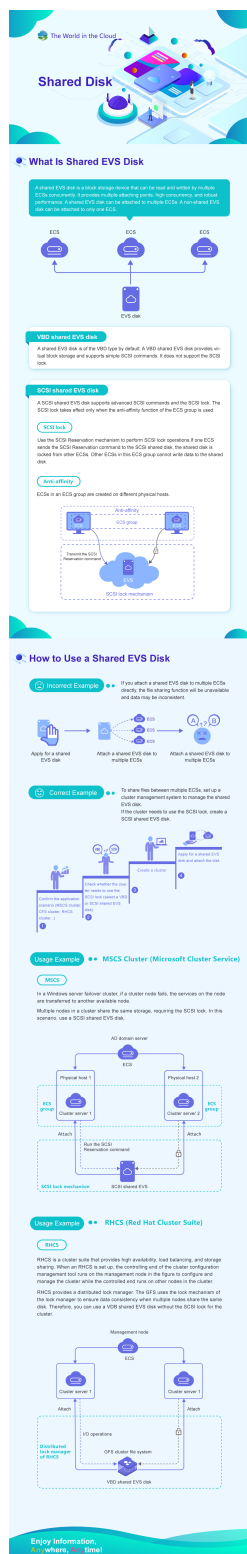
[2.2 Shared Disk](#)

[2.3 Changing the Disk Type](#)

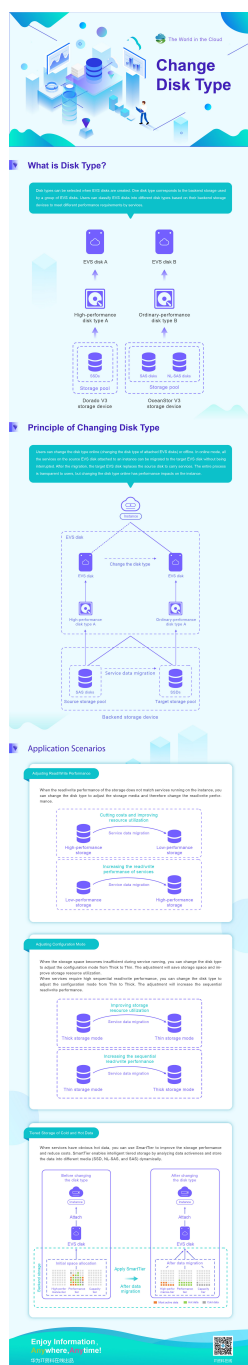
2.1 Introduction



2.2 Shared Disk



2.3 Changing the Disk Type



3 User Guide (for ECS)

- [3.1 Introduction](#)
- [3.2 Related Concepts](#)
- [3.3 Operation Process](#)
- [3.4 Applying for an EVS Disk](#)
- [3.5 Attaching an EVS Disk](#)
- [3.6 Initializing a Data Disk](#)
- [3.7 Expanding EVS Disk Capacity](#)
- [3.8 Releasing an EVS Disk](#)
- [3.9 Managing Disk Types](#)
- [3.10 Managing Snapshots](#)
- [3.11 Managing Backups](#)
- [3.12 Managing EVS Disks](#)
- [3.13 FAQs](#)

3.1 Introduction

3.1.1 What Is Elastic Volume Service?

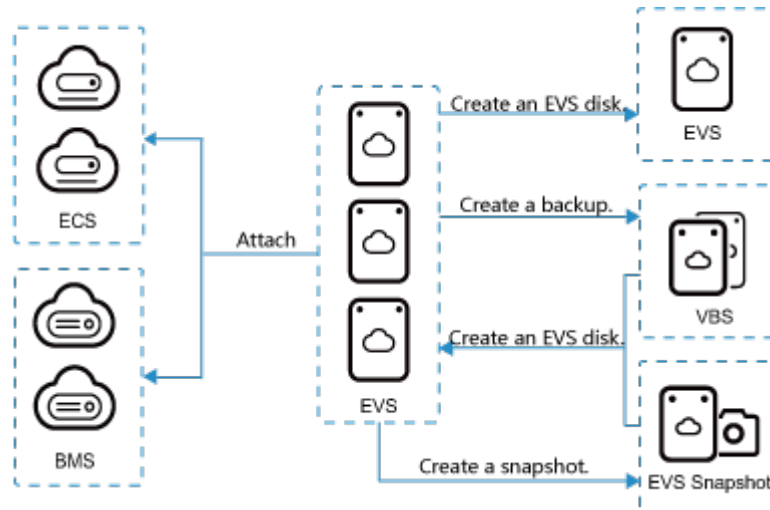
Definition

Elastic Volume Service (EVS) is a virtual block storage service, which provides block storage space for Elastic Cloud Servers (ECSs) and Bare Metal Servers (BMSs). You can create EVS disks on the console and attach them to ECSs and BMSs. The method for using EVS disks is the same as that for using disks on physical servers. EVS disks have higher data reliability and I/O throughput and are easier to use. EVS disks are suitable for file systems, databases, or system software or applications that require block storage devices. [Figure 3-1](#) shows how to use an EVS disk.

In this document, an EVS disk is also referred to as a disk.

In this document, instances refer to the ECSs or BMSs that you apply for.

Figure 3-1 EVS functions



Functions

EVS provides various persistent storage devices. You can choose disk types based on your needs and store files and build databases on EVS disks. EVS has the following major features:

- Elastic attaching and detaching
An EVS disk is like a raw, unformatted, external block device that you can attach to a single instance. Disks are not affected by the running time of instances. After attaching a disk to an instance, you can use the disk as if you were using a physical disk. You can also detach a disk from an instance and attach the disk to another instance.
- Various disk types
A disk type represents backend storage devices used by a group of disks. You can divide disk types of EVS disks based on backend storage types to meet different performance requirements of services. If the read/write performance of an upper-layer service does not match that of the storage medium used by the service, you can change the disk type to change the read/write performance of the storage medium to meet the requirements of instance storage service performance adjustment.
- Scalability
A single disk has a maximum capacity of 64 TB. You can configure storage capacity and expand the capacity on demand to deal with your service data increase.
- Snapshot
You can back up your data by taking a snapshot of your disk data at a specific point in time to prevent data loss caused by data tampering or mis-deletion and ensure a quick rollback in the event of a service fault. You can also create disks from snapshots and attach them to other instances to provide data resources for a variety of services, such as data mining, report query, and

development and test. This method protects the initial data and creates disks rapidly, meeting the diversified service data requirements.

- Shared disk

Multiple instances can access (read and write) a shared disk at the same time, meeting the requirements of key enterprises that require cluster deployment and high availability (HA).

Comparison Between EVS and SFS

Table 3-1 compares EVS and SFS.

Table 3-1 Comparison between EVS and SFS

Dimension	EVS	SFS
Usage	Provides persistent block storage for compute services such as ECS and BMS. EVS disks feature high availability, high reliability, and low latency. You can format, create file systems on, and persistently store data on EVS disks.	Provides ECSs with a high-performance shared file system that supports on-demand elastic scaling. The file system complies with the standard file protocol and delivers scalable performance, supporting massive amount of data and bandwidth-demanding applications.
Data access mode	Data access is limited within the internal network of a data center.	Data access is limited within the internal network of a data center.
Sharing mode	Supports EVS disk sharing. A shared EVS disk can be attached to a maximum of 16 ECSs in the cluster management system.	Supports data sharing. A file system can be mounted to a maximum of 256 ECSs.
Storage capacity	The maximum capacity of a single disk is 64 TB.	The maximum capacity of a single file is 240 TB, and the file system capacity can be scaled to the PB level.
Backend storage	Supports Huawei SAN storage, Huawei Distributed Block Storage, and heterogeneous storage.	OceanStor 9000 and OceanStor Dorado 6.x
Recommended scenarios	Scenarios such as database, enterprise office applications, and development and testing.	Scenarios such as media processing and file sharing.

3.1.2 Advantages

- Varying specifications
EVS disks of different performance levels are provided. You can choose and configure EVS disks of appropriate performance levels to meet your service requirements.
- Scalable
EVS disks provide ultra-large block storage and a single EVS disk has a maximum capacity of 64 TB. You can expand the EVS disk capacity on running ECSs to meet your increasing service requirements.
 - On-demand expansion
You can expand the capacity of EVS disks based on your needs, with at least 1 GB added at a time.
 - Linear performance improvement
You can expand the capacity of EVS disks on running ECSs to implement linear performance improvement, thereby meeting your service requirements.
- Secure and reliable
Distributed storage is adopted, and data is stored in multiple identical copies, ensuring zero data loss. Data durability reaches 99.9999999%.
- Backup and restoration
Functions, such as EVS disk backup and EVS disk snapshot, are supported to prevent incorrect data caused by application exceptions or attacks.
 - EVS disk backup
This function enables the system to create EVS disk backups. The backups can be used to roll back EVS disks, maximizing user data accuracy and security and ensuring service availability.
 - EVS disk snapshot
This function enables the system to create snapshots for EVS disks. A snapshot can be used to roll back an EVS disk to the state when the snapshot is created, maximizing data accuracy and security and ensuring service availability.

3.1.3 Application Scenarios

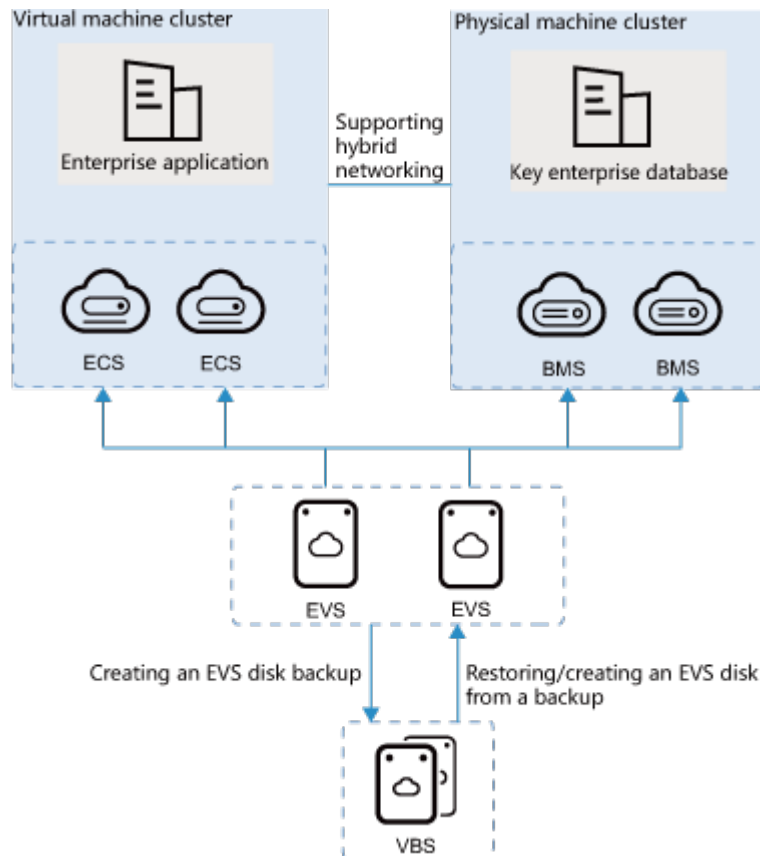
You can configure and select disk types with different service levels based on your application requirements for flexible deployment.

Relational Database

The service core database needs to support massive access at traffic peaks, and requires disks with persistent and stable high performance and low latency. You can use the disk type with ultra-high performance to implement a combination of excellent performance and superior reliability, meeting the high requirements for low latency and high I/O performance in data-intensive scenarios, such as relational databases. [Figure 3-2](#) shows the architecture in these scenarios. Disks with ultra-high performance service levels can meet the following performance requirements:

- The latency is shorter than 1 ms.
- The performance ranges from 2000 IOPS/TB to 20,000 IOPS/TB.
- Typical configurations: Enterprise storage Dorado 5000 V3 is selected for backend storage, twenty-five 1 TB, 2 TB, or 4 TB SSDs are configured for every dual controllers, and RAID 6 is configured. Deduplication and compression functions are enabled, and a maximum of four controllers and 50 disks (30 TB, 60 TB, or 120 TB) are configured for a single system.

Figure 3-2 Architecture in the relational database scenario



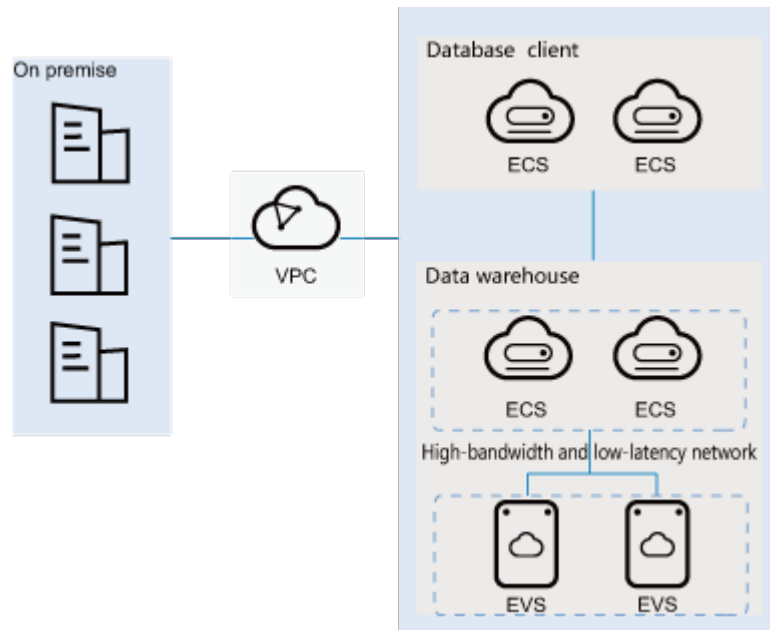
Data Warehouse

In scenarios with intensive data reads, deploy data warehouses, and it is recommended that you use the disk type with high performance to meet the application requirements for low latency, high read and write speed, and large throughput. **Figure 3-3** shows the architecture in these scenarios. Disks with high performance service levels can meet the following performance requirements:

- The delay ranges from 1 ms to 3 ms.
- The performance ranges from 500 IOPS/TB to 4000 IOPS/TB.
- Typical configuration 1: OceanStor 6800 V5 is selected for backend storage, fifty 1.92 TB, 3.84 TB, or 7.68 TB SSDs are configured for every two controllers, and RAID 5 is configured. A maximum of eight controllers and 200 disks (300 TB, 600 TB, or 1200 TB) are configured for a single system.
- Typical configuration 2: Huawei Distributed Block Storage is selected for backend storage. RH2288H V5 servers are used. Twelve 4 TB, 6 TB, 8 TB, or 10

TB SATA disks are configured. Three-duplicate mode is adopted. One 1.6 TB or 3.2 TB SSD is configured. The total available space on each node is about 15.2 TB, 22.8 TB, 30.4 TB, or 38 TB.

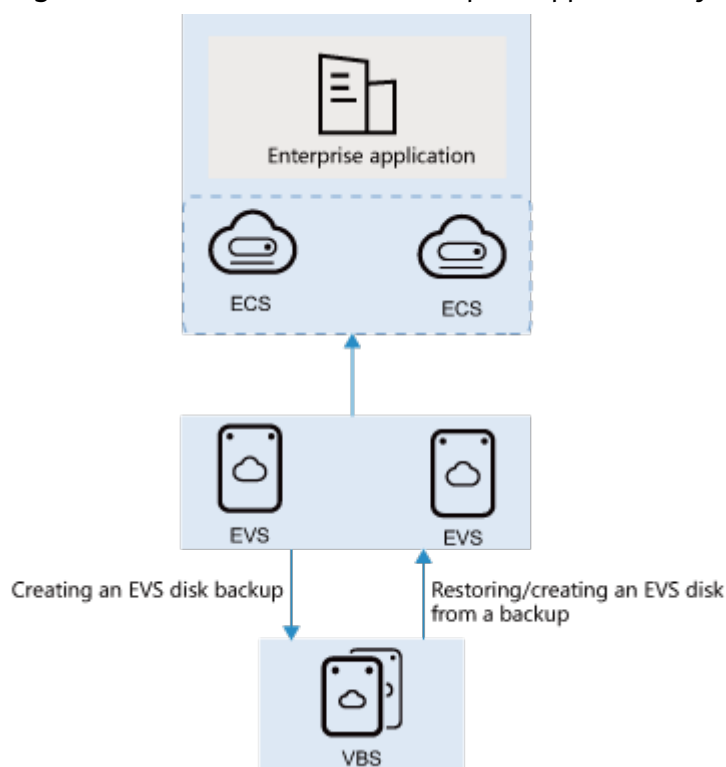
Figure 3-3 Data warehouse scenario architecture



Enterprise Application System

Mission-critical applications of enterprises are deployed in these scenarios. These scenarios, such as common databases, application VMs, and middleware VMs, require relatively low performance but rich enterprise-class features. It is recommended that you use the disk type with medium performance. [Figure 3-4](#) shows the architecture in these scenarios. Disks with medium performance service levels can meet the following performance requirements:

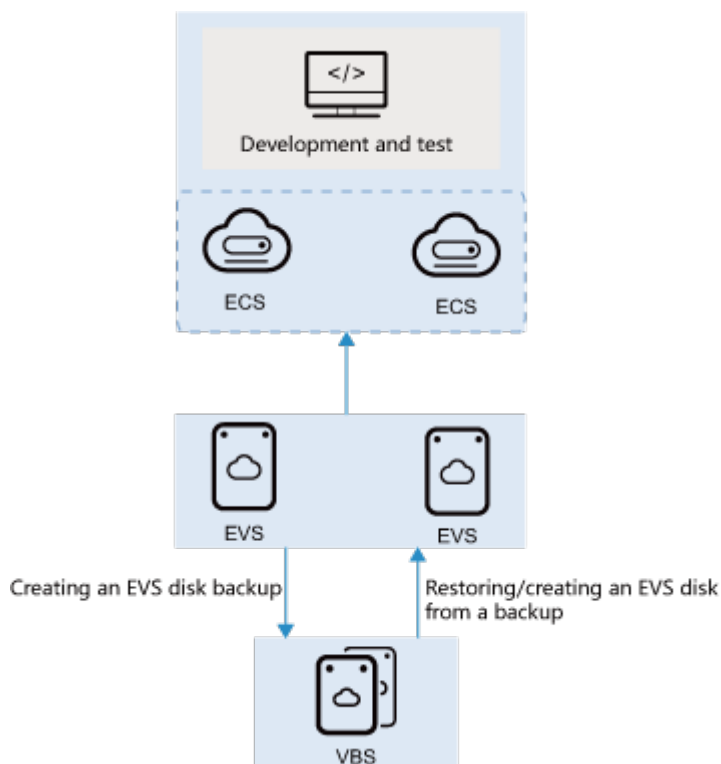
- The delay ranges from 3 ms to 10 ms.
- The performance ranges from 250 IOPS/TB to 1000 IOPS/TB.
- Typical configuration 1: OceanStor 5500 V5 is selected for backend storage. There are fewer than 250 disks for every two controllers, including ten 1.92 TB, 3.84 TB, or 7.68 TB SSDs and fewer than 240 600 GB, 1.2 TB, or 1.8 TB SAS disks. RAID 5 is configured. A single system supports a maximum of six controllers and 750 disks (360 TB, 720 TB, or 1116 TB).
- Typical configuration 2: Huawei Distributed Block Storage is selected for backend storage. 5288 V3 servers are used. Thirty-six 2 TB, 4 TB, 6 TB, or 8 TB SATA disks are configured. Three-duplicate mode is adopted. Two 1.6 TB or 3.2 TB SSDs are configured. The total available space on each node is about 22.8 TB, 45.6 TB, 68.4 TB, or 91.2 TB.

Figure 3-4 Architecture in the enterprise application system scenario

Development and Test

In these scenarios, development and test applications are deployed. It is recommended that you use the disk type with common performance to meet the requirements of development, test, deployment, and O&M. [Figure 3-5](#) shows the architecture in these scenarios. Disks with common performance service levels can meet the following performance requirements:

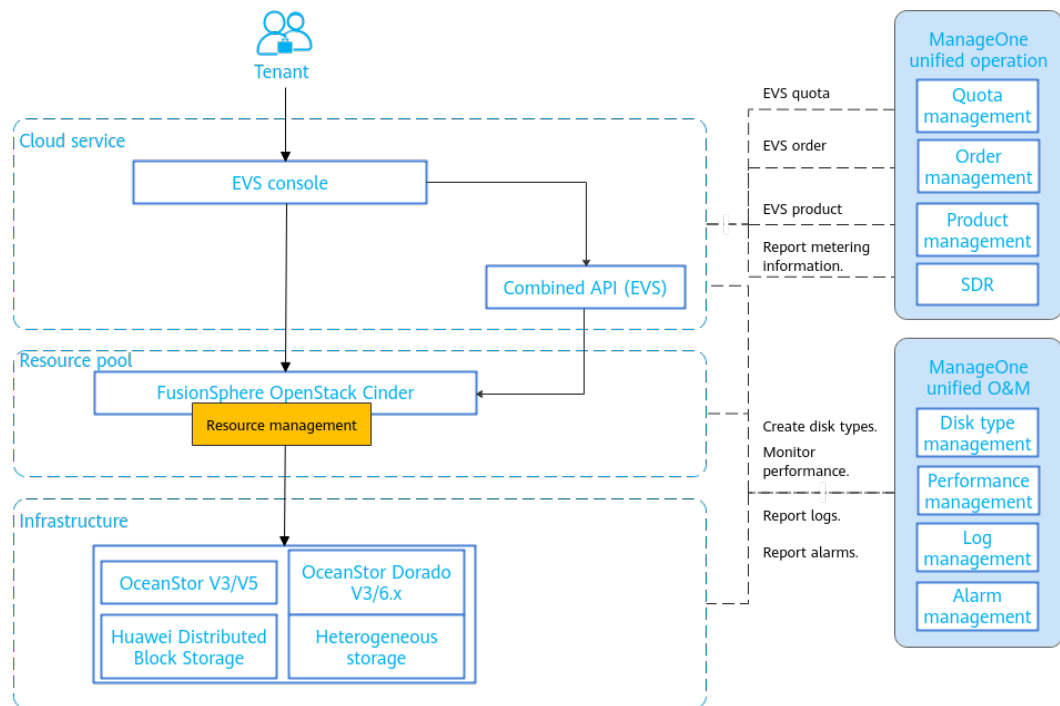
- The delay ranges from 10 ms to 20 ms.
- The performance ranges from 5 IOPS/TB to 25 IOPS/TB.
- Typical configuration: OceanStor 5300 V5 is selected for backend storage. Fewer than 396 disks (2 TB/4 TB/6 TB/8 TB/10 TB NL-SAS disks) are configured for every two controllers. RAID 6 is configured. A single system supports a maximum of two controllers (612 TB/1224 TB/1840 TB/2460 TB/3060 TB).

Figure 3-5 Development and test scenario architecture

3.1.4 Implementation Principles

Architecture

EVS includes components such as the EVS console, EVS service API, FusionSphere OpenStack Cinder, and storage device. [Figure 3-6](#) shows the logical architecture of EVS.

Figure 3-6 Logical architecture of EVS**Table 3-2** EVS component description

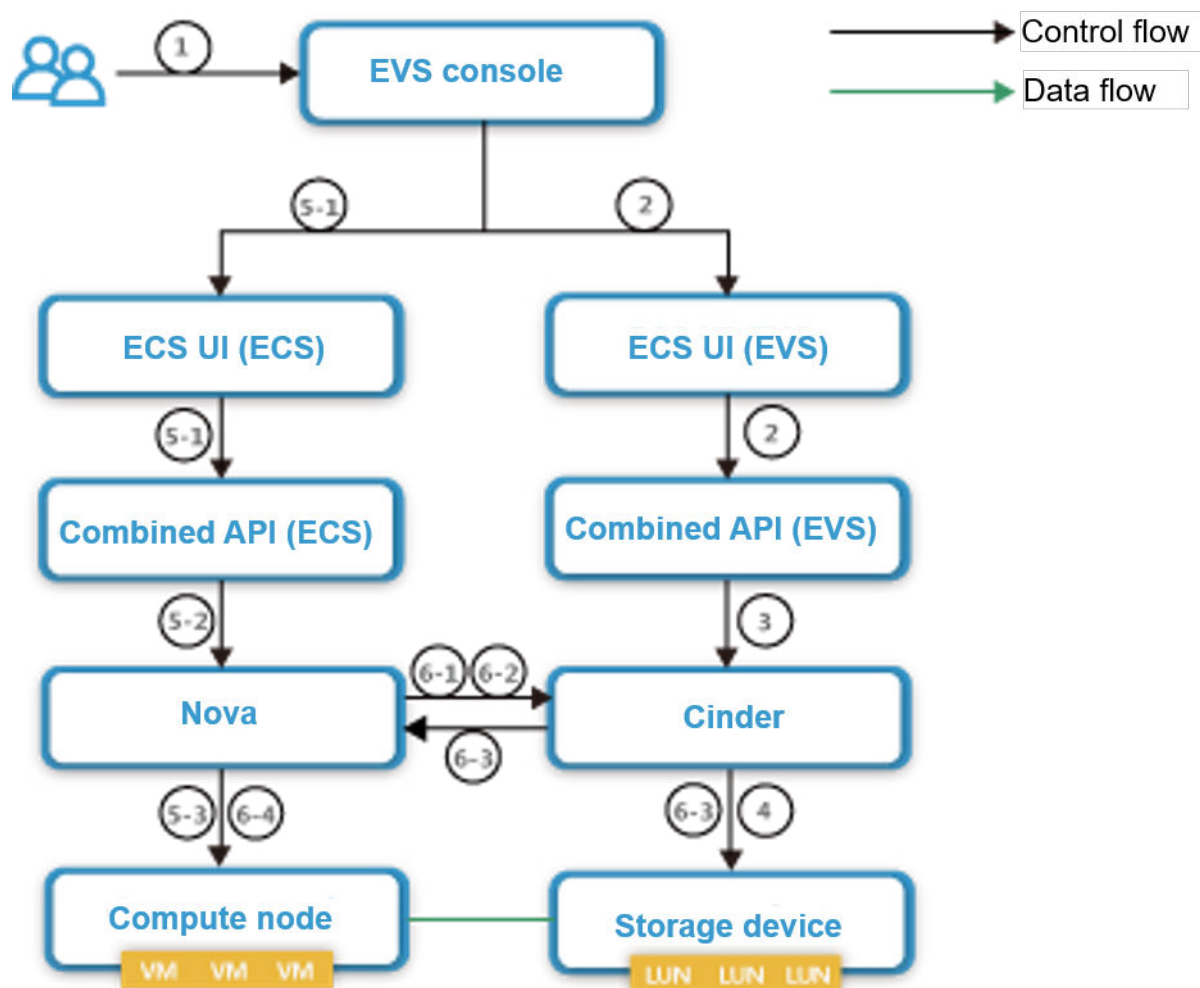
Component Name	Description
EVS console	The EVS console provides tenants with an entry to EVS. Tenants can apply for EVS disks on the console.
Combined API (EVS)	The EVS service API encapsulates or combines the logic based on the native Cinder interface to implement some EVS functions. The EVS service API can be invoked by the EVS console or tenants.
FusionSphere OpenStack Cinder	FusionSphere OpenStack Cinder provides persistent block storage to manage block storage resources. It is mainly used to create disk types in EVS. Disks are created on the storage device and attached to ECSs or BMSs.
Infrastructure	Infrastructure refers to the physical storage device that provides block storage based on physical resources. The following storage devices can function as the backend storage of EVS: Huawei SAN storage (OceanStor V3/V5 and OceanStor Dorado V3/6.x), Huawei Distributed Block Storage, and heterogeneous storage (such as HP 3PAR 8000 series).
ManageOne unified operation	ManageOne unified operation provides quota management, order management, product management, and resource metering and charging for EVS.

Component Name	Description
ManageOne unified O&M	ManageOne unified O&M provides disk type management, performance monitoring, logging, and alarm reporting for EVS.

Workflow

Figure 3-7 shows the workflow for EVS to provision EVS disks and attach the disks to ECSs.

Figure 3-7 EVS workflow



1. A VDC administrator or VDC operator applies for storage resources on the EVS console.
2. The EVS console sends the request to Combined API (EVS) through ECS UI (EVS).
3. Combined API distributes the request to Cinder.

4. Cinder creates volumes in the storage pool based on the policy for applying for storage resources. Cinder includes the following components:
 - Cinder API: receives external requests.
 - Cinder Scheduler: selects a proper backend storage server and specifies the storage server where the created volume resides.
 - Cinder Volume: connects to various storage device drivers and delivers requests to specific storage devices.
5. The VDC administrator or VDC operator attaches the requested storage resources to ECSs on the EVS console.
 - a. The EVS console sends the request to Combined API (ECS) through ECS UI (ECS).
 - b. Combined API distributes the request to Nova.
 - c. Nova processes the attachment task using Nova-compute running on the compute node.
6. Nova instructs Cinder to attach EVS disks.
 - a. Nova obtains EVS disk information and instructs Cinder to reserve EVS disks.
 - b. Nova obtains host initiator information and sends it to Cinder.
 - c. Cinder instructs the storage array to map the initiator and target and returns the Nova target information.
 - d. Nova completes the attachment task.

3.1.5 Related Services

Figure 3-8 shows the dependencies between EVS and other cloud services. **Table 3-3** provides more details.

Figure 3-8 Relationship between the EVS service and other cloud services

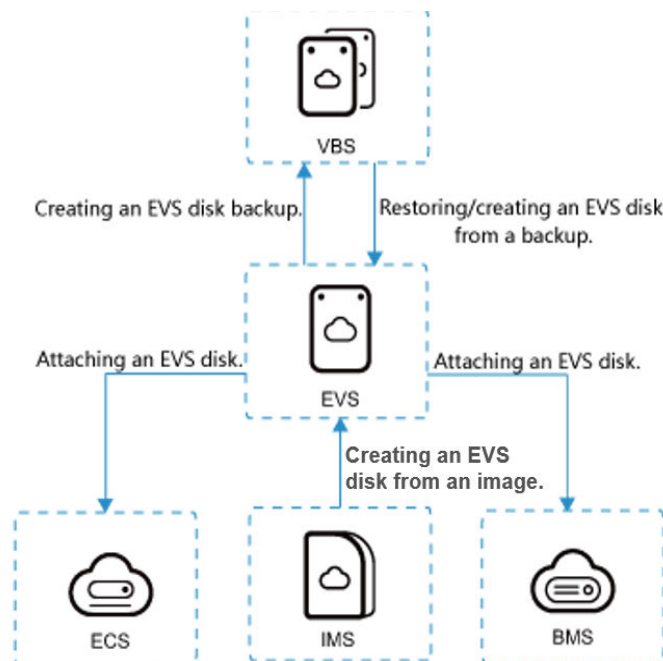


Table 3-3 Dependencies between EVS and other cloud services

Service Name	Description
ECS	You can attach EVS disks to ECSs to provide scalable block storage.
BMS	You can attach SCSI EVS disks to BMSs to provide scalable block storage.
VBS	Volume Backup Service (VBS) can be used to create backups for EVS disks. EVS disk data can be restored using the backups. Backups can be used to create EVS disks.
IMS	EVS can be used to create data disks from data disk images and system disks from system disk images. Image Management Service (IMS) can be used to create data disk images or system disk images.

3.1.6 Key Metrics

Table 3-4 lists the key metrics of EVS.

Table 3-4 Key metrics of the EVS service

Item	Metric
Maximum number of EVS disks that you can obtain at a time	This metric is related to the EVS disk quota. <ul style="list-style-type: none">If the number of EVS disks in the quota is greater than 100, a maximum number of 100 EVS disks can be applied for each time.If the number of EVS disks in a quota is less than 100, the maximum number of EVS disks that can be applied for each time is equal to the quota quantity.
Maximum number of instances to which a shared disk can be attached simultaneously	16 If Huawei SAN storage is used as backend storage and the storage version is earlier than V300R006C50, a shared disk can be attached to fewer than eight instances simultaneously.
Maximum number of snapshots that can be created for an EVS disk	32 (recommended) This item is related to the backend storage type. The maximum number of snapshots that can be created varies with the backend storage type. For details, see the product documentation of the corresponding backend storage.

3.1.7 Restrictions

Before using EVS, learn the restrictions described in [Table 3-5](#).

Virtual Resource Management (VRM) is a piece of virtualization management software developed by Huawei. It works with the Unified Virtualization Platform (UVP) to form a virtual infrastructure product. In VRM scenarios, if a datastore of the VIMS type is used, the datastore supports only some functions of the EVS, as listed in [Table 3-6](#).

Table 3-5 Restrictions on EVS

Item	Restrictions
Backend storage	<ul style="list-style-type: none">Supported Huawei storage device types include enterprise storage OceanStor V3/V5, OceanStor Dorado V3/6.x, and Huawei Distributed Block Storage. You can visit HUAWEI CLOUD Stack Information Center to query the specific storage models and versions.For details about the supported heterogeneous storage types and versions, visit HUAWEI CLOUD Stack Information Center.When OceanStor Dorado V3/6.x or heterogeneous storage is used as the backend storage, AZs with the virtualization capability of Xen cannot be accessed but AZs with the virtualization capability of KVM or IroniC can be accessed.An AZ with the virtualization capability of KVM can contain multiple Huawei SAN storage devices, such as OceanStor V3/V5 and OceanStor Dorado V3/6.x. It can also contain both Huawei Distributed Block Storage and Huawei SAN storage.An AZ with the virtualization capability of VRM or IroniC supports only multiple Huawei SAN storage devices, such as OceanStor V3/V5 and OceanStor Dorado V3/6.x. Huawei Distributed Block Storage and Huawei SAN storage cannot share the same AZ.Heterogeneous storage cannot be deployed in the same AZ with Huawei SAN storage and Huawei Distributed Block Storage.Only one set of Huawei Distributed Block Storage can be deployed in an AZ.FC SAN and IP SAN protocols cannot be used in the same AZ at the same time.It is recommended that a disk type contain only backend storage of the same storage type, ensuring that the backend storage capabilities are the same.If the backend storage is FusionStorage 8.0.0 or later, deduplication and compression are enabled by default. Therefore, the provisioned disks have the deduplication and compression functions.If the backend storage is FusionStorage 8.0.0 or later and the type of the access storage pool is self-encrypting, the provisioned disks support data encryption.If the backend storage is OceanStor V500R007C60SPC300 or later, or OceanStor Dorado 6.0.0 or later, set Max. Sessions per User to 0 on DeviceManager, indicating that the number of sessions is not limited.If the backend storage is OceanStor Dorado 6.x, manually disable the recycle bin function. For details, see "Configure" > "Basic Storage Service Configuration Guide for Block" > "Managing Basic Storage Services" > "Managing LUNs" > "Managing the Recycle Bin" > "Configuring the Recycle Bin" in

Item	Restrictions
	<i>OceanStor Dorado 8000 and Dorado 18000 Product Documentation.</i>

Item	Restrictions
Applying for an EVS disk	<ul style="list-style-type: none">• The maximum capacity of a single disk is 64 TB.• Shared disks can be used as data disks and cannot be used as system disks.• When the backend storage type is heterogeneous: Blank EVS disks can be created and EVS disks can be created from an existing disk or a snapshot. EVS disks cannot be created from a backup and the disk type cannot be SCSI or shared.• When you use an existing disk to create a disk, the restrictions are as follows:<ul style="list-style-type: none">- When OceanStor V3/V5 is used, an EVS disk can be created from an existing EVS disk only after the administrator imports the HyperCopy license onto the storage device.- If the backend storage type is Dorado V3, the version must be Dorado V300R001C21 or later.- In VRM scenarios, when you use an existing disk to create a disk, the disk capacity cannot be configured and must be the same as that of the source disk. You cannot select an attached EVS disk. The disk type and disk mode cannot be changed, which are the same as those of the source disk.- In KVM scenarios, when you use an existing disk to create a disk, the disk capacity can be configured but must be greater than or equal to that of the source disk. The disk type and disk mode cannot be changed, which are the same as those of the source disk.- In VMware scenarios, when you use an existing disk to create a disk, the disk capacity can be configured but must be greater than or equal to that of the source disk. You cannot select an attached EVS disk. The disk type and disk mode cannot be changed, which are the same as those of the source disk.- If the disk capacity and disk type have been preset for the selected product, you can choose only a disk whose capacity is less than or equal to the preset disk capacity of the source disk, and the disk type of the disk must be the same as the preset disk type.- The source disk and the disk to be created must be in the same AZ.- New disks cannot be created when the source disk is in Reserved or Maintenance state.• When creating a disk using a snapshot, if the backend storage type is OceanStor V3/V5 series, the administrator needs to import the license of the HyperCopy feature on the device in advance.• Snapshots in one AZ cannot be used to create disks in another AZ.

Item	Restrictions
Attaching an EVS disk	<ul style="list-style-type: none">• The ECS supports the attaching of disks in VBD and SCSI modes.• Regardless if a shared EVS disk or non-shared EVS disk is attached to an instance, the EVS disk and the instance must be in the same AZ.• Data disks can only be attached to ECSs as data disks. System disks can be attached to ECSs as system disks or data disks.• An EVS disk cannot be attached to an instance that has expired.• An EVS disk cannot be attached to an instance that has been soft deleted.• When a disk is attached to an ECS configured with the disaster recovery (DR) service (CSDR/CSHA/VHA), you must ensure that the disk is created using the same backend storage as the existing disk on the ECS.• An EVS disk with snapshots of a VM can be attached only to the VM and cannot be attached to any other VM.• Neither shared EVS disks nor SCSI EVS disks can be attached to an ECS that has the CSHA service configured.• If the ECS uses the Windows operating system and the administrator set Disk Device Type to ide when registering the image, shut down the ECS before attaching the EVS disk to the ECS.• If the ECS to which the EVS disk belongs has not been created, the EVS disk cannot be attached to another ECS.

Item	Restrictions
Creating a snapshot	<ul style="list-style-type: none">• If backend storage is one of OceanStor V3/V5 or OceanStor Dorado V3 series, it is necessary for the administrator to import the HyperSnap license on the device in advance.• Snapshots can be created only for disks in the Available or In-use state.• A snapshot name cannot be the same as that of the prefix of the temporary snapshot created by the backup service (VBS/CSBS), the DR service (CSDR/CSHA/VHA), or the VM snapshot.• Snapshots created using the EVS console consume the capacity quota instead of quantity quota of EVS disks.• Temporary snapshots created by the backup service (VBS/CSBS) or the DR service (CSDR/CSHA/VHA) do not consume EVS disk quotas. Snapshots created using the VM snapshot function do not consume EVS disk quotas.• Snapshots created using the EVS console, temporary snapshots created by DR and backup services, and snapshots created using the VM snapshot function consume backend storage capacity. If a large number of snapshots are created, contact the administrator to set the thin provisioning ratio of backend storage to a large value, preventing EVS disk provisioning failures caused by excessive snapshots.• If backend storage of the disk is heterogeneous storage, snapshots can be created.• In the VRM or VMware scenario, no snapshots can be created for shared EVS disks.• If an EVS disk is created from data storage of the VIMS type in the VRM scenario and the EVS disk has been attached to an ECS, a snapshot can be created for the EVS disk only when the ECS is in the stopped state.• In the VMware scenario, no snapshots can be created for an EVS disk that has been attached to an ECS.• No snapshots can be created for disks that have expired.• No snapshots can be created for disks that have been soft deleted.• Snapshots cannot be created when the disk status is Reserved or Maintenance.• If a task for creating a snapshot fails, the task is automatically deleted.

Item	Restrictions
Rolling back a disk from a snapshot	<ul style="list-style-type: none">• Temporary snapshots created by the backup service (VBS/CSBS) or the DR service (CSDR/CSHA/ VHA) cannot be rolled back.• Snapshots created for disks having any DR service (CSDR/CSHA/ VHA) configured cannot be rolled back.• Snapshots created using the VM snapshot function cannot be used for EVS disk rollback.• After an EVS disk without VM snapshots is attached to a VM with VM snapshots, the EVS disk will be detached when the VM is rolled back using a VM snapshot.• If backend storage of the disk is heterogeneous storage (excluding XSKY), EVS disk rollback from a snapshot is not supported.• A snapshot can be used to roll back its source EVS disk, and cannot be used to roll back any other EVS disk.• A rollback can be performed only when the snapshot status is Available and the status of the snapshot source disk is Available (that is, the snapshot is not attached to any instance) or Rollback failed.• When the source disk of a snapshot is in the recycle bin, EVS disk rollback from the snapshot is not supported.
Creating a backup	<ul style="list-style-type: none">• Only disks in the Available or In-use state can be backed up.• If backend storage of the disk is heterogeneous storage, backups cannot be created.
Expanding EVS disk capacity	<ul style="list-style-type: none">• When you expand the capacity of a disk online, the instance to which the disk is attached must be in the Running or Stopped state.• Shared EVS disks do not support online capacity expansion, that is, the capacity of a shared EVS disk can be expanded only when the disk is in the Available state.• The capacity of a disk configured with the DR service (CSHA/ CSDR/VHA) cannot be expanded.• When the backend storage is Huawei SAN storage (OceanStor V3/V5 series or OceanStor Dorado V3 series/6.x) or heterogeneous storage, if the EVS disk has snapshots, capacity expansion is not supported. When the backend storage is Huawei Distributed Block Storage, capacity expansion can be performed for an EVS disk with snapshots.• If backend storage of the disk is heterogeneous storage, online capacity expansion is not supported while offline capacity expansion is supported.• Capacity expansion cannot be performed when the disk status is Reserved or Maintenance.

Item	Restrictions
Changing the disk type	<ul style="list-style-type: none">• Changing the disk type is supported when the backend storage is OceanStor V3/V5, OceanStor Dorado V3/6.x, or Huawei Distributed Block Storage.• If the backend storage is OceanStor V3/V5 or OceanStor Dorado V3/6.x, the disk type can be changed between different storage pools in the same storage system. If the backend storage is Huawei Distributed Block Storage, the disk type can be changed only in the same storage pool.• If the backend storage is Huawei Distributed Block Storage, the disk type can be changed only by modifying the QoS attribute.• The administrator needs to import the SmartMigration license on the device in advance if the backend storage is OceanStor V3/V5 or Dorado V3.• When changing the disk type, if the backend storage is OceanStor Dorado 6.x, the administrator needs to check whether the SmartMigration license has been imported to the device in advance. (The basic software package of OceanStor Dorado 6.x contains the SmartMigration license.)• You can change the type of the EVS disk only in the Available or In-use state.• If a disk has snapshots or is configured with the backup service (VBS/CSBS) or the DR service (CSDR/CSHA/VHA), the disk type cannot be changed.• If backend storage of the disk is heterogeneous storage, the disk type cannot be changed.
Extending the EVS disk validity period	<ul style="list-style-type: none">• If an EVS disk is created with an instance, the validity period of the EVS disk is unlimited.• If the validity period of an EVS disk is unlimited, the validity period cannot be extended.• You can extend the EVS disk validity period only when the disk is in the available or In-use status.• If an EVS disk has expired, its snapshot cannot be used to roll back the EVS disk or create an EVS disk. To continue using this EVS disk, extend its validity period.• When an EVS disk expires, its data will not be deleted. You can continue using this EVS disk after extending its validity period.

Item	Restrictions
Detaching an EVS disk	<ul style="list-style-type: none">• ECSs of the KVM virtualization type support online data disk detachment, namely, you can detach a data disk from an ECS in Running state. For details about online detachment restrictions, see Detaching an EVS Disk.• ECSs of the Xen virtualization type support online data disk detachment. However, ECSs of the VMware virtualization type do not support online data disk detachment.• System disks cannot be detached online.• Before detaching a disk online from an instance running Windows, log in to the instance to perform the offline operation and confirm that UVP VMTools has been installed on the ECS and is running properly. At the same time, ensure that this disk is not being read and written. Otherwise, the disk will fail to be detached.• Before detaching a disk online from an instance running Linux, log in to the instance, run the umount command to cancel the relationship between the disk and the file system, and confirm that the disk is not being read and written. Otherwise, the disk will fail to be detached.
Deleting an EVS disk	<ul style="list-style-type: none">• If a disk has been attached to an instance, the disk cannot be deleted.• If a disk has snapshots, the disk can be deleted only when the snapshot status is Available or Error.• You can delete a disk only when the disk status is Available, Error, Restoration failed, or Rollback failed, and no VM snapshot has been created for the ECS where the disk resides.• Disks configured with the DR service (CSDR/CSHA/VHA) cannot be deleted.• If an EVS disk has a snapshot, the EVS disk can be soft deleted only when the snapshot is in the Available or Error state.• When an EVS disk is permanently deleted, all snapshots of the EVS disk are also deleted.• A shared disk to be deleted must have been detached from all instances.• If the ECS to which the EVS disk belongs has not been created, the EVS disk cannot be deleted.• Local disks can be used as data disks or system disks for an ECS. When a local disk is used as the system disk or a data disk, its life cycle starts and ends with the ECS, and cannot be manually detached or deleted.

Item	Restrictions
Deleting a snapshot	<ul style="list-style-type: none">• Users are allowed to delete a temporary snapshot created by the backup service (VBS/CSBS). After the snapshot is deleted, if users want to back up the EVS disk corresponding to the snapshot, full backup is performed for the first time.• Temporary snapshots created by the DR service (CSDR/CSHA/VHA) cannot be deleted.• A snapshot created using the VM snapshot function cannot be deleted, and the name of the snapshot cannot be changed.• You can delete a snapshot only when its state is Available, Deletion failed, or Error.
Creating and associating a QoS	<ul style="list-style-type: none">• The QoS function is supported only in KVM and BMS virtualization scenarios.• IOPS and bandwidth upper limits can be set only when the backend storage is OceanStor V3/V5, OceanStor Dorado V3/6.x, or Huawei Distributed Block Storage.• The I/O priority can be set only when the backend storage is OceanStor V3/V5 or OceanStor Dorado V3/6.x.• A QoS policy cannot be associated with a disk type that has disks provisioned.• One disk type can be associated with only one QoS policy. One QoS policy can be associated with multiple disk types.• Before creating a QoS policy, if the backend storage is Huawei SAN storage, check on OceanStor DeviceManager that the SmartQoS license has been activated.

Item	Restrictions
Disk migration - advanced migration	<ul style="list-style-type: none">• Advanced migration applies to Huawei SAN storage (OceanStor V3/V5 and OceanStor Dorado V3/6.x) only. The source storage and target storage must be Huawei SAN storage and must meet the version requirements.• Disks in AZs whose virtualization type is KVM can be migrated offline (not attached to ECSs) and online (attached to ECSs). Only BMS disks can be migrated offline.• Before performing online migration, ensure that the corresponding compute node uses OceanStor UltraPath V200R001 or later as the multipathing software.• Only disks in the In-use or Available state can be migrated.• The source storage and target storage must be connected. The protocols of the links between the source storage and target storage, between the host and source storage, and between the host and target storage must be the same (FC or iSCSI protocol).• During migration, the source storage and target storage must be in the same AZ.• SCSI disks can be migrated no matter the ECS is powered on or off.• Disks attached to ECSs in the Running or Stopped state can be migrated, but the ECSs cannot have other services running.• Shared disks can be migrated.• Disks that have snapshots or disks attached to ECSs that have VM snapshots cannot be migrated.• Disks that have any DR service (CSDR/CSHA/VHA) configured cannot be migrated. You can perform the migration only after canceling the DR protection for the ECS and changing the configuration item Same Storage to No on the ECS page.• Disks that have any backup service (CSBS/VBS) configured cannot be migrated. Migration can be performed only after the backup service is stopped.• Disks attached to ECSs that have the VM HA function configured cannot be migrated. To perform migration, disable the VM HA function first.• If the target backend storage after migration is OceanStor V500R007C20/V300R006C30 or later, Dorado V300R002C00 or later, or Dorado 6.0.0 or later, the ECS to which the disk is attached supports the active-active configuration. Other versions do not support the active-active configuration.• Before the migration, check on OceanStor DeviceManager that SmartMigration license has been activated in the backend storage.• After the migration is complete, the disk has all features of the target disk type.

Item	Restrictions
	<ul style="list-style-type: none">• During the migration, do not perform other operations on disks. Do not power on or off the ECS. Do not configure DR services for the disk or ECS.• No more than three sets of source storage devices can be migrated to one set of target storage device. It is recommended that one set of source storage device be migrated to one target storage device.• The remaining capacity of the storage pool to which the disk to be migrated must be greater than 1% of the total capacity of the storage pool.• During disk migration, if a resource tag has been set for the source disk type, a resource tag must be set for the target disk type. Otherwise, disk migration is not supported.

Item	Restrictions
Disk migration - general migration	<ul style="list-style-type: none">• Migration can be implemented among Huawei SAN storage (OceanStor V3/V5 and OceanStor Dorado V3/6.x), Huawei Distributed Block Storage, and heterogeneous SAN storage.• Only disks in the AZs whose virtualization type is KVM can be migrated. The source storage and target storage of the migration must be in the same AZ.• Only attached disks can be migrated.• Disks no more than 2 TB and attached to ECSs in the Running or Stopped state can be migrated, but the ECSs cannot have other services running.• SCSI disks can be migrated only when ECSs are shut down.• Disks that have snapshots or disks attached to ECSs that have VM snapshots cannot be migrated.• Shared disks cannot be migrated.• Disks that have any DR service (CSDR/CSHA/VHA) configured cannot be migrated. You can perform the migration only after canceling the DR protection for the ECS and changing the configuration item Same Storage to No on the ECS page.• Disks that have any backup service (CSBS/VBS) configured cannot be migrated. Migration can be performed only after the backup service is stopped.• Migrate the selected disks on the same VM one by one. A maximum of two VMs can be migrated at a time on one physical host. This number of VMs can be changed. Options are one VM or two VMs (default). For details, see "Changing the Number of Volumes per Host to be Concurrently Cold or Live Migrated Across Storage" in <i>HUAWEI CLOUD Stack 8.1.0 O&M Guide</i>.• After the migration is complete, the disk has all features of the target disk type.• During the migration, do not perform other operations on disks. Do not power on or off the ECS. Do not configure DR services for the disk or ECS.• If the administrator sets Disk Device Type to ide when registering an image, the ECS provisioned using the image does not support disk migration.• During disk migration, if a resource tag has been set for the source disk type, a resource tag must be set for the target disk type. Otherwise, disk migration is not supported.


Table 3-6 Data storage functions of the VIMS type in VRM scenarios

Function	Sub-function	Supported or Not
Applying for an EVS disk	Creating a blank EVS disk	Yes
	Applying for an EVS disk from a snapshot	Yes
	Applying for an EVS disk from a backup	No
	Creating an EVS disk using an existing disk	Yes
	Specifying the shared EVS disk	Yes
	Specifying the EVS disk of the SCSI type	No
	Specifying the AZ, capacity, name, disk type, and quantity	Yes
	Specifying the validity period	Yes
Managing an EVS disk	Querying the EVS disk list and details	Yes
	Changing the disk type	No
	Changing the name and description of an EVS disk	Yes
	Soft-deleting the EVS disk to the recycle bin and deleting it from the recycle bin	Yes
	Attaching and uninstalling EVS disks	Yes
	Extending the validity period of an EVS disk	Yes
Expanding EVS disk capacity	Expanding the disk capacity online	No
	Expanding the disk capacity offline	Yes
Snapshot	Creating a snapshot	Yes
	Restoring the EVS disk from a snapshot	No
	Deleting a snapshot	Yes
Backup	Creating a backup	No
Quota	Specifying the capacity and number of EVS disks	Yes
Specifying the EVS disks metering	Performing metering by capacity, disk type, or duration	Yes

Function	Sub-function	Supported or Not
Snapshot metering	Performing metering by capacity or duration	Yes

3.1.8 Accessing and Using the Cloud Service

Two methods are available:

- Web UI
Log in to ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios) as a tenant, click  in the upper left corner of the page, select a region, and select the cloud service.
- API
Use this mode if you need to integrate the cloud service into a third-party system for secondary development. For details, see "Operation Help Center" > "Storage Services" > "Elastic Volume Service" > "API Reference".

3.2 Related Concepts

3.2.1 Device Type

Definition

Device types of EVS disks are divided based on whether advanced SCSI commands are supported. The device type can be VBD or SCSI.

- VBD: EVS disks of this type support only basic SCSI read and write commands. They are used in common scenarios, for example, OA, tests, Linux clusters such as RHCS.
- SCSI: EVS disks of this type support transparent SCSI command transmission and allow the ECS operating system to directly access the underlying storage media. SCSI EVS disks support advanced SCSI commands (such as SCSI-3 persistent pre-lock) in addition to basic SCSI read and write commands. They can be used in cluster scenarios where data security is ensured by using the SCSI lock mechanism, such as the Windows MSCS cluster.

NOTE

For details about ECS operating systems supported and ECS software required by SCSI EVS disks, see [Requirements and Restrictions on Using SCSI EVS Disks](#).

Requirements and Restrictions on Using SCSI EVS Disks

- If the VM HA function, storage plane anti-split-brain function, and data disk protection function are enabled, only non-SCSI disks can be locked to prevent data disks from being written in dual-write mode. (For details about how to check whether the functions are enabled, see "Configuring the VM HA Function" in *HUAWEI CLOUD Stack 8.1.0 O&M Guide*). Lock protection is not

supported on SCSI data disks. Lock protection is implemented based on SCSI commands, but SCSI disks support transparent SCSI command transmission. As a result, lock protection on SCSI disks may conflict with user operations on SCSI disks, resulting in task failure.

- When SCSI disks are attached to an ECS, check whether the ECS supports SCSI disks based on the following description:
 - Windows OS
 - i. Check whether the ECS operating system supports the SCSI feature. Obtain the ECS operating system version by referring to *FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Private Cloud)* and check whether the SCSI (virtio-scsi) or raw device mapping feature is supported. Obtain the document by referring to the following note.
 - ii. The Windows operating system must have UVP VMTools installed to support SCSI.
Generally, the administrator has installed the UVP VMTools when creating a public image. You do not need to install it manually.

If the Windows operating system supports the SCSI feature and UVP VMTools has been installed in the operating system, you can attach SCSI disks to the ECS.

- Linux OS

The SCSI feature of the Linux operating system is not provided by the UVP VMTools but the driver in the operating system. Therefore, you only need to check whether the current ECS operating system supports the SCSI feature.

The check method is as follows: Check the ECS operating system version by referring to *FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Private Cloud)* and check whether the SCSI (virtio-scsi) or raw device mapping feature is supported.

If the Linux operating system supports the SCSI feature, you can attach SCSI disks to the ECS. Obtain the document by referring to the following note.

NOTE

Obtain *FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Private Cloud)*:

- x86
 - Carrier users: [Click here](#). Search for **FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Private Cloud)**.
 - Enterprise users: [Click here](#). Search for **FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Private Cloud)**.
- Arm
 - Carrier users: [Click here](#). Search for **FusionSphere SIA Huawei Guest OS Compatibility Guide (ARM)**.
 - Enterprise users: [Click here](#). Search for **FusionSphere SIA Huawei Guest OS Compatibility Guide (ARM)**.

3.2.2 Disk Type

Definition

A disk type can be selected during disk creation. A disk type represents backend storage devices used by a group of disks. You can divide disk types of EVS disks based on backend storage types to meet different performance requirements of services.

Based on performance differences of backend storage used by disks, typical disk types and their application scenarios are as follows:

- Common performance: EVS disks of this type are suitable for scenarios that require large capacity, medium-level read and write speed, and relative fewer transactions, such as the scenario for deploying development and test applications.
- Medium performance: EVS disks of this type are suitable for scenarios that require common performance but rich enterprise-class features. They can be used in common databases, application VMs, and middleware VMs.
- High performance: EVS disks of this type are suitable for scenarios that require high performance, fast read and write speed, and large throughput, such as data warehouses.
- Ultra-high performance: EVS disks of this type are suitable for data-intensive scenarios that require very high I/O performance, such as NoSQL and relational databases.

Changing a Disk Type

When the read and write performance of the storage device where the upper-layer service resides no longer suits the service, you can change the disk type to alter the type of the storage device to change the read and write performance, meeting the requirements of varying service performance of the instance. Examples are as follows:

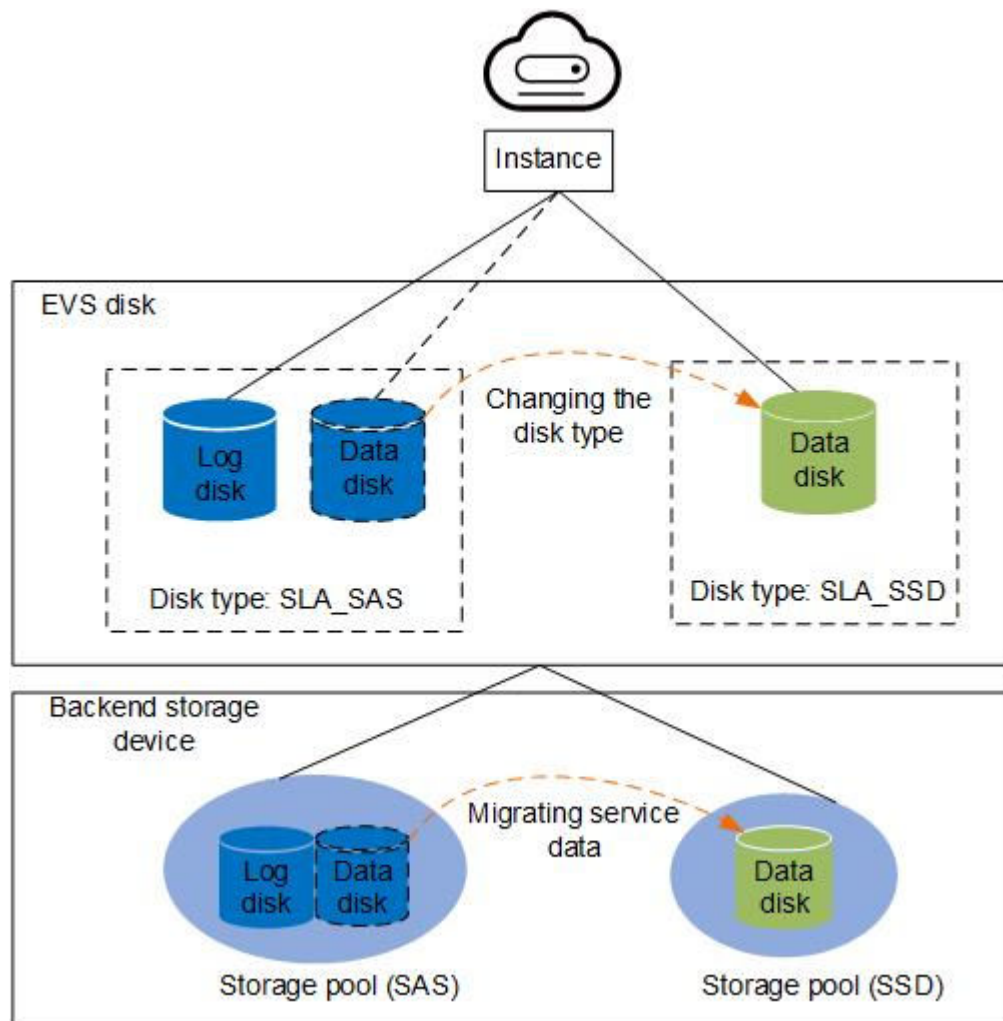
- When your service requires a higher read and write performance, you can migrate your service from disks created on low-speed storage media to disks created on high-speed storage media to improve the read and write performance.
- If the priority of the performance of a service degrades, you can migrate your service to disks created on low-performance storage media. This helps release storage resources for high-performance disks for other services.

You can change the disk type of an in-use disk (a disk that has been attached to an instance). You can also detach a disk from the instance, and then change the disk type of the disk.

If you change the disk type of an in-use EVS disk, the service of the source EVS disk on the instance will be migrated to the destination EVS disk without interrupting host services. After service migration, the destination EVS disk replaces the source EVS disk to run the service, without any adverse impact on customer experience. However, when you change the disk type of an in-use EVS disk, the performance of the instance is adversely affected to some extent.

Figure 3-9 shows the implementation principle of changing a disk type. In the following figure, two disks are attached to an instance. One of the disks serves as a log disk, and the other serves as a data disk. The original disk type of the two disks is SLA_SAS. Because the service has a higher performance requirement on the data disk, the disk type of the data disk is changed from SLA_SAS to SLA_SSD, seamlessly migrating service data to a disk of the target disk type. The backend storage device performs service data migration. After service data migration, the system automatically attaches the destination disk to the instance, without service interruption. In addition, the source disk will be deleted to release storage resources for other services.

Figure 3-9 Implementation principle of changing a disk type



3.2.3 EVS Disk Data Redundancy Protection

The backend storage of EVS disks supports multiple storage types, including Huawei SAN, Huawei Distributed Block Storage, and heterogeneous storage. A storage array provides data redundancy protection. When a physical device is faulty, data on the faulty device can be automatically restored and data can still be accessed. For example: Huawei SAN storage uses the RAID technology and allows one or two disks to fail at the same time. Huawei Distributed Block Storage uses the three-copy technology to create two identical copies for each piece of

data. The data and copies are stored on different storage nodes. The Erasure Coding (EC) technology is added to OceanStor Pacific series 8.1.1. When the storage type is OceanStor Pacific series 8.1.1, EC is used as the default data redundancy mode. This section uses the EC technology as an example to describe EVS disk data redundancy protection. For more information about storage device redundancy protection, see the product documentation of the corresponding storage device.

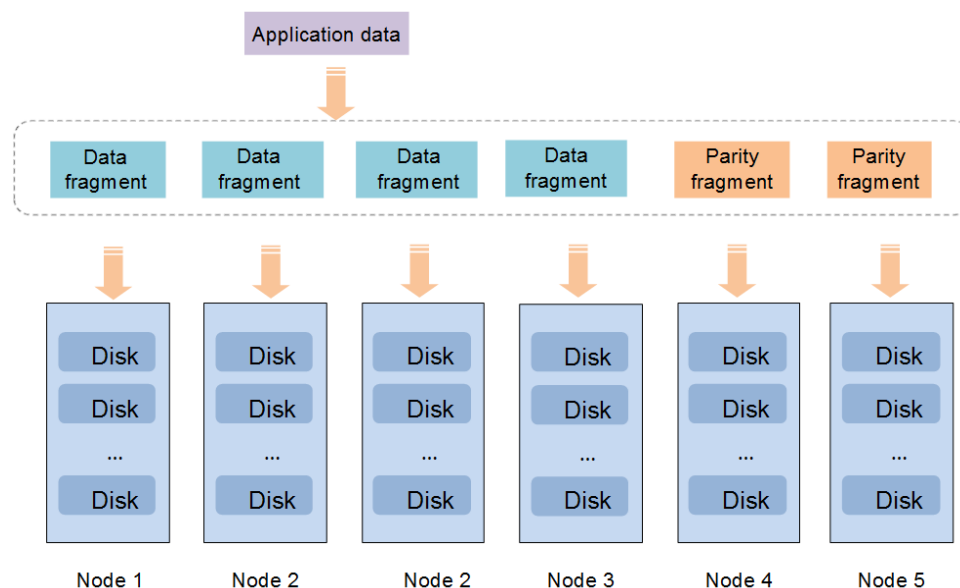
EC Principle

EC protection levels are expressed in N+M mode. N indicates the number of data fragments (data strips) and M indicates the number of parity fragments (parity strips). Huawei Distributed Block Storage supports N+2, N+3, and N+4 protection levels and allows a maximum of four storage nodes or cabinets to fail at the same time. The minimum node quantity and corresponding disk utilization rate vary by protection level. For details, see "General Information" > "Product Description" > "Functions and Features" > "Key Features" > "Block Service" > "EC and Multi-Copy" in *OceanStor Pacific Series 8.1.1 Product Documentation*.

The following uses the N+2 level (N is 4) as an example to describe the EC implementation process.

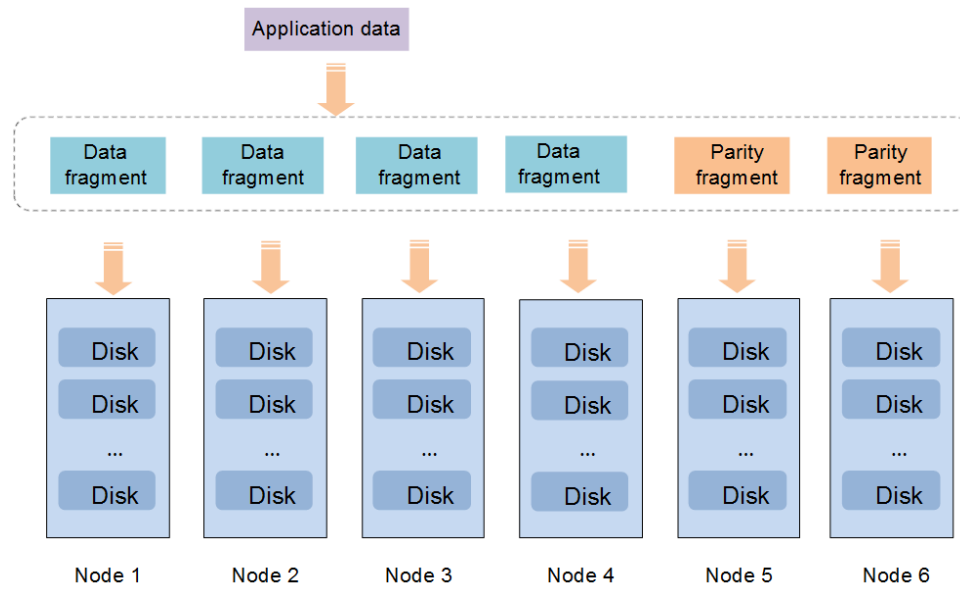
- **Figure 3-10** shows how data protection is implemented when the number of storage nodes is greater than or equal to N+2. In this example, two storage nodes or cabinets can be faulty at the same time.
- **Figure 3-11** shows how data protection is implemented when the number of storage nodes is greater than or equal to (N+2)/2 and less than N+2. In this example, one storage node or cabinet is allowed to fail at the same time, and two disks or storage nodes are allowed to fail at the same time on a single storage node.

Figure 3-10 Huawei Distributed Block Storage EC technology (N+M configuration mode)



1. Divide written data into four data fragments.
2. Calculate and generate two parity fragments based on the four data fragments.
3. Write all the data fragments and parity fragments to the six storage nodes in the redundancy ratio mode.

Figure 3-11 Huawei Distributed Block Storage EC technology (N+M:1 configuration mode)



1. Divide written data into four data fragments.
2. Calculate and generate two parity fragments based on the four data fragments.
3. Write all the data fragments and parity fragments to the three storage nodes in the redundancy ratio mode.

Huawei Distributed Block Storage EC data fragmentation improves data read and write performance and ensures high disk utilization. If no more than M disks in a Huawei Distributed Block Storage cluster are faulty, the system can quickly restore lost data by reconstructing data among nodes. This ensures high data reliability and service availability.

3.2.4 Shared Disk

In the traditional cluster architecture, multiple computing nodes need to access the same data. This ensures that when a single or multiple computing nodes are faulty, the HA cluster can continue providing services, which means that a faulty component will not cause service interruption. Therefore, important data files need to be stored on shared block storage, and shared block storage is centrally managed using the cluster file system. When multiple frontend computing nodes access data, the data will be the same on the multiple computing nodes.

The shared disk is designed for the core service HA architecture of enterprise customers. The shared disk is suitable for scenarios that require shared block storage access in the share-everything architecture. The scenarios include the HA

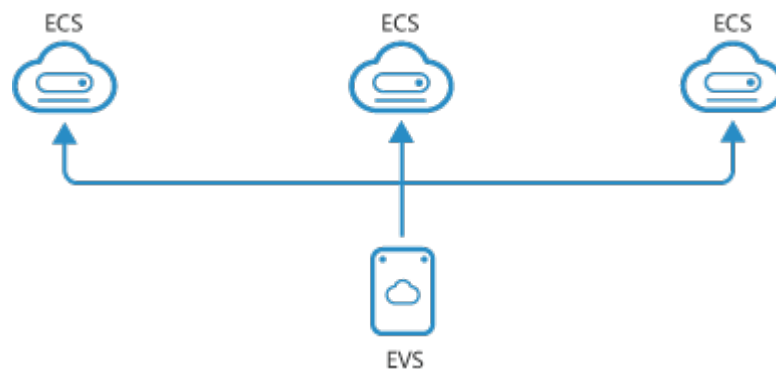
Oracle RAC database architecture for government, enterprise, and finance customers and the HA server cluster architecture.

Definition

Shared EVS disks are block storage devices that support concurrent read/write operations of multiple ECSs/BMSs. Shared EVS disks feature multiple attachments, high-concurrency, high-performance, and high-reliability. A shared EVS disk can be attached to a maximum of 16 ECSs/BMSs. A non-shared EVS disk can be attached to only one ECS/BMS. This document uses ECS as an example, as shown in [Figure 3-12](#).

Currently, shared EVS disks can be used as data disks only and cannot be used as system disks. Shared EVS disks of the VBD or SCSI type can be created.

Figure 3-12 Shared EVS disk



SCSI Reservation

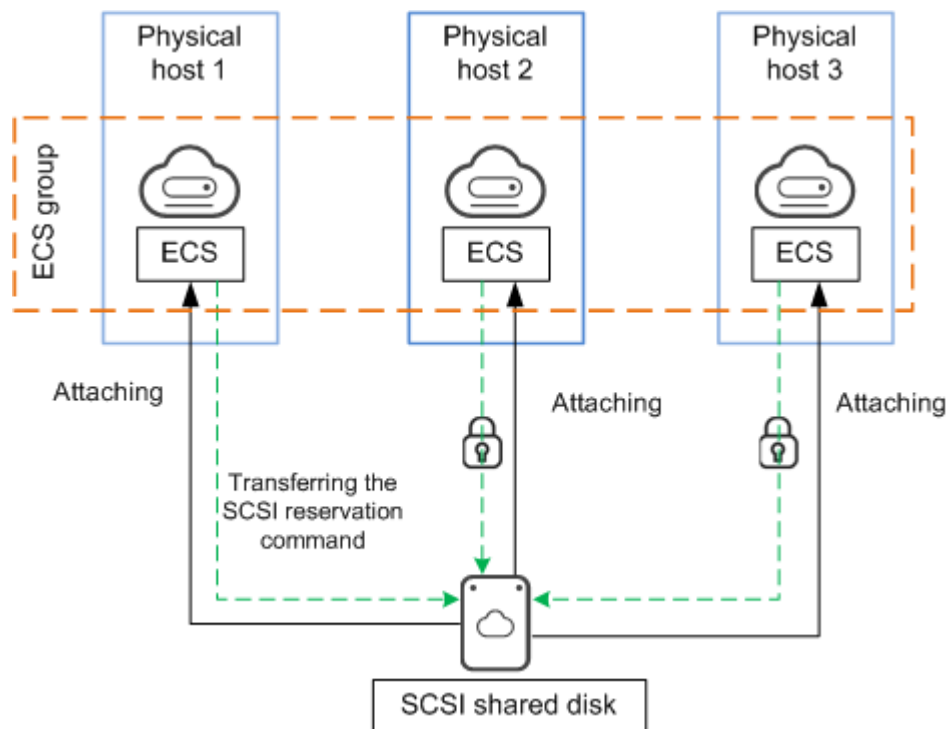
Shared EVS disks of the VBD type do not support SCSI locks. SCSI shared EVS disks support SCSI reservation. If SCSI reservation is required for your applications, create SCSI shared EVS disks.

SCSI reservation is the basic mechanism for multiple hosts to use disks. In a shared storage environment, multiple service hosts may access a disk simultaneously. If multiple hosts perform the write operation on the disk at the same time, the disk does not know data from which host will be written first. To prevent this problem that may cause data damage, SCSI reservation is introduced.

SCSI reservation for an EVS disk cannot distinguish multiple ECSs on a single physical host, and SCSI reservation is supported only when ECSs are deployed on different physical hosts. If a SCSI shared EVS disk is attached to ECSs, use anti-affinity ECS groups with SCSI reservation so that SCSI reservation takes effect.

The ECSs in an anti-affinity ECS group will be created on different physical hosts, thereby improving service reliability. You can add an ECS to an ECS group only when creating the ECS. An existing ECS cannot be added to any ECS group.

[Figure 3-13](#) shows how SCSI reservation is implemented. When a SCSI shared disk is attached to multiple ECSs in an anti-affinity ECS group, if one of the ECSs sends a SCSI reservation command to the SCSI shared disk, the SCSI shared disk is locked for the other ECSs. In this case, the other ECSs cannot write data into the SCSI shared disk.

Figure 3-13 SCSI reservation implementation mechanism**NOTICE**

If an ECS does not belong to any ECS group, the SCSI shared EVS disk should better not be attached to the ECS. Otherwise, SCSI reservations may not work properly, which puts your data at risk.

Precautions for Using the Shared EVS Disk

A shared EVS disk is essentially the disk that can be attached to multiple instances 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 and write rules, such as the read and write sequence and meaning, between these servers are not defined, data read and write interference between servers or other unpredictable errors may occur.

Shared EVS disks provide block storage devices whose data can be randomly read or written and allows shared access. Shared EVS disks do not provide the cluster file system. You need to install the cluster file system to manage shared EVS disks.

If a shared EVS disk is attached to multiple instances but is managed using a common file system, disk space allocation conflict will occur and data files will be inconsistent. The details are as follows:

- **Disk space allocation conflict**
Suppose that a shared EVS disk is attached to multiple instances. When a process on instance A writes files into the shared EVS disk, it checks the file system and available disk space. After files are written into the shared EVS disk, instance A will change its own space allocation records, but will not

change the space allocation records on the other instances. Therefore, when instance B attempts to write files to the shared EVS disk, it may allocate disk space addresses that have been allocated by instance A, resulting in disk space allocation conflict.

- Inconsistent data files

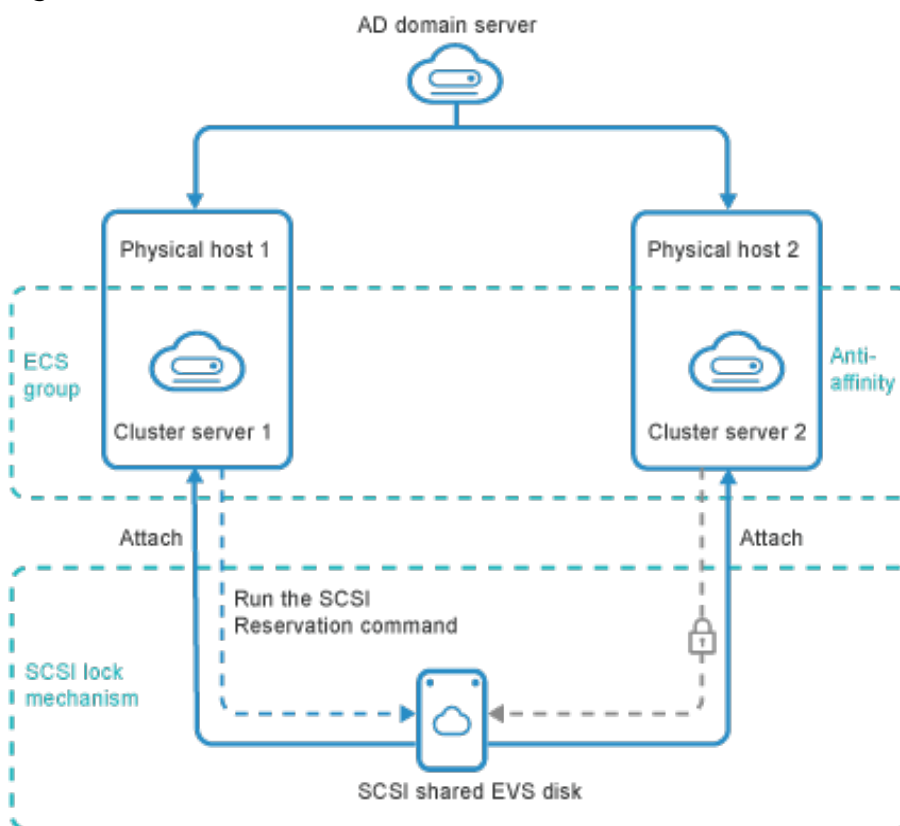
Suppose instance A reads data and records it in the cache. When another process on instance A accesses the same data, the process will read the data directly from the cache. If instance B changes the data, instance A will not know and will read the data from the cache. As a result, service data will be inconsistent on instance A and instance B.

The cluster management system is used to manage the shared EVS disks. If the cluster needs to use SCSI reservations, you need to apply for a shared EVS disk of SCSI type. Example enterprise applications include Windows MSCS (Microsoft Cluster Service) and Linux RHCS (Red Hat Cluster Suite). Windows MSCS and Linux RHCS are used as an example to describe how to use shared EVS disks.

- Windows MSCS

Figure 3-14 shows the Windows MSCS diagram. Multiple nodes in the cluster share the same storage. The cluster needs to use the SCSI reservation. Therefore, the shared EVS disk of the SCSI type is used. When a node in the cluster is faulty, services on the node are switched to another available node.

Figure 3-14 Windows MSCS

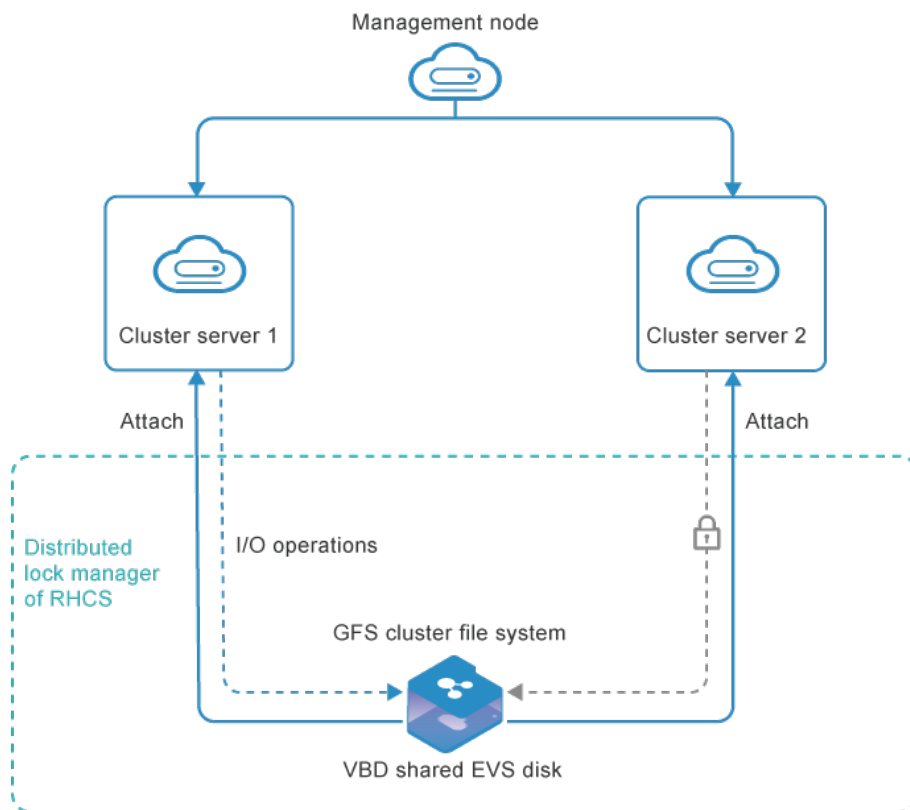


- Linux RHCS

Figure 3-15 shows the Linux RHCS diagram. Linux RHCS is a cluster suite that provides high availability, load balancing, and storage sharing. Linux RHCS provides a distributed lock manager. The GFS uses the lock mechanism of the

lock manager to ensure data consistency when multiple nodes share the same disk. Therefore, a shared EVS disk of VBD type without SCSI reservation can be used in the cluster.

Figure 3-15 Linux RHCS



3.2.5 EVS Disk Snapshot

Definition

EVS disk snapshot is an important data recovery method that records the status of EVS disk data at a specific point in time. The snapshot created for an EVS disk at a certain point in time is independent from the life cycle of the EVS disk. The snapshot can be used to roll back and restore data of the EVS disk at the time when the snapshot was taken.

You can create an EVS disk from a snapshot. The created EVS disk contains the data of the snapshot, and is a precise copy of the source EVS disk. An EVS disk created from a snapshot does not need to be partitioned or formatted, and no file system needs to be created. When the EVS disk is attached to an instance, the EVS disk can read and write data. Therefore, the snapshot is an important way of sharing and migrating data.

Comparison Between EVS Snapshot and VBS

Both EVS snapshot and backup can protect EVS data. [Table 3-7](#) describes the differences between them.

Table 3-7 Comparison between EVS snapshot and VBS

Item	EVS Disk Snapshot	VBS
DR methods	A snapshot records the status of an EVS disk at a specific point in time. You can roll back and restore data of the EVS disk at the time when the snapshot was taken. A snapshot is not an actual disk data copy. If the disk is physically damaged, data cannot be restored using the snapshot rollback function. In this case, backup can be used.	A backup is a copy of EVS disk data at a certain point in time Backup is implemented based on the snapshot technology and snapshot comparison technology. A disk is backed up in full backup mode only when they are initially backed up. After the full backup, the disk is backed up in incremental backup mode.
Storage mode	Snapshot data is stored with disk data, consuming the capacity quota instead of quantity quota of EVS disks. The storage space consumed by snapshots is small. Therefore, snapshots can be created faster than backups.	The backup data is stored in the backup storage. It consumes no quota of EVS disks.
Periodic snapshot/backup	Periodic snapshots are not supported. Snapshots have to be created manually.	Periodic backup is supported. You can create a backup policy to automatically back up disks based on the backup policy.
DR range	Snapshots are region-specific. Snapshots, and the corresponding source disks belong to the same AZ. When you create an EVS disk using a snapshot, the new EVS disk and the snapshot are in the same AZ.	EVS disks and backups can be in different AZs. Backup copies can be used to restore EVS disk data across AZs.
Service restoration	You can use the snapshot to roll back the disk to the time when the snapshot was created, or create an EVS disk using the snapshot to retrieve the data at the time when the snapshot was created.	You can restore the backup to the original disk or another disk, or create an EVS disk from the backup to retrieve the data at the backup time point.
Deployment mode	The snapshot function can be used after EVS is deployed.	The advanced cloud service VBS must be deployed.

Application Scenarios

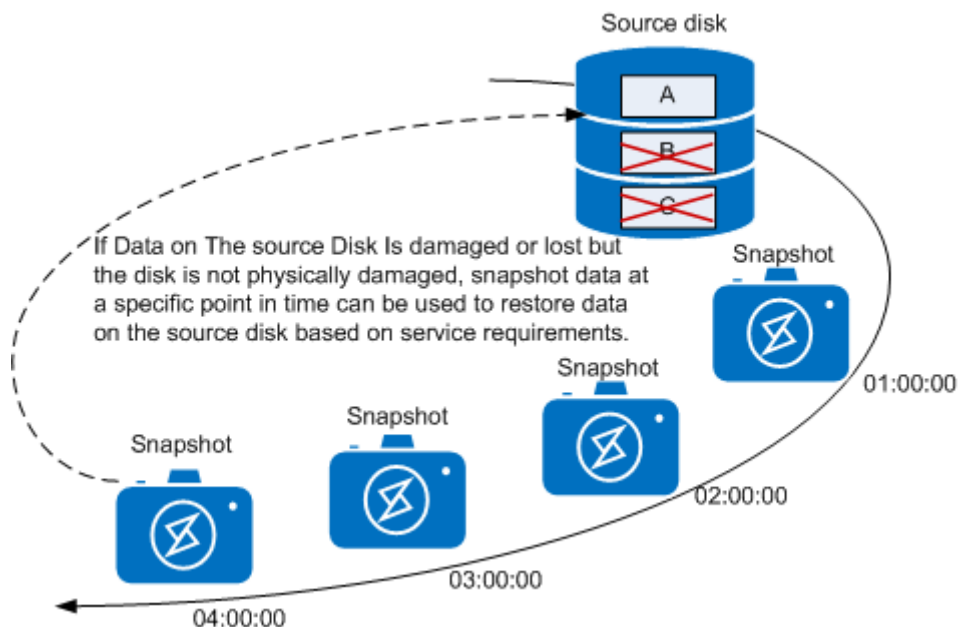
The snapshot is a convenient and efficient means of data protection, and it is recommended that you use this means of data protection in the following scenarios:

- Routine data backup and restoration

Snapshots are used to periodically back up important service data on system disks and data disks to prevent data loss caused by misoperations, attacks, or viruses.

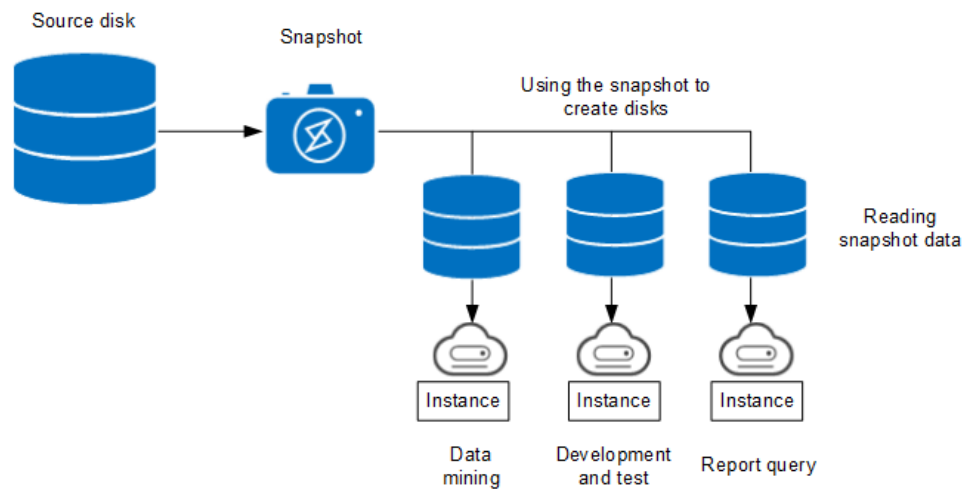
When data loss or data inconsistency occurs on an EVS disk due to misoperations, viruses, or hacker attacks, you can use a snapshot to restore a previous normal status of the EVS disk. In addition, you are advised to create disk snapshots before a big change (such as application software upgrade and service data migration). If the operation fails, you can roll back the snapshots to restore service data, as shown in [Figure 3-16](#).

Figure 3-16 Using snapshots for routine data backup and restoration



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, such as data mining, report query, and development and test. This method protects the initial data and creates EVS disks rapidly, meeting the diversified service data requirements. [Figure 3-17](#) shows the procedure for using a snapshot to deploy multiple services.

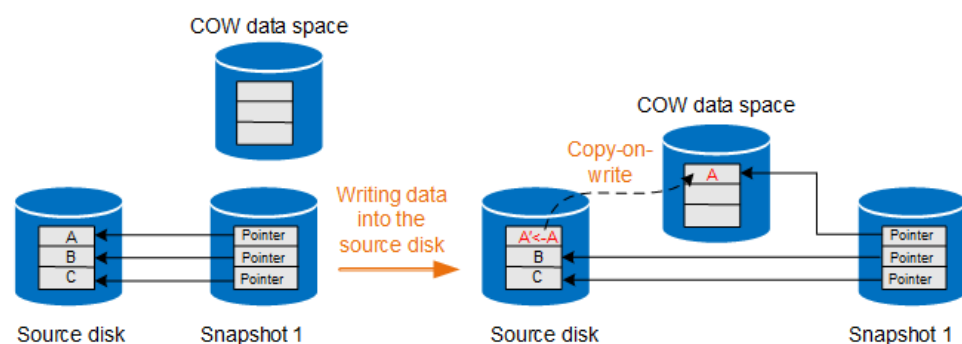
Figure 3-17 Using a snapshot to deploy multiple services

Implementation Principles

The snapshot implementation principle varies with the type of backend storage where the disk resides. Snapshot implementation principles for different backend storage types are described as follows:

- **OceanStor V3 or OceanStor V5 series as backend storage**

A snapshot is a copy of source disk data, which is generated at a specific time. A snapshot consists of a source disk, Copy-on-Write (COW) data space, and snapshot data. Snapshots are implemented using the mapping table and COW technology. [Figure 3-18](#) shows the snapshot implementation principle.

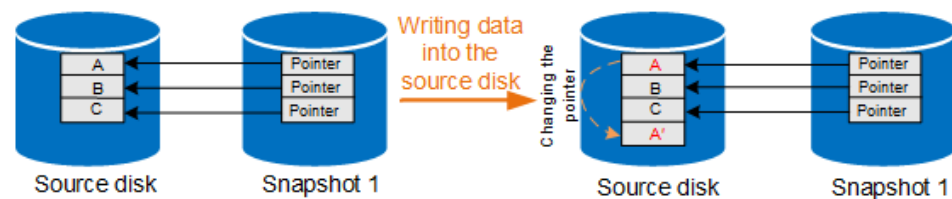
Figure 3-18 Snapshot implementation principle (OceanStor V3 or OceanStor V5 series as backend storage)

- Before creating a snapshot: When no snapshot is created for a disk, the procedure for writing data into the disk is the same as the procedure for writing data into other disks. Data changes will be directly written into disk data blocks, overwriting the original data, and the original data will not be retained.
- After creating a snapshot: After a snapshot is created, a data copy that is identical to the source disk is generated. In this step, the backend storage system dynamically allocates COW data space in the storage pool where the source disk resides, and automatically generates a snapshot. The pointer of the snapshot points to the storage location of source disk data.

- Writing data into the source disk: When an instance sends a request to write data into the source disk, the backend storage system will not write the new data immediately. Instead, the backend storage system employs the COW mechanism to copy the original data from the source disk to the COW data space, modifies the mapping in the mapping table, and writes the new data to the source disk. As shown in [Figure 3-18](#), when data A of the source disk needs to be changed, data A will be copied to the COW data space, and then the snapshot pointer will be changed to point to the storage location of data A in the COW data space. Finally, data A' will be written into the source disk.
- **Dorado V3 series as backend storage**

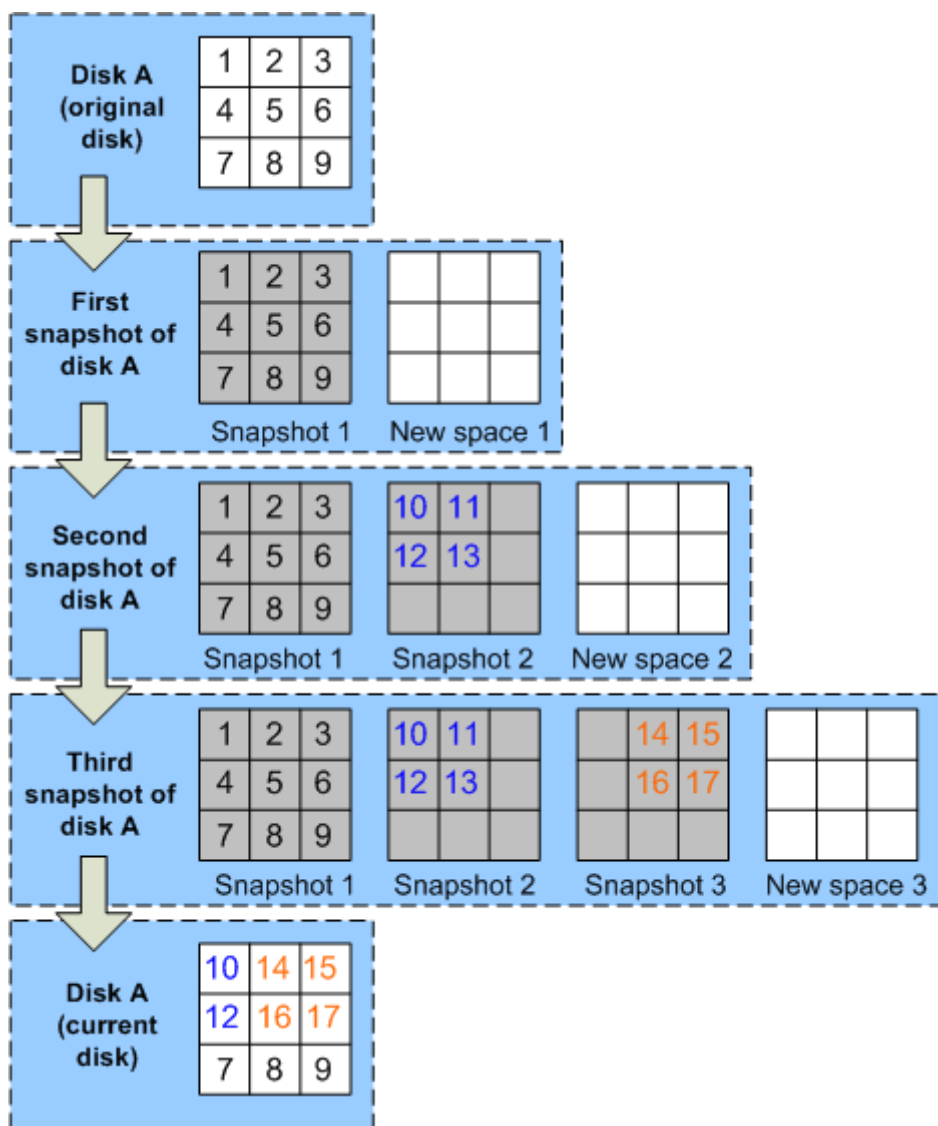
The core technology in snapshot implementation is Redirect-on-Write (ROW). [Figure 3-19](#) shows the snapshot implementation principle.

Figure 3-19 Snapshot implementation principle (Dorado V3 series as backend storage)



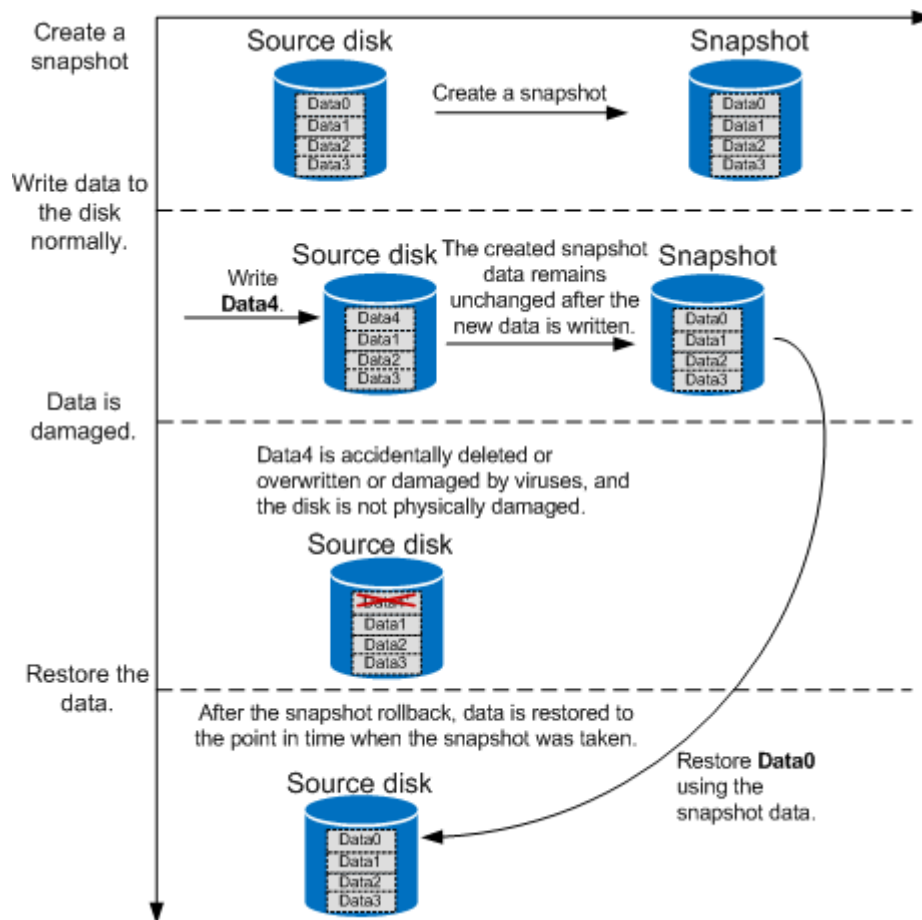
- Before creating a snapshot: When no snapshot is created for a disk, the procedure for writing data into the disk is the same as the procedure for writing data into other disks. Data changes will be directly written into disk data blocks, overwriting the original data, and the original data will not be retained.
- After creating a snapshot: After a snapshot is created, a data copy that is identical to the source disk is generated. In this step, the backend storage system copies the pointer of the source disk to the snapshot, and the pointer of the snapshot points to the storage location of source disk data.
- Writing data into the source disk: When an instance sends a request to write data into the source disk after a snapshot is created, the storage system uses the ROW technology to save the new data to a new location and changes the pointer of the source disk to point to the storage location of the new data. The pointer of the snapshot still points to the storage location of the original data. The source disk data at the time when the snapshot was created is saved. As shown in [Figure 3-19](#), when data A of the source disk needs to be changed, data A' (new data) will be written into a new location, and the pointer of the source disk will be changed to point to the storage location of data A'. The pointer of the snapshot still points to the storage location of data A (original data).
- **Huawei Distributed Block Storage as backend storage**

Snapshot data is based on the Distributed Hash Table (DHT) mechanism. The following shows the snapshot implementation principle.

Figure 3-20 Snapshot implementation principle (Huawei Distributed Block Storage as backend storage)

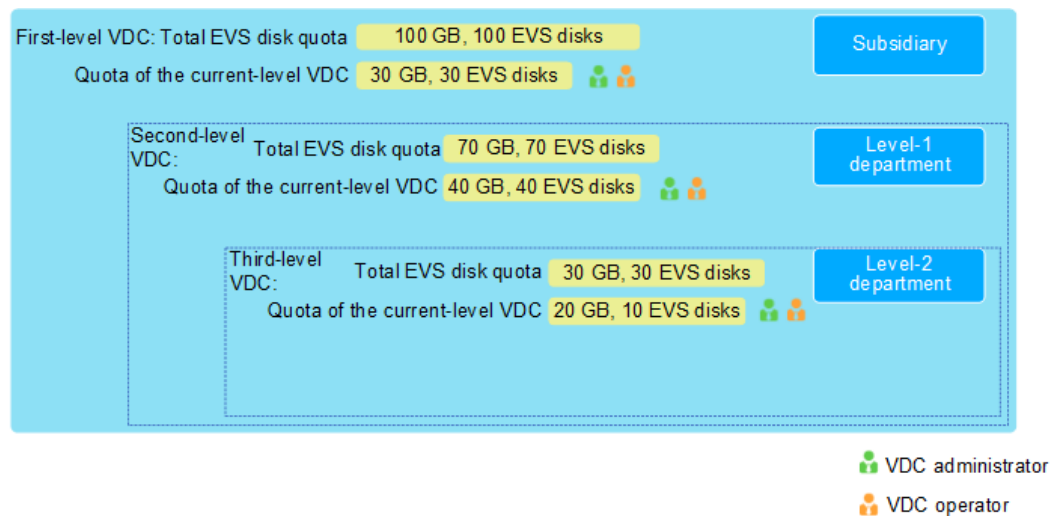
Rolling Back a Disk from a Snapshot

Snapshot rollback is a mechanism for quickly restoring data on the source disk by using the snapshot of the source disk at a certain point in time. If the data on the source disk is accidentally deleted, damaged, or infected by viruses and the source disk is not physically damaged, you can use the snapshot rollback function to quickly restore data on the source disk at the point in time when the snapshot was taken, reducing the amount of data lost. [Figure 3-21](#) shows snapshot rollback process.

Figure 3-21 Snapshot rollback

3.2.6 EVS Disk Quota

A Quota is a resource management and control technology that limits the maximum number of resources (including the resource capacity and number of resources) that can be used by a single VDC, preventing resources from being overused by users in some VDCs and affecting other VDCs. When creating a first-level VDC, the operation administrator can set the total quota (capacity and quantity) of EVS disks in the VDC and the EVS disk quota of the current-level VDC. When creating a lower-level VDC, the VDC administrator can set the total quota of EVS disks in the lower-level VDC and the EVS disk quota of the current-level VDC. [Figure 3-22](#) shows the quota of EVS disks in VDCs of different levels.

Figure 3-22 EVS disk quota

There are three levels of VDCs in the figure.

- Users in the VDC of each level can use EVS disk resources in the quota of the current-level VDC.
- The maximum quota of second-level VDCs is the maximum quota of first-level VDCs minus the quota allocated to the first-level VDCs.
- The maximum quota of third-level VDCs is the maximum quota of second-level VDCs minus the quota allocated to the second-level VDCs.

NOTE

You can create multiple VDCs at the same level and allocate quota to each VDC. For example, when multiple second-level VDCs are created, the maximum quota of all the second-level VDCs is the maximum quota of first-level VDCs minus the quota allocated to the first-level VDCs.

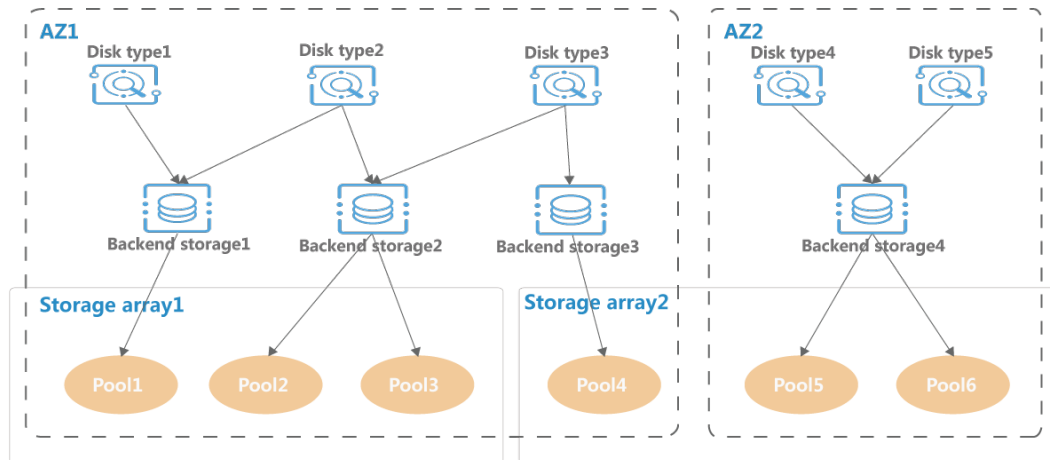
3.2.7 Relationship Between the Disk Type, Backend Storage, and Storage Array

Definition

- **Disk Type**
A disk type can be selected during disk creation. A disk type represents backend storage devices used by a group of disks.
For more information about disk types, see [Disk Type](#).
- **Backend Storage**
Backend storage is a logical storage device that stores EVS disk resources. A backend storage device contains one or more storage pools on a storage array (Huawei Distributed Block Storage, Huawei SAN storage, or heterogeneous storage).

Relationship

Figure 3-23 shows the relationship between the disk type, backend storage, and storage array

Figure 3-23 Relationship between the disk type, backend storage, and storage array

- A storage array may belong to different availability zones. For example, storage pool 4 on storage array 2 belongs to AZ 1 while storage pools 5 and 6 belong to AZ 2.
- A backend storage device belongs to only one AZ.
- A storage pool belongs to only one backend storage device.
- A disk type belongs to only one AZ.
- A backend storage device can contain one or more storage pools, but the storage pools must be on the same storage array, for example, backend storage 2.
- A disk type can contain multiple backend storage devices from the same AZ. The storage devices can be on the same or different storage arrays but the arrays must be in the same AZ.
- Multiple disk types can be created for a backend storage device, for example, backend storage 1. Different disk types can be configured with different value-added features, such as SmartThin, SmartTier, and SmartDedupe.

3.2.8 Mapping Between Mount Points and Device Names

A block storage device is a storage device that moves data in sequences by bytes or bits (blocks). These devices support random access and wide use of cache I/O, including disks, CD-ROM, and flash drives. A block storage device can be attached to a computer or remotely accessed as it is attached to a computer. The instance supports the following block storage devices:

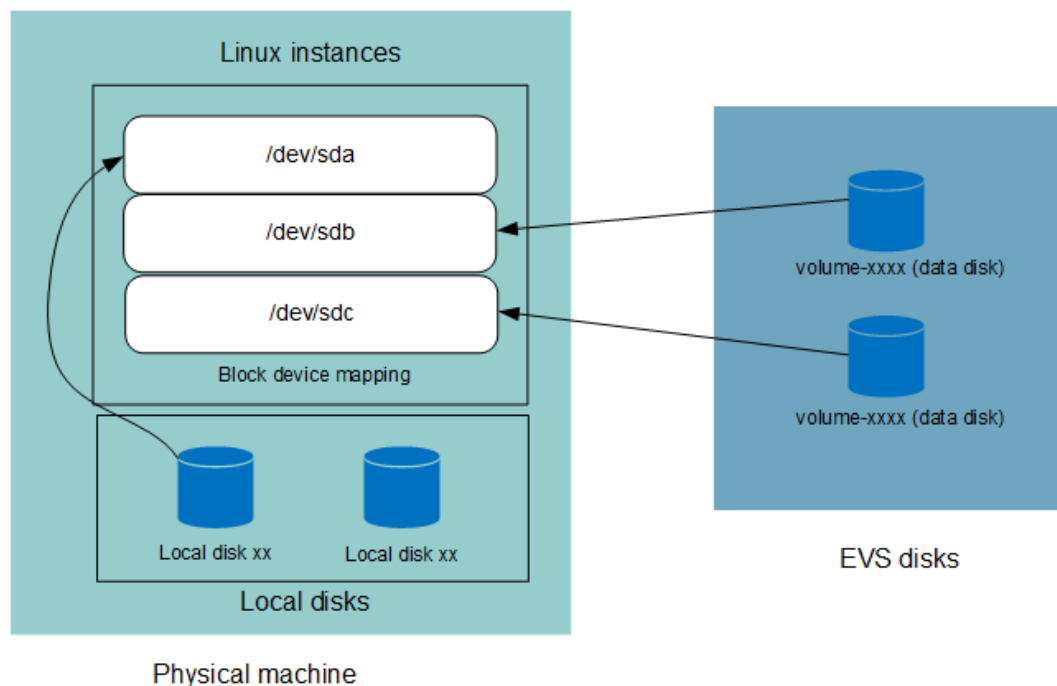
- Local disk: is the disk that is attached to the physical machine (host machine) where the instance is located and is a temporary block storage device.
- EVS disk: is a cloud disk that is attached to an instance and is a persistent block storage device.

The attachment point is the entry directory of the disk file system in Linux. It is similar to the drive letters, such as **C:**, **D:**, and **E:**, which are used to access different partitions in Windows. Each attachment point corresponds to a device name. You can attach the corresponding disk to an instance by specifying the device name of the attachment point.

Block Storage Device Mapping

The instance uses the device name (for example, `/dev/sdb`) to describe the block storage device and uses the block storage device mapping to specify the block storage device to be attached to the instance. [Figure 3-24](#) shows an example of mapping between EVS disks as well as local disks and instances. In the preceding figure, one local disk is attached to the Linux instance, the local disk is mapped to `/dev/sda` as the system disk, and two EVS disks are mapped to `/dev/sdb` and `/dev/sdc`, respectively, as data disks.

Figure 3-24 Example of mapping between EVS disks as well as local disks and instances



3.2.9 Disk Migration

Disk migration ensures that disks can be migrated within a storage array or between storage arrays in the same AZ without interrupting services to precisely match service requirements. Two disk migration methods are supported:

- **Advanced migration**
The migration is performed based on the data copy capability of Huawei SAN storage, which only applies to scenarios where Huawei SAN storage is used and efficient batch migration is required. Using this method, attached, unattached, and shared disks can be migrated. Before the migration, ensure that a separate storage link has been configured between the source backend storage where the disks reside and target backend storage.
- **General migration**
The migration is performed based on the data copy capability of compute hosts and applies to Huawei SAN storage, Huawei Distributed Block Storage, and third-party SAN storage. As data copy requires compute host resources, this method can be used only when a small number of disks attached to ECS

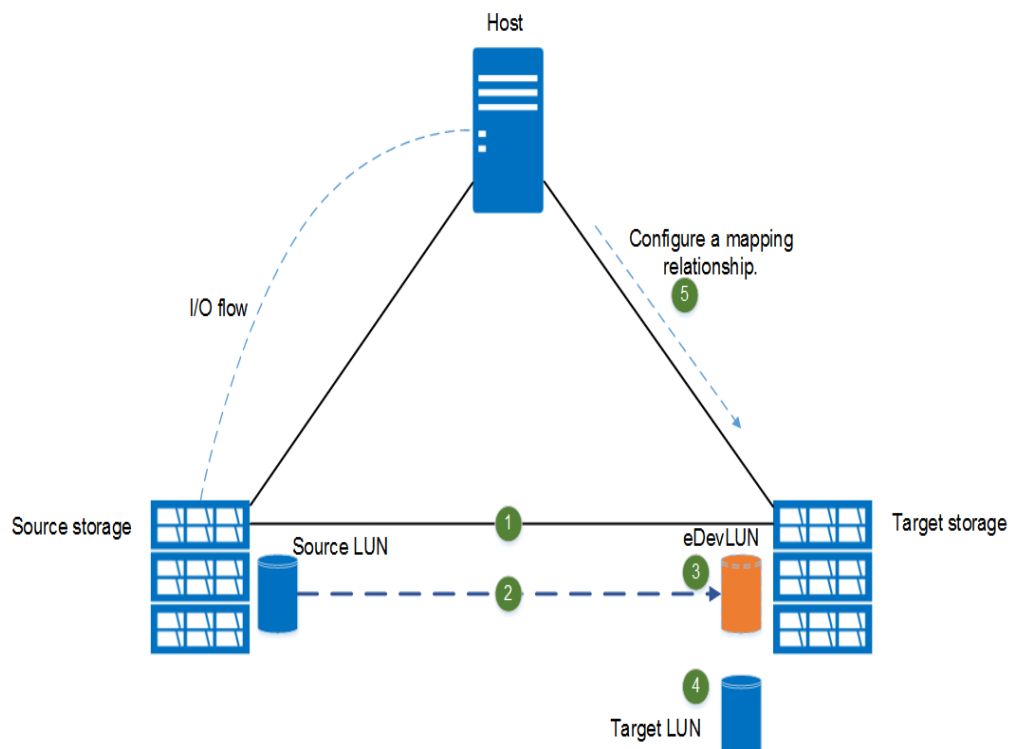
need to be migrated. It does not apply to storage device replacement. Shared disks cannot be migrated using this mode. SCSI disks can be migrated only when ECSs are shut down.

Advanced Migration Principles

Advanced migration is implemented based on SmartMigration of Huawei SAN storage.

Figure 3-25 to **Figure 3-28** show how to migrate the source LUN from the source storage to the target storage using SmartMigration.

Figure 3-25 Phase 1: Target storage taking over the source disk



1. Configure the link between the source storage and target storage. The link protocols between the source storage and target storage, between the host and source storage, and between the host and target storage are the same, which is Fibre Channel or iSCSI.

NOTE

The link between the source storage and target storage must be manually configured to ensure connectivity.

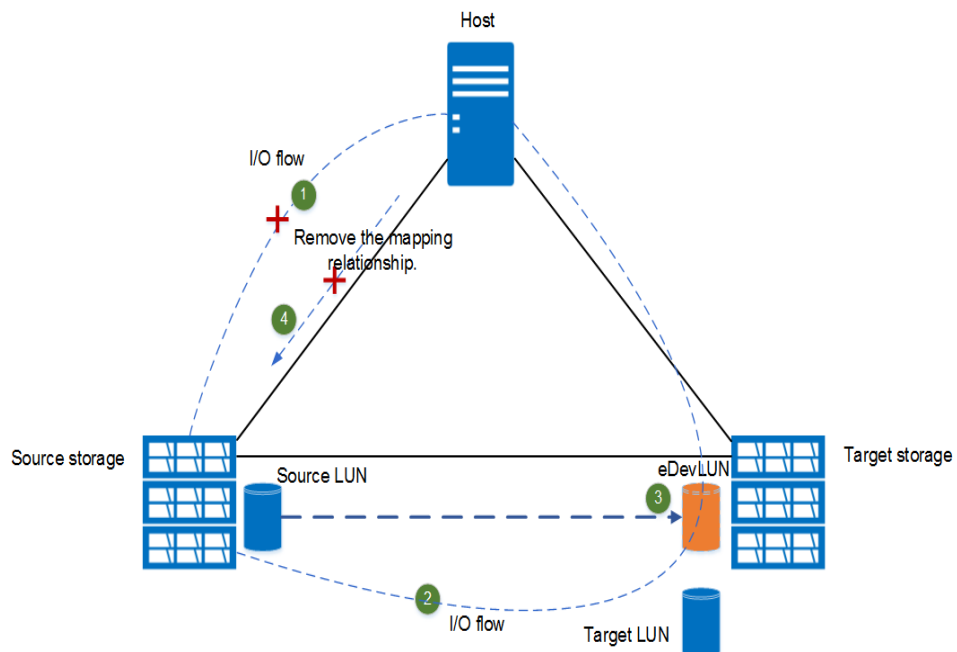
2. Map the source LUN of the source storage to the target storage.
3. Create an eDevLUN based on the mapped LUN on the target storage and configure the masquerading property for the eDevLUN. The host will identify the eDevLUN as the source LUN of the source storage.

The physical storage space required by the data volume of the eDevLUN is provided by the source storage. The target storage provides only the storage

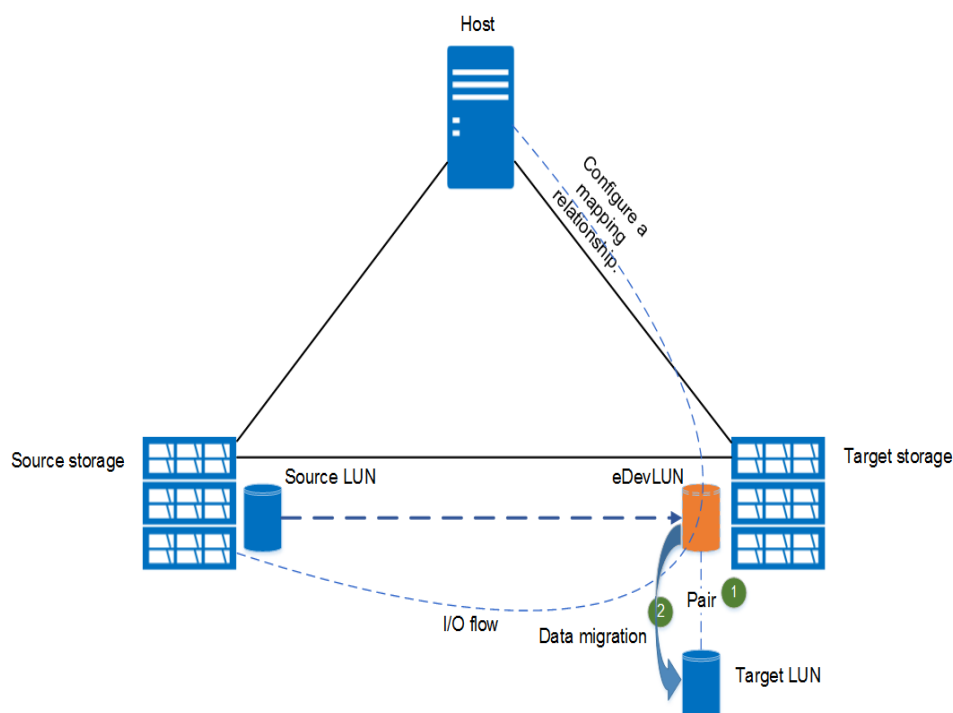
space of the metadata volume. Therefore, the storage space occupied by the target storage is small.

4. Create a target LUN on the target storage and migrate service data to the LUN.
5. Create the mapping between the eDevLUN and the host. The host scans the mapped eDevLUN.

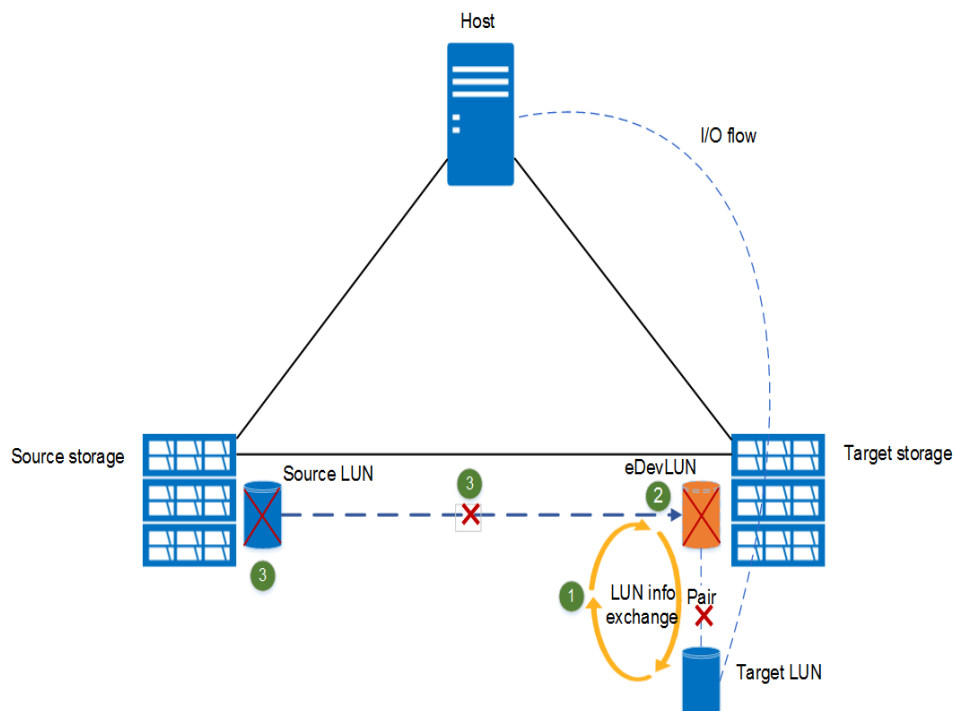
Figure 3-26 Phase 2: Configuring host path switchover



1. Configure I/O suspension.
2. Switch the I/O path of the host to the path between the host and the target storage so that the host can deliver I/Os to the source storage through the eDevLUN and cancel the I/O suspension.
3. Clear the path switchover attribute of the masquerading eDevLUN to ensure that the next migration can be performed properly.
4. Remove the mapping between the source LUN and the host.

Figure 3-27 Phase 3: Migrating service data

1. Create a SmartMigration task and create a pair relationship between the eDevLUN and target LUN.
In SmartMigration, a pair indicates the data migration relationship between the source LUN and target LUN.
2. Perform initial data synchronization and change data synchronization between the source LUN and target LUN.
The two synchronization modes are independent and can be performed at the same time to ensure that the service data change of the host can be completely synchronized to the source LUN and target LUN during data synchronization.

Figure 3-28 Phase 4: Removing the relationship

1. Stop service data synchronization between the eDevLUN and target LUN, exchange LUN information, and remove the data migration relationship after LUN information exchange. In this way, the target LUN can replace the eDevLUN to carry services.

In this process, host services are suspended. After information is exchanged, services are delivered to the target LUN after migration. In this way, services can be migrated without user experience.

2. Delete the eDevLUN.
3. Delete the mapping between the source LUN and target storage, and delete the source LUN to release the storage space.

3.2.10 EVS Disk Performance

When creating a disk type, you need to select backend storage for the disk type. The I/O performance of the disk type depends on the backend storage. When applying for an EVS disk, you can select a disk type based on application performance requirements.

EVS Disk Performance Indicators

EVS disk performance indicators include Input/Output Per Second (IOPS), throughput, and latency.

- **IOPS** indicates I/O requests the system can process per unit of time, usually, per second. I/O requests typically mean data read or write operation requests. The most commonly used IOPS indicators are sequential I/O and random I/O:
 - **Random I/O:** Access addresses are randomly distributed in the addressing space of disks. Services that generate random I/O include OLTP, SQL, and instant messaging services.

- Sequential I/O: Read/write operations access data continuously from adjacent addresses and are performed one by one based on logical blocks. During sequential I/O access, the disk seek time is greatly shortened because read and write heads can access the next block without moving. For example, most services such as data backup and log recording generate sequential I/O.
- **Throughput:** indicates the 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. The unit is MB/s.
- **Latency:** indicates the minimum interval between two consecutive read/write operations of an EVS disk. The unit is second.

Performance Test Methods

You can use different tools to test the EVS disk performance of instances using different operating systems:

- Linux instance: You can use tools such as DD, FIO, or Sysbench to test the disk performance. For details about the test method, see [How Can I Test the Performance of a Linux EVS Disk?](#).
- Windows instance: You can use tools such as FIO and Iometer to test the disk performance. For details about the test method, see [How Can I Test the Performance of a Windows EVS Disk?](#).

3.3 Operation Process

[Figure 3-29](#) shows the process of applying for and using an EVS disk.

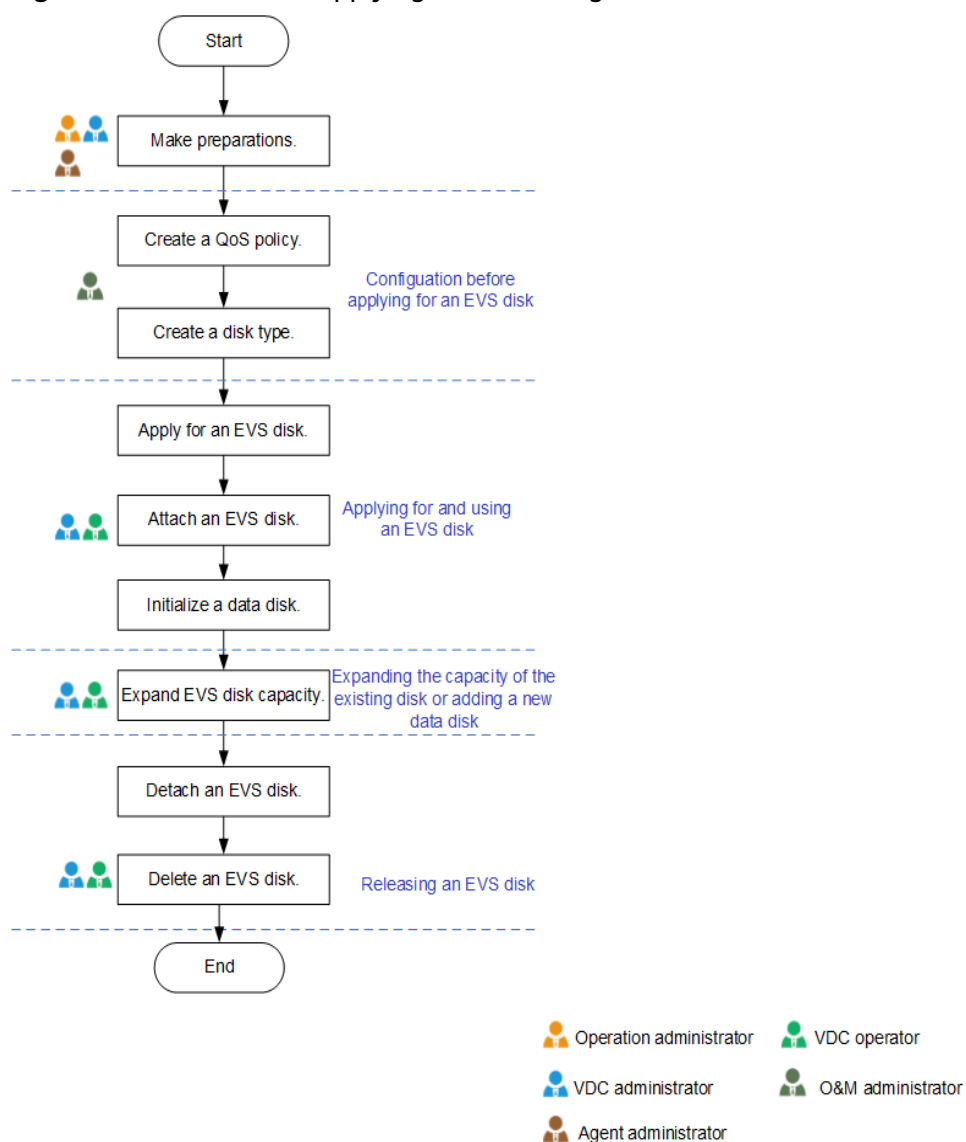
Figure 3-29 Process for applying for and using an EVS disk

Table 3-8 lists the operators and user interfaces (UIs) required for each step.

Table 3-8 Description of the EVS disk operation process

Operation	Description and Reference	Operator and UI Portal
Preparations	<p>Before using the EVS service, make the following preparations:</p> <ul style="list-style-type: none">• Obtain a VDC administrator account or a VDC operator account before creating an EVS disk. If no such account is available, contact an operation administrator for creating a VDC and VDC administrator, and then use the VDC administrator account to create a VDC operator. For details, see "VDC Tenant Modeling" in <i>HUAWEI CLOUD Stack 8.1.0 Resource Provisioning Guide</i>.• If you want to use a non-default service to apply for an EVS disk, you need to apply for the EVS service on ManageOne Operation Portal (ManageOne Operation Management Portal in B2B scenarios) and publish it. For details, see "Managing Services" in <i>HUAWEI CLOUD Stack 8.1.0 Resource Provisioning Guide</i>.	<ul style="list-style-type: none">• ManageOne Operation Portal (ManageOne Operation Management Portal in B2B scenarios): operation administrator• ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios): VDC administrator and agent administrator
Configuration before applying for an EVS disk	<ul style="list-style-type: none">• If you want to limit the maximum performance (IOPS or bandwidth) of burst services to prevent burst traffic from impacting other service traffic in the system, you can create a QoS policy and associate the QoS policy with a disk type. Then, the QoS policy can be applied to the disks provisioned using the disk type.• A disk type is created during installation and deployment using HUAWEI CLOUD Stack Deploy. If the disk type does not meet your requirements, contact the administrator to create a new one. "Storage Services" > "Elastic Volume Service (EVS for ECS)" > "Configuration Before Applying for an EVS Disk" > "(Optional) Creating a Disk Type" in <i>HUAWEI CLOUD Stack 8.1.0 Resource Provisioning Guide</i>	ManageOne Maintenance Portal: O&M administrator

Operation	Description and Reference	Operator and UI Portal
Applying for and using an EVS disk	You can apply for a blank EVS disk that does not contain data or apply for an EVS disk from a disk, snapshot, or backup and attach the EVS disk to an instance. Applying for a Data Disk Applying for a System Disk Applying for an EVS Disk Using an Existing Disk Applying for an EVS Disk Using a Snapshot Applying for an EVS Disk Using a Backup Applying for an EVS Disk Using an Image	ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios): VDC operator and VDC administrator
	An EVS disk can be used by an instance only after being attached to the instance. The EVS disk created together with an instance is automatically attached to the instance. Attaching an EVS Disk	ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios): VDC operator and VDC administrator
	After an EVS disk is attached to an ECS as a data disk of the ECS, you need to initialize the data disk. Initializing a Data Disk	ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios): VDC operator and VDC administrator
Expanding the capacity of the existing disk or adding a new data disk	Expand the capacity of the EVS disk whose available space is insufficient. Expanding EVS Disk Capacity	ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios): VDC operator and VDC administrator
Releasing an EVS disk	If the instance to which an EVS disk is attached does not need the EVS disk anymore, you can detach the EVS disk. Detaching an EVS Disk	ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios): VDC operator and VDC administrator

Operation	Description and Reference	Operator and UI Portal
	After an EVS disk is detached, delete the EVS disk to release storage space. Deleting an EVS Disk	ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios): VDC operator and VDC administrator

3.4 Applying for an EVS Disk

3.4.1 Applying for a Data Disk

A data disk can be created together with an instance. During the instance creation, you can configure parameters such as the disk type, capacity, and sharing attribute. After the instance is created, the disk will be automatically attached to the instance. This section describes how to apply for and add a blank data disk to an instance separately.

Context

ECSs do not support the merging of EVS disk spaces. Each EVS disk is independent, and the spaces of multiple EVS disks cannot be merged through formatting. You are advised to plan the number and capacity of EVS disks before disk creation. It is not recommended that logical volumes managed by the LVM be created on the disks, because snapshots are created for independent EVS disks and creating such logical volumes will generate differential data after a snapshot rollback.

Procedure

- Step 1** Log in to the EVS console as a VDC administrator or VDC operator. For details, see [Logging In to the EVS Console as a VDC Administrator or VDC Operator](#).
- Step 2** Click **Apply for EVS Disk**.
- Step 3** In the **Select Service** dialog box, select the target service and then click **Apply Now**.

NOTE

Services are created by administrators based on operation requirements. When creating a service, the administrator can lock service parameters (for example, apply for a fixed disk type or apply for resources in a specified AZ), specify the service publishing scope (for example, visible to only specific VDCs), and set the product approval process. For details about how to create more services, see "Creating a Common Service" in *HUAWEI CLOUD Stack 8.1.0 Resource Provisioning Guide*.

- Step 4** In the **Apply for EVS Disk** dialog box, set the parameters as prompted and based on [Table 3-9](#).

← Apply for EVS Disk

1 Specify Details

2 Confirm Specifications

3 Finish

AZ

Do not specify

Create from snapshot

Create from disk

Create from image

Data disk

System disk

Disk Type

If you need to create a disk type, contact the administrator to create it on Service OM. If -- is displayed after a feature, it indicates that the feature is not configured when the disk type is created. If you need to configure this feature, contact the administrator.

Disk Type	Configuration Mode	SmartTier	Deduplication	IOPS Upper Limit	Bandwidth Upper Limit	I/O Priority
<input checked="" type="radio"/> 001	Thin	Initial Allocation Policy Relocation policy	Deduplication Compression	IOPS Upper Limit:GB Max IOPS Upper Limit Min IOPS Upper Limit	Bandwidth Upper Limit Max Bandwidth Upper Limit Min Bandwidth Upper Limit	
<input type="radio"/> SSD	Thin	Initial Allocation Policy Relocation policy	Deduplication Compression	IOPS Upper Limit:GB Max IOPS Upper Limit Min IOPS Upper Limit	Bandwidth Upper Limit Max Bandwidth Upper Limit Min Bandwidth Upper Limit	
<input type="radio"/> ssd	Thin	Initial Allocation Policy Relocation policy	Deduplication Compression	IOPS Upper Limit:GB Max IOPS Upper Limit Min IOPS Upper Limit	Bandwidth Upper Limit Max Bandwidth Upper Limit Min Bandwidth Upper Limit	
<input type="radio"/> SSD3	Thin	Initial Allocation Policy Relocation policy	Deduplication Compression	IOPS Upper Limit:GB Max IOPS Upper Limit Min IOPS Upper Limit	Bandwidth Upper Limit Max Bandwidth Upper Limit Min Bandwidth Upper Limit	
<input type="radio"/> Type Qcs	Thin	Initial Allocation Policy Relocation policy	Deduplication Compression	IOPS Upper Limit:GB Max IOPS Upper Limit Min IOPS Upper Limit	Bandwidth Upper Limit Max Bandwidth Upper Limit Min Bandwidth Upper Limit	

5

Total Records: 6

1

2

Capacity (GB)

10

The remaining capacity is 100 GB. Contact the administrator if you want to apply for more quota.

IOPS Upper Limit of the Disk

Bandwidth Upper Limit of the Disk

(MB/s)

Device Type

Virt

SCSI

Disk Sharing

Disable

Enable

Disk Name

Quantity

1

You can apply for 100 more disks. Only 100 disk can be applied for at one time. Contact the administrator if you want to apply for more quota.

Required Duration

Unlimited

1 year

Custom

Description

NOTE

If the service you selected has parameters **AZ**, **Capacity (GB)**, **Disk Sharing**, or **Disk Type** configured, the parameter values configured for the service will be displayed for the EVS disk you apply for.

Table 3-9 Parameters required when applying for an EVS disk

Parameter	Description	Example Value
AZ	<p>Specifies the availability zone (AZ) where an EVS disk is to be created.</p> <p>NOTE</p> <ul style="list-style-type: none">EVS disks can be attached to instances only in the same AZ.If an AZ is bound to a tenant upon the tenant creation, the resource pools in this AZ are used as dedicated storage pools for this tenant and invisible to other tenants not bound to the AZ. Services of this tenant can run on dedicated physical devices without interference from other tenants.	az1.dc1
Data Source	Select Do not specify to apply for a blank EVS disk that does not contain data.	Do not specify
Disk	<p>Specifies whether to apply for a data disk or system disk.</p> <p>This section describes how to apply for a data disk.</p>	Data disk

Issue 01 (2021-09-30)

Copyright © Huawei Technologies Co., Ltd.

73

Parameter	Description	Example Value
Disk Type	Select a disk type. You can select the disk type created in "Storage Services" > "Elastic Volume Service (EVS for ECS)" > "Configuration Before Applying for an EVS Disk" > "(Optional) Creating a Disk Type" in <i>HUAWEI CLOUD Stack 8.1.0 Resource Provisioning Guide</i> . You can also select the disk type created during automatic installation and deployment using HUAWEI CLOUD Stack Deploy. The name of the automatically created disk type is the value of business_volume_type in deployment parameter summary file <i>xxx_export_all_v2_EN.xlsx</i> .	-
Capacity (GB)	EVS disk capacity. The EVS disk capacity can neither exceed the total capacity quota of EVS disks nor the capacity quota of the current disk type.	10 GB
Device Type	<ul style="list-style-type: none">• VBD indicates a VBD EVS disk.• SCSI indicates a SCSI EVS disk. SCSI-type disks allow the ECS OS to directly access the underlying storage media and support SCSI commands more advanced than VBD-type disks. For details about how to use SCSI disks, see Device Type. NOTE Device Type can be set only to VBD for the selected AZ using a heterogeneous storage device.	VBD
Disk Sharing	<ul style="list-style-type: none">• Disable indicates that the new disk will be a non-shared EVS disk.• Enable indicates that the new disk will be a shared EVS disk. Such a disk can be attached to multiple ECSs. NOTE If the backend storage device is a heterogeneous storage device, no shared EVS disk can be created.	Enable

Parameter	Description	Example Value
Disk Name	<p>The disk name can contain only letters, digits, underscores (_), and hyphens (-). When applying for a single EVS disk, ensure that the disk name contains less than or equal to 63 characters. When applying for EVS disks in batches, ensure that the disk name contains less than or equal to 58 characters.</p> <ul style="list-style-type: none">• If you apply for a single EVS disk, the value of this parameter is used as the name of the EVS disk.• If you apply for multiple EVS disks in batches, the value of this parameter is used as the prefix of the names of the EVS disks. The name of each EVS disk resembles <i>Disk name-A four-digit number</i>. <p>NOTE For example, if you apply for two EVS disks and set Disk Name to volume, the names of the two EVS disks will be volume-0001 and volume-0002.</p>	volume-0001
Quantity	<p>Specifies the number of EVS disks that you apply for. The default value is 1, which means that you apply for one EVS disk.</p> <p>By default, a maximum of 100 EVS disks can be created at a time. The number of EVS disks that can be applied for in a batch varies with the current EVS disk quantity quota.</p>	1
Required Duration	<p>Specifies the validity period of the EVS disk that you apply for.</p> <ul style="list-style-type: none">• If you select Unlimited, the new EVS disk has no expiration date.• If you select 1 year, the validity period of the new EVS disk will be one year, which is subject to the expiration date displayed on the console.• If you select Custom, specify an expiration date for the EVS disk that you apply for.	Unlimited
Description	<p>Describes the EVS disk that you apply for.</p> <p>The length cannot exceed 63 characters.</p>	-

Step 5 Click **Next**.

Step 6 Confirm that the information is correct and click **Add to Cart** or **Apply Now**.

If the configuration is incorrect, click **Back**.

- If you click **Add to Cart**, go to [Step 7](#).
- If you click **Apply Now**, go to [Step 8](#).

Step 7 Submit an application for the service in the shopping cart.

1. Click the shopping cart in the upper right corner of the page.
2. Select the service and click **Apply Now**.
3. Enter the order information and click **OK** to submit the application.

Step 8 If the application for an EVS disk requires approval, contact the administrator for approval. Otherwise, skip this step.

 **NOTE**

The application for an EVS disk can be rejected by the administrator. If you have entered incorrect configuration when applying for an EVS disk, you can contact the administrator to reject the application, correct the configuration, and submit the application again.

Step 9 On the **Elastic Volume Service** page, view the status of the EVS disk. After the EVS disk is created and its status changes to **Available**, the EVS disk is successfully created.

----End

Follow-up Procedure

If you want to use the new EVS disk, [attach it to an ECS](#) first.

3.4.2 Applying for a System Disk

If the system disk is unavailable because your instance is deleted by mistake or any other fault occurs, perform the following: You can apply for a new blank system disk and restore backup data of the unavailable system disk on the new disk. Alternatively, you can apply for a system disk from the existing system disk, system disk snapshot, system disk backup, or image, and use the new system disk to start an ECS instance. This section describes how to apply for a new blank system disk and apply for a system disk from a system disk image. To apply for a system disk from the existing system disk, snapshot, or backup, see [Applying for an EVS Disk Using an Existing Disk](#), [Applying for an EVS Disk Using a Snapshot](#), or [Applying for an EVS Disk Using a Backup](#).

Procedure

- Step 1** Log in to the EVS console as a VDC administrator or VDC operator. For details, see [Logging In to the EVS Console as a VDC Administrator or VDC Operator](#).
- Step 2** Click **Apply for EVS Disk**.
- Step 3** In the **Select Service** dialog box, select the target service and then click **Apply Now**.

 NOTE

Services are created by administrators based on operation requirements. When creating a service, the administrator can lock service parameters (for example, apply for a fixed disk type or apply for resources in a specified AZ), specify the service publishing scope (for example, visible to only specific VDCs), and set the product approval process. For details about how to create more services, see "Creating a Common Service" in *HUAWEI CLOUD Stack 8.1.0 Resource Provisioning Guide*.

Step 4 In the **Apply for EVS Disk** dialog box, set the parameters as prompted, as shown in [Table 3-10](#).

< Apply for EVS Disk

1 Specify Details 2 Confirm Specifications 3 Finish

AZ

Data Source

Disk

This system disk cannot be used to start an ECS directly. If you want to do so, copy the backup data of a normal system disk to it.

Image

Other

Disk Type

If you need to create a disk type, contact the administrator to create it on Service OM. If -- is displayed after a feature, it indicates that the feature is not configured when the disk type is created. If you need to configure this feature, contact the administrator.

Disk Type	Configuration Mode	SmartTier	Deduplication	IOPS Upper Limit	Bandwidth Upper Limit	I/O Priority
<input checked="" type="radio"/> 001	Thin	Initial Allocation Policy Relocation policy	Deduplication Compression	IOPS Upper Limit/CB Max IOPS Upper Limit Min IOPS Upper Limit	Bandwidth Upper Limit/CB Max Bandwidth Upper Li Min Bandwidth Upper Li	
<input type="radio"/> 134	--	Initial Allocation Policy Relocation policy	Deduplication Compression	IOPS Upper Limit/CB Max IOPS Upper Limit Min IOPS Upper Limit	Bandwidth Upper Limit/CB Max Bandwidth Upper Li Min Bandwidth Upper Li	
<input type="radio"/> SSD	Thin	Initial Allocation Policy Relocation policy	Deduplication Compression	IOPS Upper Limit/CB Max IOPS Upper Limit Min IOPS Upper Limit	Bandwidth Upper Limit/CB Max Bandwidth Upper Li Min Bandwidth Upper Li	
<input type="radio"/> wal	Thin	Initial Allocation Policy Relocation policy	Deduplication Compression	IOPS Upper Limit/CB Max IOPS Upper Limit Min IOPS Upper Limit	Bandwidth Upper Limit/CB Max Bandwidth Upper Li Min Bandwidth Upper Li	
<input type="radio"/> SSD3	Thin	Initial Allocation Policy Relocation policy	Deduplication Compression	IOPS Upper Limit/CB Max IOPS Upper Limit Min IOPS Upper Limit	Bandwidth Upper Limit/CB Max Bandwidth Upper Li Min Bandwidth Upper Li	

5 Total Records: 7 1 2

Capacity (GB) The remaining capacity is 90 GB. Contact the administrator if you want to apply for more quota.

IOPS Upper Limit of the Disk --

Bandwidth Upper Limit of the Disk (MB/s) --

Device Type

Disk Sharing

Disk Name

Quantity You can apply for 99 more disks. Only 100 disk can be applied for at one time. Contact the administrator if you want to apply for more quota.

Required Duration

Description

 NOTE

If the service you selected has parameters **AZ**, **Capacity (GB)**, **Disk Sharing**, or **Disk Type** configured, the parameter values configured for the service will be displayed for the EVS disk you apply for.

Table 3-10 Parameters required when applying for an EVS disk

Parameter		Description	Example Value
AZ		<p>Specifies the availability zone (AZ) where an EVS disk is to be created.</p> <p>NOTE</p> <ul style="list-style-type: none">• EVS disks can be attached to instances only in the same AZ.• If the selected AZ uses a heterogeneous storage device, no system disk can be created.• If an AZ is bound to a tenant upon the tenant creation, the resource pools in this AZ are used as dedicated storage pools for this tenant and invisible to other tenants not bound to the AZ. Services of this tenant can run on dedicated physical devices without interference from other tenants.	az1.dc1
Data Source	Do not specify	<p>Apply for a blank system disk.</p> <p>The system disk created using this method cannot be used to start the ECS. If you need to use this function, copy the backup data of the normal system disk to the new system disk.</p>	Do not specify
	Create from image	<p>Create a system disk from an image.</p> <p>Click Select. On the Select Image page, select a system disk image to create a system disk.</p> <p>The created system disk can only be attached to an ECS with the same image.</p> <p>NOTE</p> <p>In VMware and VRM scenarios, VIMS data storage do not support creating a system disk from image.</p>	
Disk		Apply for a system disk.	System disk
Image		<p>Specifies the image used to create the system disk.</p> <p>This parameter is mandatory when Data Source is Do not specify and Disk is System disk. Options include Public Image, Private Image, and Shared Image. For details about the three types of images and how to create them, see Operation Help Center > Image Management Service.</p>	-

Parameter	Description	Example Value
Disk Type	Select a disk type. You can select the disk type created in "Storage Services" > "Elastic Volume Service (EVS for ECS)" > "Configuration Before Applying for an EVS Disk" > "(Optional) Creating a Disk Type" in <i>HUAWEI CLOUD Stack 8.1.0 Resource Provisioning Guide</i> . You can also select the disk type created during automatic installation and deployment using HUAWEI CLOUD Stack Deploy. The name of the automatically created disk type is the value of business_volume_type in deployment parameter summary file <i>xxx_export_all_v2_EN.xlsx</i> .	business
Capacity (GB)	EVS disk capacity. The EVS disk capacity can neither exceed the total capacity quota of EVS disks nor the capacity quota of the current disk type.	10 GB
Device Type	Only system disks of the VBD type can be created.	VBD
Disk Sharing	Only system disks that are not shared can be created.	Disable
Disk Name	The disk name can contain only letters, digits, underscores (_), and hyphens (-). When applying for a single EVS disk, ensure that the disk name contains less than or equal to 63 characters. When applying for EVS disks in batches, ensure that the disk name contains less than or equal to 58 characters. <ul style="list-style-type: none">If you apply for a single EVS disk, the value of this parameter is used as the name of the EVS disk.If you apply for multiple EVS disks in batches, the value of this parameter is used as the prefix of the names of the EVS disks. The name of each EVS disk resembles <i>Disk name-A four-digit number</i>. NOTE For example, if you apply for two EVS disks and set Disk Name to volume , the names of the two EVS disks will be volume-0001 and volume-0002 .	volume-0001

Parameter	Description	Example Value
Quantity	<p>Specifies the number of EVS disks that you apply for. The default value is 1, which means that you apply for one EVS disk.</p> <p>By default, a maximum of 100 EVS disks can be created at a time. The number of EVS disks that can be applied for in a batch varies with the current EVS disk quantity quota.</p> <p>NOTE You cannot batch create system disks from an image. Quantity must be 1.</p>	1
Required Duration	<p>Specifies the validity period of the EVS disk that you apply for.</p> <ul style="list-style-type: none">• If you select Unlimited, the new EVS disk has no expiration date.• If you select 1 year, the validity period of the new EVS disk will be one year, which is subject to the expiration date displayed on the console.• If you select Custom, specify an expiration date for the EVS disk that you apply for.	30 days
Description	<p>Describes the EVS disk that you apply for. The length cannot exceed 63 characters.</p>	-

Step 5 Click **Next**.

Step 6 Confirm that the information is correct and click **Add to Cart** or **Apply Now**.

If the configuration is incorrect, click **Back**.

- If you click **Add to Cart**, go to [Step 7](#).
- If you click **Apply Now**, go to [Step 8](#).

Step 7 Submit an application for the service in the shopping cart.

1. Click the shopping cart in the upper right corner of the page.
2. Select the service and click **Apply Now**.
3. Enter the order information and click **OK** to submit the application.

Step 8 If the application for an EVS disk requires approval, contact the administrator for approval. Otherwise, skip this step.

 **NOTE**

The application for an EVS disk can be rejected by the administrator. If you have entered incorrect configuration when applying for an EVS disk, you can contact the administrator to reject the application, correct the configuration, and submit the application again.

- Step 9** On the **Elastic Volume Service** page, view the status of the EVS disk. After the EVS disk is created and its status changes to **Available**, the EVS disk is successfully created.

----End

Follow-up Procedure

After applying for a system disk, perform the following operations so that the system disk can be used by an ECS instance:

- System disk applied when **Data Source** is **Do not specify**: You cannot directly start an ECS using the system disk. You need to restore backup data of the unavailable system disk on the new system disk and attach the new system disk to the ECS by referring to [attach the new system disk to the ECS](#). Before attaching the new system disk, detach the existing system disk from the ECS.
- System disk applied when **Data Source** is **Create from image**: You need to attach the new system disk to the ECS by referring to [attach the new system disk to the ECS](#). Before attaching the new system disk, detach the existing system disk from the ECS.

3.4.3 Applying for an EVS Disk Using an Existing Disk

You can use a specific disk to apply for an EVS disk that contains the same data quickly. If the source disk is a data disk, you can apply for a data disk that contains the same data. If the source disk is a system disk, you can apply for a system disk that contains the same data. The applied EVS disk can be attached to an ECS instance. For details about how to apply for a system disk using a disk, see "Storage Services" > "Elastic Volume Service (EVS for ECS)" > "Best Practice of Restoring ECSs" > "Restoring ECS Services By Creating a Disk Using an Existing Disk" in *HUAWEI CLOUD Stack 8.1.0 Resource Provisioning Guide*.

Restrictions

- When you apply for an EVS disk using an existing disk, the source disk must be in **Available** or **In-use** state. If VRM or VMware products are used, the source disk attached to the instance cannot be selected.
- When OceanStor V3/V5 is used, an EVS disk can be applied for from an existing EVS disk only after the administrator imports the HyperCopy license onto the storage device.
- When you use a specific system disk to apply for a new system disk, the image of the new system disk is the same as that of the source system disk, and the new system disk can be attached only to the ECS with the same image.

Procedure

- Step 1** Log in to the EVS console as a VDC administrator or VDC operator. For details, see [Logging In to the EVS Console as a VDC Administrator or VDC Operator](#).
- Step 2** In the EVS disk list, locate the disk used for applying for a new EVS disk.

 **NOTE**

You can also click **Apply for EVS Disk** and choose **Create from disk** to apply for an EVS disk. The restrictions on the source disk are as follows:

- Only EVS disks in the current AZ can be selected as source disks.
- If the disk capacity and type have been specified in the selected service, the capacity of the source disk must be less than or equal to that of the specified one, and the disk type must be the same as the specified one. Disks not meeting this requirement cannot be used as source disks.

Step 3 In the row where the source disk resides, choose **More > Apply for Disk**.

Step 4 In the **Select Service** dialog box, select the target service and then click **Apply Now**.

 **NOTE**

- If the disk capacity has been specified in the service, you can only select the service whose disk capacity is greater than or equal to the source disk capacity.
- If the disk type has been specified in the service, you can only select the service whose disk type is the same as the source disk type.
- If the disk capacity and disk type have been specified in the service, you can only select the service whose disk capacity is greater than or equal to the source disk capacity and whose disk type is the same as the source disk type.

Step 5 In the **Apply for EVS Disk** dialog box, configure EVS disk information as prompted.

Some parameters must be consistent with those of the source disk and cannot be modified on the GUI. Configure other parameters by referring to [Table 3-11](#). Parameters not listed in [Table 3-11](#) cannot be modified.

Table 3-11 Parameters for applying for an EVS disk using a disk

Parameter	Description	Example Value
Capacity (GB)	<p>EVS disk capacity.</p> <p>If you do not specify the capacity when applying for an EVS disk using a disk, the capacity of the new disk is the same as that of the source disk.</p> <p>The EVS disk capacity can neither exceed the total capacity quota of EVS disks nor the capacity quota of the current disk type.</p>	10 GB

Parameter	Description	Example Value
Disk Sharing	<ul style="list-style-type: none">• Disable indicates that the new disk will be a non-shared EVS disk.• Enable indicates that the new disk will be a shared EVS disk. Such a disk can be attached to multiple ECSs. NOTE <ul style="list-style-type: none">- If the backend storage device is a heterogeneous storage device, no shared EVS disk can be created.- If the source disk is a system disk, a shared disk cannot be created.	Enable
Disk Name	The disk name can contain only letters, digits, underscores (_), and hyphens (-). When applying for a single disk, ensure that the disk name contains not more than 63 characters.	volume-0001
Quantity	Number of applied EVS disks. You cannot apply for EVS disks in batches using a disk. You can apply for only one EVS disk at a time.	1
Required Duration	Specifies the validity period of the EVS disk that you apply for. <ul style="list-style-type: none">• If you select Unlimited, the new EVS disk has no expiration date.• If you select 1 year, the validity period of the new EVS disk will be one year, which is subject to the expiration date displayed on the console.• If you select Custom, specify an expiration date for the EVS disk that you apply for.	Unlimited
Description	Describes the EVS disk that you apply for. The length cannot exceed 63 characters.	-

Step 6 Click **Next**.

Step 7 Confirm the information and click **Add to Cart** or **Apply Now**.

If the configuration is incorrect, click **Back**.

- If you click **Add to Cart**, go to [Step 8](#).
- If you click **Apply Now**, go to [Step 9](#).

Step 8 Submit an application for the service in the shopping cart.

1. Click the shopping cart in the upper right corner of the page.

2. Select the service and click **Apply Now**.
3. Enter the order information and click **OK** to submit the application.

Step 9 If the application for an EVS disk requires approval, contact the administrator for approval. Otherwise, skip this step.

 **NOTE**

The application for an EVS disk can be rejected by the administrator. If you have entered incorrect configuration when applying for an EVS disk, you can contact the administrator to reject the application, correct the configuration, and submit the application again.

Step 10 On the **Elastic Volume Service** page, view the status of the EVS disk. After the EVS disk is created and its status changes to **Available**, the EVS disk is successfully created.

 **NOTE**

The time required for applying for an EVS disk using an existing disk depends on the capacity of the existing disk. A larger capacity indicates a longer time. The default speed is 10 MB/s to 20 MB/s.

----End

Follow-up Procedure

If you want to use the new EVS disk, [attach it to an ECS](#) first.

3.4.4 Applying for an EVS Disk Using a Snapshot

You can use a snapshot to create multiple disks (data disk or system disk) containing the same initial data, and these disks can be used as data resources for various services, such as data mining, report query, and development and test. This method protects the initial data and creates EVS disks rapidly, meeting the diversified service data requirements.

Context

You can create EVS disks using a snapshot. The new EVS disks containing the data in the snapshot are copies of the original EVS disk. You can but are advised not to use the snapshot of a small-capacity data disk (< 2 TB) to create large-capacity data disks (> 2 TB).

Instead, you can create large-capacity blank data disks or use the snapshot of a large-capacity data disk to create other large-capacity data disks. The reasons are as follows:

- When a large-capacity data disk is created using a snapshot of a small-capacity data disk, the disk capacity is only expanded at the block device level, but the partition format and file system are not automatically converted.
- If the snapshot of a small-capacity data disk uses the MBR partition format, the partitioning tool (**parted** on the Linux OS or the disk management module on the Windows OS) cannot convert the partition format from MBR to GPT while retaining data. Even if you use a snapshot of a small-capacity data disk to create a large-capacity data disk, you need to delete the original

data and then partition the disk in the GPT format during partition formatting.

- If a large-capacity data disk has been created using the snapshot of a small-capacity data disk, the partition format of the created disk is MBR. Before formatting, you need to delete the original disk, convert it into a GPT partition, and then partition and format the data again. That is, the data in the original snapshot will be deleted.

Restrictions

When applying for an EVS disk using a snapshot, if the backend storage type is OceanStor V3/V5 series, the administrator needs to import the license of the HyperCopy feature on the device in advance.

Procedure

Step 1 Log in to the EVS console as a VDC administrator or VDC operator. For details, see [Logging In to the EVS Console as a VDC Administrator or VDC Operator](#).

Step 2 In the navigation pane on the left, choose **Elastic Volume Service > Snapshot**.

NOTE

You can also click **Apply for EVS Disk** and choose **Create from snapshot** to apply for an EVS disk. In this case, only snapshots in the current AZ can be used to apply for EVS disks.

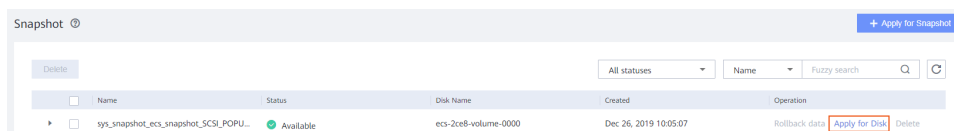
Step 3 In the snapshot list, locate the row that contains the snapshot, and click **Apply for EVS Disk** in the **Operation** column.

Step 4 In the **Select Service** dialog box, select the target service and then click **Apply Now**.

NOTE

- If the disk capacity has been specified in the service, you can only select the service whose disk capacity is greater than or equal to the snapshot capacity.
- If the disk type has been specified in the service, you can only select the service whose disk type is the same as the source disk type of the snapshot.
- If the disk capacity and disk type have been specified in the service, you can only select the service whose disk capacity is greater than or equal to the snapshot capacity and whose disk type is the same as the source disk type of the snapshot.

Step 5 The **Apply for Disk** page is displayed. Parameters are described in [Table 3-12](#).



NOTE

Some parameters must be consistent with those of the source disk of the snapshot and cannot be modified on the GUI. Configure other parameters by referring to [Table 3-12](#). Parameters not listed in [Table 3-12](#) cannot be modified.

Table 3-12 Parameters for applying for an EVS disk using a snapshot

Parameter	Description	Example Value
Capacity (GB)	<p>EVS disk capacity.</p> <ul style="list-style-type: none">In VRM scenarios, if an EVS disk is created from a snapshot, the EVS disk capacity cannot be set as it must be the same as the source disk capacity. In other scenarios, the EVS disk capacity must be greater than or equal to the source disk capacity.When the disk capacity is not specified, the default capacity is the same as that of the source disk of the snapshot.The EVS disk capacity can neither exceed the total capacity quota of EVS disks nor the capacity quota of the current disk type.	10 GB
Disk Sharing	<ul style="list-style-type: none">Disable indicates that the new disk will be a non-shared EVS disk.Enable indicates the new disk will be a shared EVS disk. Such a disk can be attached to multiple instances. <p>NOTE If the source disk is a system disk, a shared disk cannot be created.</p>	Disable
Disk Name	The disk name can contain only letters, digits, underscores (_), and hyphens (-). When applying for a single disk, ensure that the disk name contains not more than 63 characters.	volume_001
Quantity	<p>Number of applied EVS disks.</p> <p>You cannot create EVS disks in batches using a snapshot. You can apply for only one EVS disk at a time.</p>	N/A
Required Duration	<p>Specifies the validity period of the EVS disk that you apply for.</p> <ul style="list-style-type: none">If you select Unlimited, the new EVS disk has no expiration date.If you select 1 year, the validity period of the new EVS disk will be one year, which is subject to the expiration date displayed on the console.If you select Custom, specify an expiration date for the EVS disk that you apply for.	Unlimited
Description	<p>Describes the EVS disk that you apply for.</p> <p>The length cannot exceed 63 characters.</p>	N/A

Step 6 Click **Next**.

Step 7 Confirm that the information is correct and click **Add to Cart** or **Apply Now**.

If the configuration is incorrect, click **Back**.

- If you click **Add to Cart**, go to [Step 8](#).
- If you click **Apply Now**, go to [Step 9](#).

Step 8 Submit an application for the service in the shopping cart.

1. Click the shopping cart in the upper right corner of the page.
2. Select the service and click **Apply Now**.
3. Enter the order information and click **OK** to submit the application.

Step 9 If the application for an EVS disk requires approval, contact the administrator for approval. Otherwise, skip this step.

 **NOTE**

The application for an EVS disk can be rejected by the administrator. If you have entered incorrect configuration when applying for an EVS disk, you can contact the administrator to reject the application, correct the configuration, and submit the application again.

Step 10 On the **Elastic Volume Service** page, view the status of the EVS disk. After the EVS disk is created and its status changes to **Available**, the EVS disk is successfully created.

 **NOTE**

The time required for applying for an EVS disk using an existing disk depends on the capacity of the existing disk. A larger capacity indicates a longer time. The default speed is 10 MB/s to 20 MB/s.

----End

Follow-up Procedure

If you want to use the new EVS disk, attach it to an ECS first. For details, see [attach it to an ECS](#).

3.4.5 Applying for an EVS Disk Using a Backup

You can create an EVS disk using the backup of an EVS disk (data disk or system disk). Data on the new disk is the same as that in the backup.

You can use a backup to create an EVS disk in either of the following ways:

- EVS console: To use the specified backup to create an EVS disk, perform the related operations on the EVS console.
- VBS console: To view the detailed information about a backup and then use it to create an EVS disk, perform the related operations on the VBS console.

EVS Console

Step 1 Log in to the EVS console as a VDC administrator or VDC operator. For details, see [Logging In to the EVS Console as a VDC Administrator or VDC Operator](#).

Step 2 Click **Apply for EVS Disk**.

Step 3 In the **Select Service** dialog box, select the target service and then click **Apply Now**.

 **NOTE**

Services are created by administrators based on operation requirements. When creating a service, the administrator can lock service parameters (for example, apply for a fixed disk type or apply for resources in a specified AZ), specify the service publishing scope (for example, visible to only specific VDCs), and set the product approval process. For details about how to create more services, see "Creating a Common Service" in *HUAWEI CLOUD Stack 8.1.0 Resource Provisioning Guide*.

Step 4 In the **Apply for EVS Disk** dialog box, set the parameters as prompted and based on [Table 3-13](#).

Table 3-13 Parameters for applying for an EVS disk using a backup

Parameter	Description	Example Value
AZ	Specifies the availability zone (AZ) where an EVS disk is to be created. NOTE <ul style="list-style-type: none">EVS disks can be attached to instances only in the same AZ.If an AZ is bound to a tenant upon the tenant creation, the resource pools in this AZ are used as dedicated storage pools for this tenant and invisible to other tenants not bound to the AZ. Services of this tenant can run on dedicated physical devices without interference from other tenants.	az1.dc1
Data Source	Select Create from backup , click Select , and select a backup as the data source of the new EVS disk. NOTE <ul style="list-style-type: none">If the VBS service is not configured or the VBS service is configured but disabled, this parameter is not displayed. To use the VBS, configure and enable it by referring to "Storage Services" > "Elastic Volume Service (EVS for ECS)" > "FAQs" > "How Do I Enable the VBS?" in <i>HUAWEI CLOUD Stack 8.1.0 Resource Provisioning Guide</i>.If the backend storage is heterogeneous storage, you cannot select Create from backup.	Create from backup
Disk	Apply for a data disk or system disk that has the same attributes as the source disk.	Data disk

Parameter	Description	Example Value
Disk Type	Select a disk type. You can select the disk type created in "Storage Services" > "Elastic Volume Service (EVS for ECS)" > "Configuration Before Applying for an EVS Disk" > "(Optional) Creating a Disk Type" in <i>HUAWEI CLOUD Stack 8.1.0 Resource Provisioning Guide</i> . You can also select the disk type created during automatic installation and deployment using HUAWEI CLOUD Stack Deploy. The name of the automatically created disk type is the value of business_volume_type in deployment parameter summary file <i>xxx_export_all_v2_EN.xlsx</i> .	-
Capacity (GB)	EVS disk capacity. When you apply for an EVS disk using a backup, the capacity of the new EVS disk must be greater than or equal to that of the source disk of the backup. The EVS disk capacity can neither exceed the total capacity quota of EVS disks nor the capacity quota of the current disk type.	10 GB
Device Type	<ul style="list-style-type: none">• VBD indicates a VBD EVS disk.• SCSI indicates a SCSI EVS disk. SCSI-type disks allow the ECS OS to directly access the underlying storage media and support SCSI commands more advanced than VBD-type disks. For details about how to use SCSI disks, see Device Type. NOTE Device Type can be set only to VBD for the selected AZ using a heterogeneous storage device.	VBD
Disk Sharing	<ul style="list-style-type: none">• Disable indicates that the new disk will be a non-shared EVS disk.• Enable indicates that the new disk will be a shared EVS disk. Such a disk can be attached to multiple ECSs.	Disable
Disk Name	The disk name can contain only letters, digits, underscores (_), and hyphens (-). When applying for a single disk, ensure that the disk name contains not more than 63 characters.	-
Quantity	Specifies the number of applied EVS disks. You cannot create EVS disks in batches using a backup. You can apply for only one EVS disk at a time.	1

Parameter	Description	Example Value
Required Duration	Specifies the validity period of the EVS disk that you apply for. <ul style="list-style-type: none">If you select Unlimited, the new EVS disk has no expiration date.If you select 1 year, the validity period of the new EVS disk will be one year, which is subject to the expiration date displayed on the console.If you select Custom, specify an expiration date for the EVS disk that you apply for.	Unlimited
Description	Describes the EVS disk that you apply for. The length cannot exceed 63 characters.	-

Step 5 Click **Next**.

Step 6 Confirm that the information is correct and click **Add to Cart** or **Apply Now**.

If the configuration is incorrect, click **Back**.

- If you click **Add to Cart**, go to [Step 7](#).
- If you click **Apply Now**, go to [Step 8](#).

Step 7 Submit an application for the service in the shopping cart.

- Click the shopping cart in the upper right corner of the page.
- Select the service and click **Apply Now**.
- Enter the order information and click **OK** to submit the application.

Step 8 If the application for an EVS disk requires approval, contact the administrator for approval. Otherwise, skip this step.

 **NOTE**

The application for an EVS disk can be rejected by the administrator. If you have entered incorrect configuration when applying for an EVS disk, you can contact the administrator to reject the application, correct the configuration, and submit the application again.

Step 9 On the **Elastic Volume Service** page, view the status of the EVS disk.

During EVS disk creation, the four intermediate states **Creating**, **Available**, **Restoring**, and **Available** will be displayed in sequence. After the state has changed from **Creating** to **Available**, the disk has been successfully created. After the state has changed from **Restoring** to **Available**, backup data has been successfully restored to the created disk.

 **NOTE**

When a new EVS disk created using a VBS backup is attached to an instance, the architecture of the instance must be the same as that of the instance to which the source disk of the backup is attached. For example, they use the ARM architecture.

----End

VBS Console

For details about how to create an EVS disk using a backup, see **Operation Help Center > DR & Backup > User Guide > Restoring a Disk Using a Backup > Creating a Disk Using Backup Data**.

3.4.6 Applying for an EVS Disk Using an Image

You can apply for a data disk using a data disk image or apply for a system disk using a system disk image. The common application scenario is as follows: When you need to quickly deploy multiple identical services for development and testing, you can create an image for the EVS disk in the current environment (for example, AZ 1) and apply for an EVS disk that contains the same data in the local environment or another environment (for example, AZ 2) using the image.

Prerequisites

You have applied for data disk image or system disk image.

For details about how to apply for an image, see "Creating a Public Image", "Creating a Private Image", or "Creating a Shared Image" in "Operation Help Center" > "Compute" > "User Guide".

Procedure

- Step 1** Log in to the EVS console as a VDC administrator or VDC operator. For details, see [Logging In to the EVS Console as a VDC Administrator or VDC Operator](#).
- Step 2** Click **Apply for EVS Disk**.
- Step 3** In the **Select Service** dialog box, select the target service and then click **Apply Now**.

NOTE

Services are created by administrators based on operation requirements. When creating a service, the administrator can lock service parameters (for example, apply for a fixed disk type or apply for resources in a specified AZ), specify the service publishing scope (for example, visible to only specific VDCs), and set the product approval process. For details about how to create more services, see "Creating a Common Service" in *HUAWEI CLOUD Stack 8.1.0 Resource Provisioning Guide*.

- Step 4** In the **Apply for EVS Disk** dialog box, set the parameters as prompted, as shown in [Table 3-14](#).

Table 3-14 Parameters for applying for an EVS disk using an image

Parameter	Description	Example Value
AZ	<p>Specifies the availability zone (AZ) where an EVS disk is to be created.</p> <p>NOTE</p> <ul style="list-style-type: none">EVS disks can be attached to instances only in the same AZ.If an AZ is bound to a tenant upon the tenant creation, the resource pools in this AZ are used as dedicated storage pools for this tenant and invisible to other tenants not bound to the AZ. Services of this tenant can run on dedicated physical devices without interference from other tenants.	az1.dc1
Data Source	<p>Create an EVS disk using an image. Select Create from image, click Select, and select a system disk image or data disk image.</p> <ul style="list-style-type: none">System disk image You can apply for a system disk using a public system disk image, private system disk image, or shared system disk image.Data disk image You can apply for a data disk using a private or shared data disk image. <p>NOTE</p> <ul style="list-style-type: none">If the backend storage is heterogeneous storage, you cannot select Create from image.In VMware and VRM scenarios, VIMS data storage do not support creating a system disk from image.The system disk applied for using a system disk image can be attached only to the ECS with the same image.	Create from image
Disk	Apply for a system disk or data disk.	Data disk
Disk Type	<p>Select a disk type.</p> <p>You can select the disk type created in "Storage Services" > "Elastic Volume Service (EVS for ECS)" > "Configuration Before Applying for an EVS Disk" > "(Optional) Creating a Disk Type" in <i>HUAWEI CLOUD Stack 8.1.0 Resource Provisioning Guide</i>. You can also select the disk type created during automatic installation and deployment using HUAWEI CLOUD Stack Deploy. The name of the automatically created disk type is the value of business_volume_type in deployment parameter summary file <code>xxx_export_all_v2_EN.xlsx</code>.</p>	-

Parameter	Description	Example Value
Capacity (GB)	EVS disk capacity. The capacity of the new EVS disk must be greater than or equal to that of the source disk. The EVS disk capacity can neither exceed the total capacity quota of EVS disks nor the capacity quota of the current disk type.	10 GB
Device Type	Only EVS disks of the VBD type can be applied for from images.	VBD
Disk Sharing	<ul style="list-style-type: none">• Disable indicates that the new disk will be a non-shared EVS disk.• Enable indicates that the new disk will be a shared EVS disk. Such a disk can be attached to multiple ECSs. You cannot apply for shared system disks.	Disable
Disk Name	The disk name can contain only letters, digits, underscores (_), and hyphens (-). When applying for a single disk, ensure that the disk name contains not more than 63 characters.	-
Quantity	Specifies the number of applied EVS disks. You cannot apply for EVS disks in batches using a data disk image. Apply for only one disk at a time.	1
Required Duration	Specifies the validity period of the EVS disk that you apply for. <ul style="list-style-type: none">• If you select Unlimited, the new EVS disk has no expiration date.• If you select 1 year, the validity period of the new EVS disk will be one year, which is subject to the expiration date displayed on the console.• If you select Custom, specify an expiration date for the EVS disk that you apply for.	Unlimited
Description	Describes the EVS disk that you apply for. The length cannot exceed 63 characters.	-

Step 5 Click **Next**.

Step 6 Confirm that the information is correct and click **Add to Cart** or **Apply Now**.

If the configuration is incorrect, click **Back**.

- If you click **Add to Cart**, go to [Step 7](#).
- If you click **Apply Now**, go to [Step 8](#).

Step 7 Submit an application for the service in the shopping cart.

1. Click the shopping cart in the upper right corner of the page.
2. Select the service and click **Apply Now**.
3. Enter the order information and click **OK** to submit the application.

Step 8 If the application for an EVS disk requires approval, contact the administrator for approval. Otherwise, skip this step.

 **NOTE**

The application for an EVS disk can be rejected by the administrator. If you have entered incorrect configuration when applying for an EVS disk, you can contact the administrator to reject the application, correct the configuration, and submit the application again.

Step 9 On the **Elastic Volume Service** page, view the status of the EVS disk. After the EVS disk is created and its status changes to **Available**, the EVS disk is successfully created.

----End

3.5 Attaching an EVS Disk

A created EVS disk can be used by an instance only after being attached to the instance. The EVS disk created together with an instance is automatically attached, requiring no manual attaching operations.

Restrictions

- The ECS supports the attaching of disks in VBD and SCSI modes.
- Regardless if a shared EVS disk or non-shared EVS disk is attached to an instance, the EVS disk and the instance must be in the same AZ.
- Data disks can only be attached to ECSs as data disks. System disks can be attached to ECSs as system disks or data disks.
- An EVS disk cannot be attached to an instance that has expired.
- An EVS disk cannot be attached to an instance that has been soft deleted.
- When a disk is attached to an ECS configured with the disaster recovery (DR) service (CSDR/CSHA/VHA), you must ensure that the disk is created using the same backend storage as the existing disk on the ECS.
- An EVS disk with snapshots of a VM can be attached only to the VM and cannot be attached to any other VM.
- Neither shared EVS disks nor SCSI EVS disks can be attached to an ECS that has the CSHA service configured.
- If the ECS uses the Windows operating system and the administrator set **Disk Device Type** to **ide** when registering the image, shut down the ECS before attaching the EVS disk to the ECS.
- If the ECS to which the EVS disk belongs has not been created, the EVS disk cannot be attached to another ECS.

Context

- If you want to attach a SCSI EVS disk to an ECS, check whether the ECS supports SCSI disks. For details, see [Requirements and Restrictions on Using SCSI EVS Disks](#).
- ECSs to which SCSI shared disks are attached must be selected from the same anti-affinity ECS group. For details about the reasons, see [SCSI Reservation](#).
- You can attach the disks in the following ways:
 - EVS console: To attach multiple EVS disks to different ECSs, perform related operations on the EVS console.
 - ECS console: To attach multiple EVS disks to one instance, perform related operations on the ECS console.

Operations on the ECS Console

Step 1 Log in to the EVS console as a VDC administrator or VDC operator. For details, see [Logging In to the EVS Console as a VDC Administrator or VDC Operator](#).

Step 2 Try either of the following to go to the **Attach Disk** page:

Method 1: On the **Elastic Cloud Server** page, locate the row that contains the target ECS, click **More** in the **Operation** column, and select **Attach Disk**.

Method 2: Click the name of the target ECS. On the displayed ECS details page, choose **EVS > Attach Disk**.

Step 3 On the **Attach Disk** page, determine whether to specify an attachment point.

- If you select **No**, the system automatically assigns an address.
Select **No** and select the disk to be attached. A data disk can only be attached to an ECS as a data disk. A system disk can be attached to an ECS as a system disk or a data disk.
- If you select **Yes**, you need to specify the device address.
Select **Yes** and select the disk to be attached and the attachment point.

NOTE

- Make sure that the ECS has no ongoing disk attaching task. Otherwise, disk attaching may fail.
- The device addresses displayed on the page are classified into the following types based on **Device Address** of images used by the ECS: a-x:x:x, a:x, and xxxx:xx:xx.x. If the disk is successfully attached and the ECS is running properly, you can view the details about the disk on the ECS details page, including **Device Name** and **Device Address**. **Device Name**: indicates the actual device name of the disk on the ECS. **Device Address**: indicates the address allocated by the system to the disk. If VMTools is not installed on the ECS or the version of VMTools is too early, only the device address is displayed on the page. For details about how to query the attachment point of a disk on an ECS based on the device address, see [Operation Help Center > Elastic Cloud Server > User Guide > Best Practices > Manually Viewing the Disk Mount Point](#).

If you want to apply for a new EVS disk, click **Apply for Disk**. For details, see section [Applying for an EVS Disk](#).


Step 4 Click **OK**.

 **NOTE**

If a disk is attached to a stopped ECS, the device name is not displayed on the page after the disk is attached. Instead, it is displayed several minutes after the ECS is started.

----End

Operations on the EVS Console

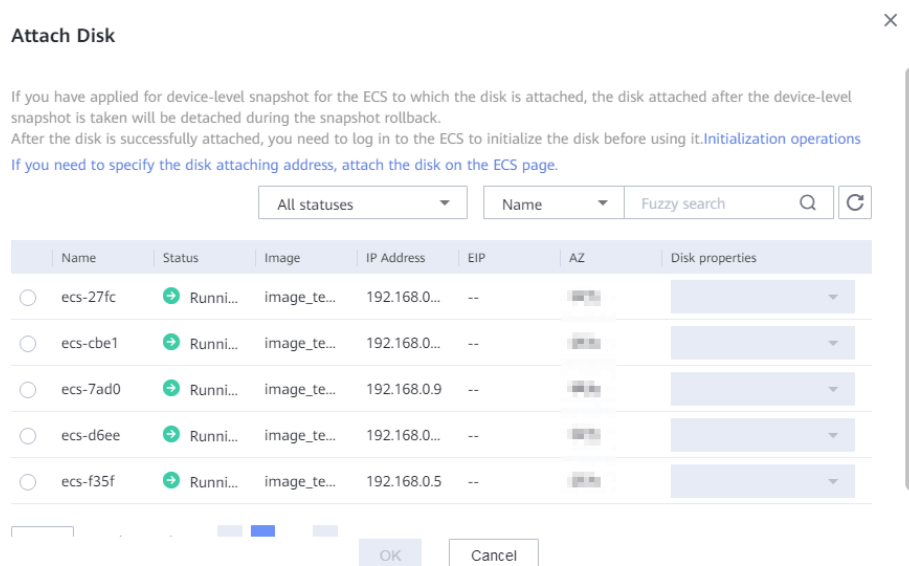
Step 1 Click  in the upper left corner, select a region and resource set, and choose **Storage > Elastic Volume Service**. The EVS console is displayed.

Step 2 On the **Elastic Volume Service** page, locate the row that contains the target EVS disk, and click **Attach** in the **Operation** column.

The disk can be either a shared disk or a non-shared disk.

Step 3 In the **Attach Disk** dialog box displayed, select the instance to which the disk is to be attached. A data disk can be attached to an instance only as a data disk. A system disk can be attached to an instance as a data disk or a system disk.

If you are attaching a shared disk, you can select multiple ECSs. A shared disk can be attached to a maximum of 16 ECSs by default.

 **NOTE**

To specify an address, you need to attach the disk on the ECS console. For details, see [Operations on the ECS Console](#).

Step 4 Click **OK**.

On the EVS disk list page, if the disk status is **In-use**, the EVS disk has been successfully attached to an instance.

----End

Follow-up Procedure

After the EVS disk is attached to the instance, perform subsequent operations by referring to [Table 3-15](#).

Table 3-15 Operations after the EVS disk is attached

EVS Disk	Instance OS	Follow-up Procedure
New blank data disk (Data Source is Do not specify)	Windows	To initialize the data disk of a Windows ECS, see Initializing a Windows Data Disk .
	Linux	<ul style="list-style-type: none">When the data disk capacity is less than 2 TB, see Initializing a Linux Data Disk (fdisk).When the data disk capacity is greater than and equal to 2 TB, see Initializing a Linux Data Disk (parted).
New blank system disk (Data Source is Do not specify)	Windows/ Linux	The system disk cannot be used to directly start an ECS. If you need to use it to start the ECS, copy the backup data of a normal system disk to the system disk.
Data disk created from a snapshot, an image, a backup, or an existing disk	Windows	<p>Check whether the source disk corresponding to snapshots, images, existing disks, and backups has been initialized.</p> <ul style="list-style-type: none">If yes, you can directly use the EVS disk without performing any other operations.If no, initialize the disk before using it. The operations are the same as those for a new blank data disk.

EVS Disk	Instance OS	Follow-up Procedure
	Linux	<p>Check whether the source disk corresponding to snapshots, images, existing disks, and backups has been initialized.</p> <ul style="list-style-type: none">• If yes, log in to the instance and run the mount command (mount <i>partition mount path</i>) to mount the partition. Configure automatic mounting upon system startup. Then, the EVS disk can be used properly. For details about how to set automatic mounting upon system startup, see Setting Automatic Disk Attachment Upon Instance Start.• If no, initialize the disk before using it. The operations are the same as those for a new blank data disk.
System disk created from an image, a snapshot, an existing disk, or a backup	Windows/Linux	You can directly use the system disk without performing any other operations.

3.6 Initializing a Data Disk

3.6.1 Initialization Overview

After attaching an EVS disk to an instance, you need to log in to the instance to partition and initialize the disk so that the data disk can be used for the instance. [Table 3-16](#) provides common partition styles.

Table 3-16 Disk partition style

Disk Partition Style	Maximum Disk Capacity Supported	Maximum Number of Partitions Supported	Partitioning Tool Supported
Main Boot Record (MBR)	2 TB	<ul style="list-style-type: none">• Four primary partitions• Three primary partitions and one extended partition	<ul style="list-style-type: none">• For Linux OS fdisk or parted• For Windows OS Disk management

Disk Partition Style	Maximum Disk Capacity Supported	Maximum Number of Partitions Supported	Partitioning Tool Supported
GUID Partition Table (GPT)	18 EB NOTE 1 EB = 1,048,576 TB	Unlimited	<ul style="list-style-type: none">• For Linux OS parted• For Windows OS Disk management

3.6.2 Initializing a Windows Data Disk

A data disk attached to an instance or created together with an instance can be used by the instance only after being initialized. This section uses the Windows Server 2008 R2 Enterprise OS as an example. Specific initialization operations vary with OSs.

Prerequisites

- For details about how to log in to an ECS, see **Operation Help Center > Compute > Elastic Cloud Server > User Guide > Logging In to an ECS**.
- A disk has been attached to an instance and has not been initialized.

Context

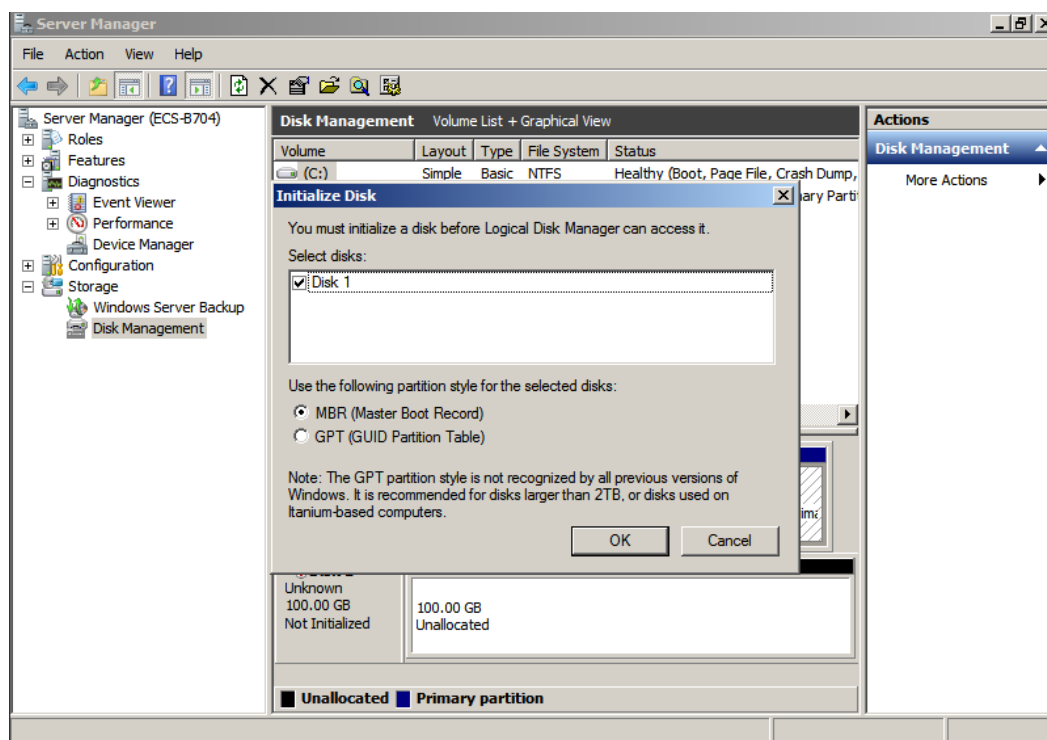
Initializing a data disk is a high-risk operation. If the data disk contains useful data, perform the operations described in [Applying for a Snapshot](#) or [Creating a Backup](#) for the data disk first.

Procedure

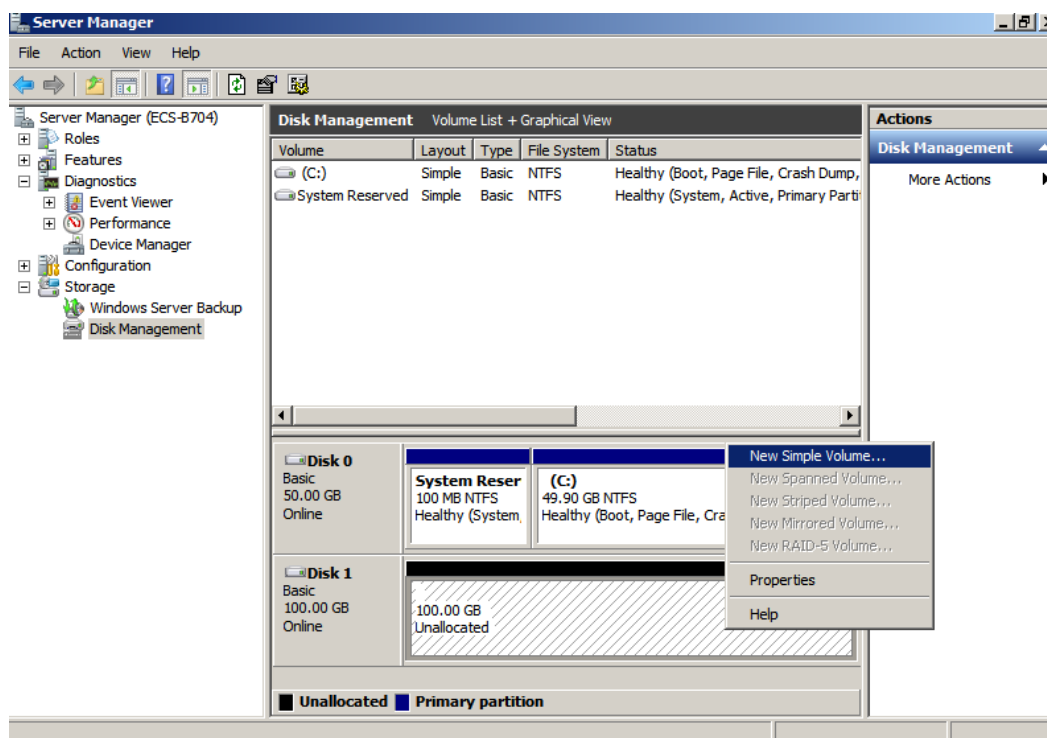
- Step 1** In desktop, right-click **Computer** and choose **Manage** from the shortcut menu.
The **Server Manager** page is displayed.
- Step 2** In the navigation pane, choose **Storage > Disk Management**.
- Step 3** If the disk to be initialized in the disk list is in **Offline** state, right-click in the disk area and choose **Online** from the shortcut menu.
Then, the disk status changes from **Offline** to **Uninitialized**.
- Step 4** Right-click in the disk area and choose **Initialize Disk** from the shortcut menu. In the displayed **Initialize Disk** dialog box, select **MBR (Master Boot Record)** and click **OK**.

NOTE

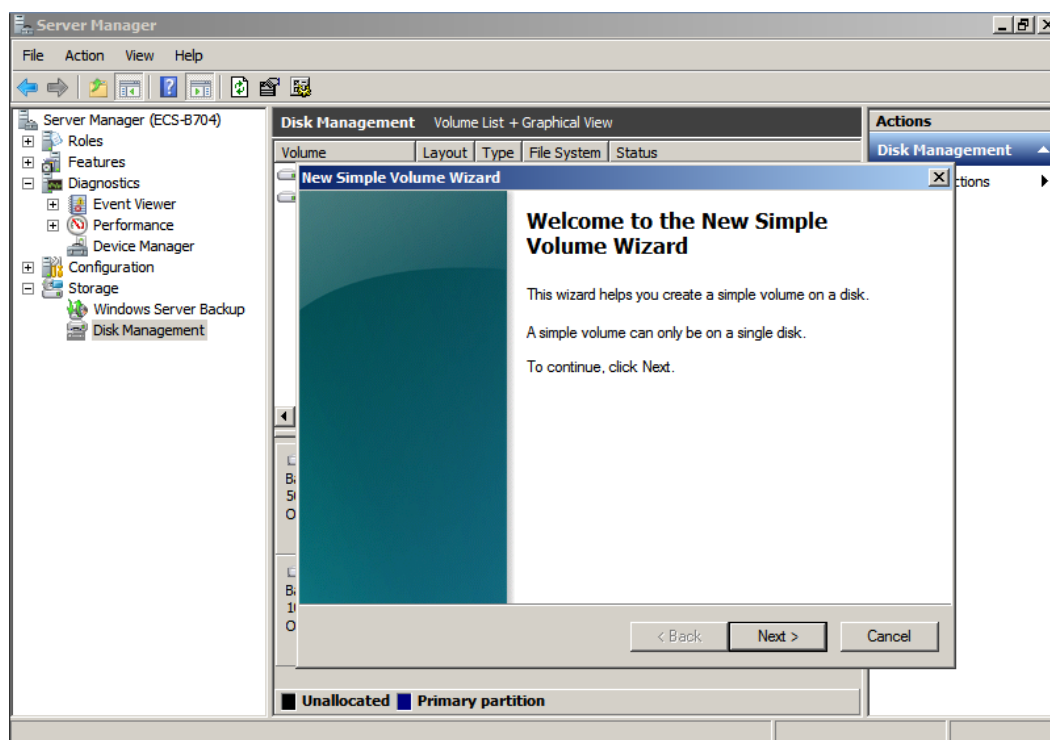
If the data disk to be initialized is larger than 2 TB, select **GPT (GUID Partition Table)** in the dialog box.



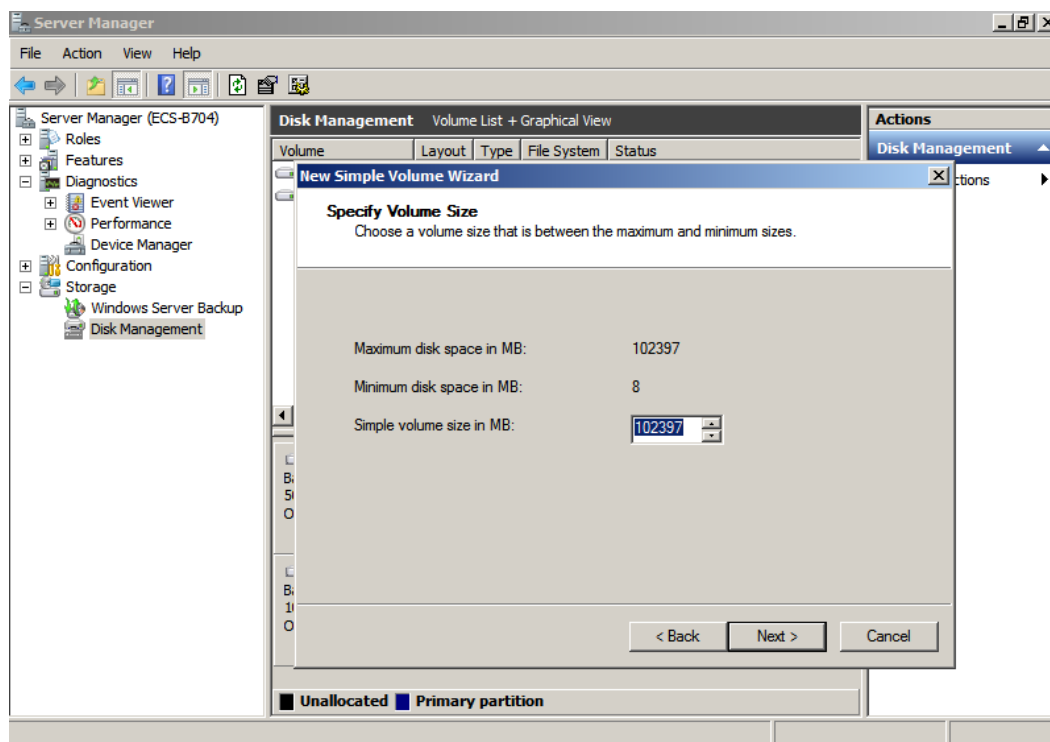
Step 5 Right-click at the unallocated disk space and choose **New Simple Volume** from the shortcut menu.



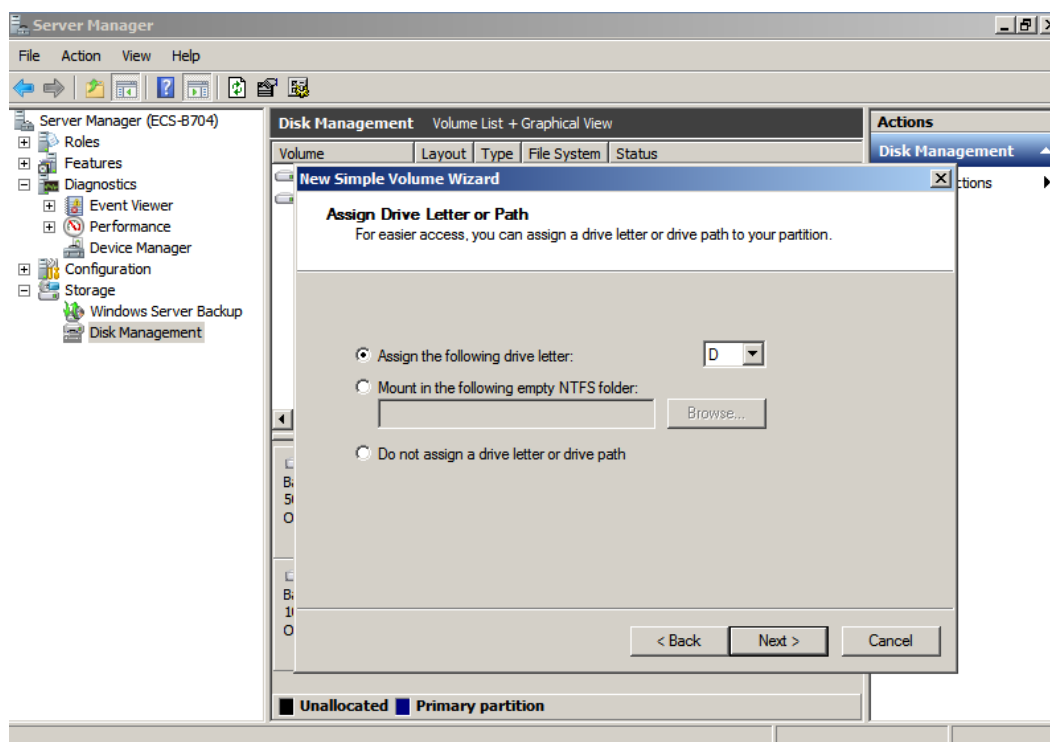
Step 6 On the displayed **New Simple Volume Wizard** page, click **Next**.



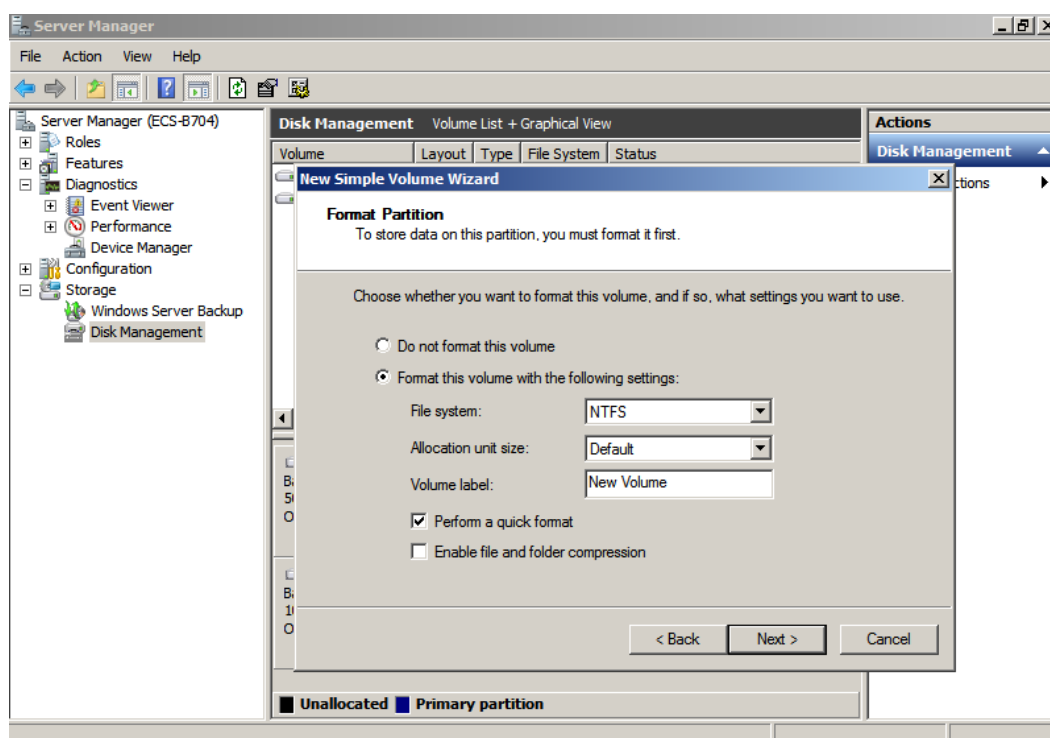
Step 7 Specify the simple volume size as required (the default value is the maximum) and click **Next**.



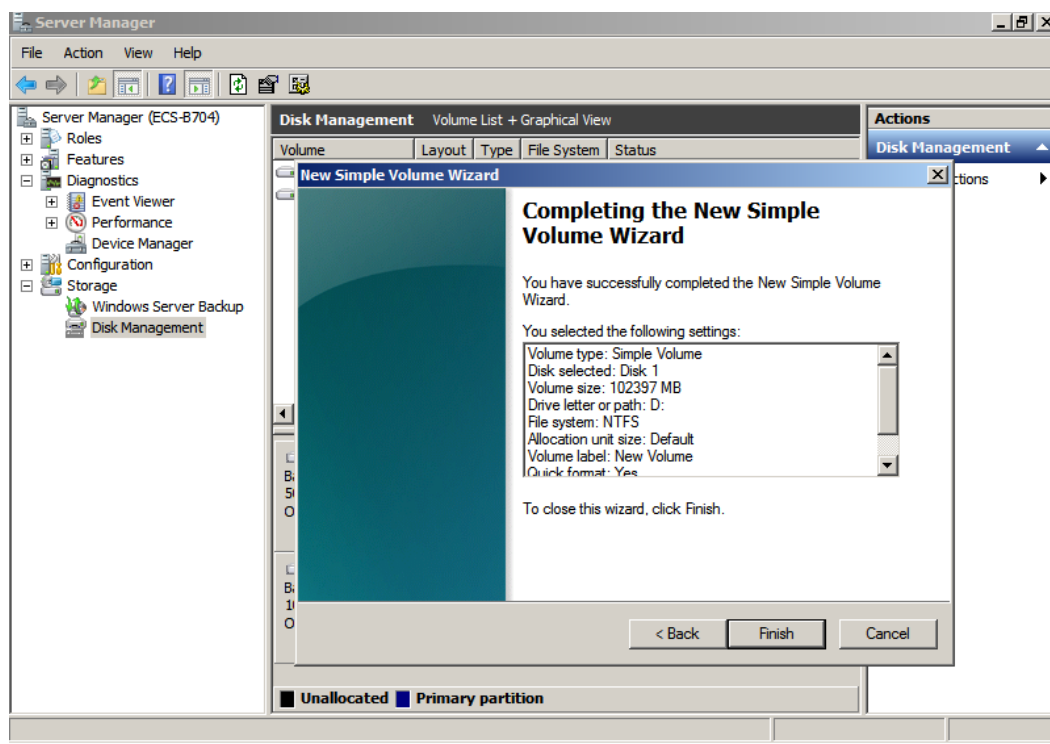
Step 8 Assign the driver letter and click **Next**.



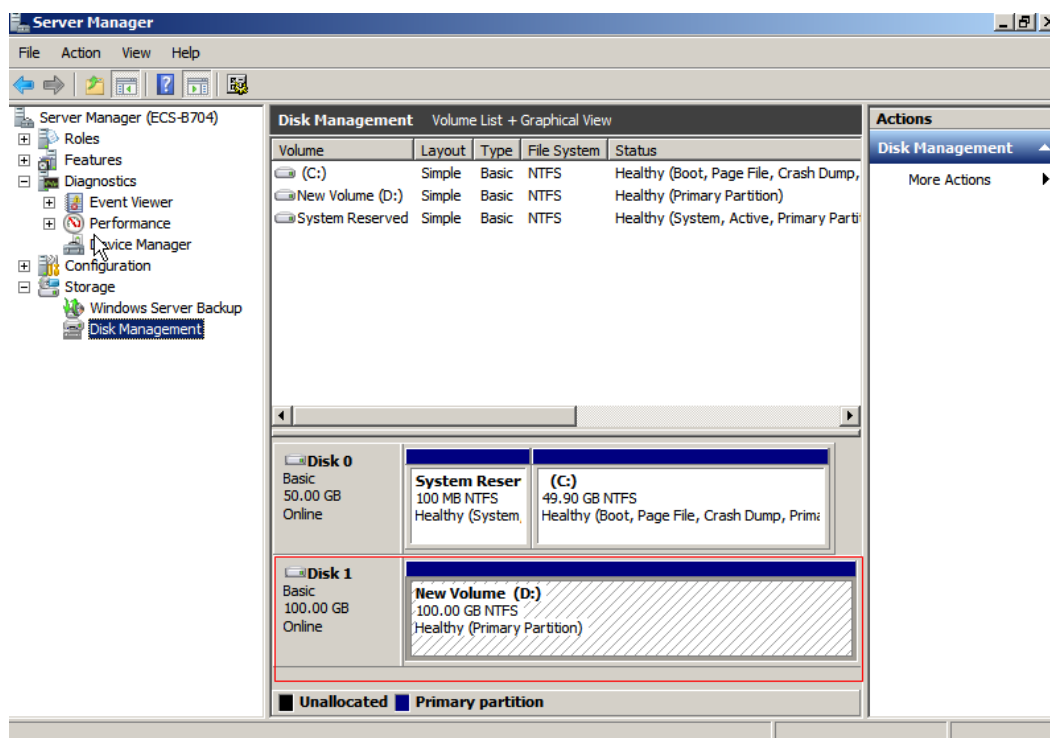
Step 9 Select **Format this volume with the following settings**, set parameters based on the actual requirements, and select **Perform a quick format**. Then click **Next**.



Step 10 Click **Finish**.



Wait for the initialization to complete. When the volume status changes to **Healthy**, the initialization has finished.



----End

3.6.3 Initializing a Linux Data Disk (fdisk)

A data disk attached to an instance or created together with an instance can be used by the instance only after being initialized. This section uses an instance running CentOS 7.0 (64-bit) as an example, and uses the fdisk partition tool to set up partitions for the data disk. Initialization operations vary with operating systems.

Prerequisites

- For details about how to log in to an ECS, see **Operation Help Center > Compute > Elastic Cloud Server > User Guide > Logging In to an ECS**.
- A disk has been attached to an instance and has not been initialized.

Context

Both the fdisk and parted can be used to partition a Linux data disk. For a disk larger than 2 TB, only parted can be used because fdisk cannot partition such a large disk. For details, see [Using the parted tool](#).

Creating Partitions and Mounting a Disk

The following example shows how to create a new primary partition on a new data disk that has been attached to an instance. The primary partition will be created using fdisk, and MBR is the default partition style. Furthermore, the partition will be formatted using the ext4 file system, mounted on the `/mnt/sdc` directory, and set to be automatically mounted upon a system start.

Step 1 Run the following command to view information about the added data disk:

fdisk -l

Information similar to the following is displayed: (In the command output, the server contains two disks. `/dev/xvda` is the system disk, and `/dev/xvdb` is the added data disk.)

NOTE

If you do not log in to the ECS and run the **umount** command but directly detach the `/dev/xvdb` or `/dev/vdb` EVS disk on the management console, the disk name in the ECS may encounter a release delay. When you attach the disk to the server again, the mount point displayed on the management console may be inconsistent with that in the server. For example, device name `/dev/sdb` or `/dev/vdb` is selected for attachment, but `/dev/xvdc` or `/dev/vdc` may be displayed as the disk name in the OS. This issue does not adversely affect services.

```
[root@ecs-b656 test]# fdisk -l
```

```
Disk /dev/xvda: 42.9 GB, 42949672960 bytes, 83886080 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x000cc4ad
```

Device	Boot	Start	End	Blocks	Id	System
/dev/xvda1	*	2048	2050047	1024000	83	Linux
/dev/xvda2		2050048	22530047	10240000	83	Linux
/dev/xvda3		22530048	24578047	1024000	83	Linux

```
/dev/xvda4 24578048 83886079 29654016 5 Extended
/dev/xvda5 24580096 26628095 1024000 82 Linux swap / Solaris
```

Disk /dev/xvdb: 10.7 GB, 10737418240 bytes, 20971520 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes

NOTE

The capacity displayed here is inconsistent with the capacity of the EVS disk applied for on ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios). The reason is as follows: The capacity of EVS disks is calculated using the unit of GiB (Gibibyte), while the capacity unit in Linux OS is GB (Gigabyte). The GiB is calculated in binary mode, and the GB is calculated in decimal format. 1 GiB = 1,073,741,824 Bytes and 1 GB = 1,000,000,000 Bytes.

Step 2 Run the following command to allocate partitions for the added data disk using `fdisk`:

fdisk *Newly added data disk*

In this example, `/dev/xvdb` is the newly added data disk.

fdisk /dev/xvdb

Information similar to the following is displayed:

```
[root@ecs-b656 test]# fdisk /dev/xvdb
Welcome to fdisk (util-linux 2.23.2).
Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.
Device does not contain a recognized partition table
Building a new DOS disklabel with disk identifier 0xb00005bd.
Command (m for help):
```

Step 3 Enter `n` and press **Enter**.

Entering `n` creates a partition.

There are two types of disk partitions:

- Choosing `p` creates a primary partition.
- Choosing `e` creates an extended partition.

```
Command (m for help): n
Partition type:
 p primary (0 primary, 0 extended, 4 free)
 e extended
```

Step 4 Enter `p` and press **Enter**.

The following describes how to create a primary partition.

Information similar to the following is displayed: (**Partition number** indicates the serial number of the primary partition. The value can be **1** to **4**.)

```
Select (default p): p
Partition number (1-4, default 1):
```

Step 5 Enter the primary partition number **1** and press **Enter**.

For example, select **1** as the partition number.

Information similar to the following is displayed: (**First sector** indicates the first sector number. The value can be **2048** to **20971519**, and the default value is **2048**.)

```
Partition number (1-4, default 1): 1
First sector (2048-20971519, default 2048):
```

Step 6 Press **Enter**.

The default start sector number 2048 is used as an example.

Information similar to the following is displayed: (**Last sector** indicates the last sector number. The value can be from **2048** to **20971519**, and the default value is **20971519**.)

```
First sector (2048-20971519, default 2048):
Using default value 2048
Last sector, +sectors or +size{K,M,G} (2048-20971519, default 20971519):
```

Step 7 Press **Enter**.

The default last sector number 20971519 is used as an example.

Information similar to the following is displayed, indicating that a primary partition is created for a 10 GB data disk.

```
Last sector, +sectors or +size{K,M,G} (2048-20971519, default 20971519):
Using default value 20971519
Partition 1 of type Linux and of size 10 GiB is set
Command (m for help):
```

Step 8 Enter **p** and press **Enter** to view the details about the created partition.

Information similar to the following is displayed: (Details about the **/dev/xvdb1** partition are displayed.)

```
Command (m for help): p

Disk /dev/xvdb: 10.7 GB, 10737418240 bytes, 20971520 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0xb00005bd

   Device Boot      Start         End      Blocks   Id  System
/dev/xvdb1        2048     20971519    10484736    83   Linux

Command (m for help):
```

Step 9 Enter **w** and press **Enter** to write the changes into the partition table.

Information similar to the following is displayed: (The partition is successfully created.)

```
Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.
```

Step 10 Run the following command to synchronize the new partition table to the data disk:

partprobe

Step 11 Run the following command to set the format for the file system of the newly created partition:

mkfs -t *File system format* /dev/xvdb1

For example, run the following command to set the **ext4** file system for the **/dev/xvdb1** partition:

```
mkfs -t ext4 /dev/xvdb1
```

Information similar to the following is displayed:

```
[root@ecs-b656 test]# mkfs -t ext4 /dev/xvdb1
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
655360 inodes, 2621184 blocks
131059 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2151677952
80 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632

Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

NOTE

The formatting takes a period of time. Observe the system running status and do not exit.

Step 12 Run the following command to create a mount directory:

```
mkdir Mount directory
```

/mnt/sdc is used in this example.

```
mkdir /mnt/sdc
```

Step 13 Run the following command to mount the new partition to the mount directory created in [Step 12](#):

```
mount /dev/xvdb1 Mount directory
```

For example, run the following command to mount the newly created partition on **/mnt/sdc**:

```
mount /dev/xvdb1 /mnt/sdc
```

Step 14 Run the following command to view the mount result:

```
df -TH
```

Information similar to the following is displayed. The newly created **/dev/xvdb1** partition has been mounted on **/mnt/sdc**.

```
[root@ecs-b656 test]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/xvda2      xfs       11G   7.4G  3.2G  71% /
devtmpfs        devtmpfs  4.1G   0  4.1G   0% /dev
tmpfs           tmpfs     4.1G  82k  4.1G   1% /dev/shm
tmpfs           tmpfs     4.1G  9.2M  4.1G   1% /run
tmpfs           tmpfs     4.1G   0  4.1G   0% /sys/fs/cgroup
/dev/xvda3      xfs       1.1G   39M  1.1G   4% /home
```

```
/dev/xvda1  xfs    1.1G 131M 915M 13% /boot
/dev/xvdb1  ext4    11G  38M 9.9G  1% /mnt/sdc
```

----End

Setting Automatic Disk Attachment Upon Instance Start

If you require a disk to be automatically attached to an instance when the instance is started, enable automatic disk attachment upon an instance start by referring to operations provided in this section. When enabling automatic disk attachment, you cannot directly specify **/dev/xvdb1** in **/etc/fstab**. This is because the sequence codes of the instance may change during an instance stop or start process. You are advised to use the universally unique identifier (UUID) in **/etc/fstab** to automatically attach the disk at a system start.

NOTE

The UUID of a disk is a character string that uniquely identifies a storage device in a Linux system.

Step 1 Run the following command to query the partition UUID:

blkid *Disk partition*

For example, run the following command to query the UUID of **/dev/xvdb1**:

blkid /dev/xvdb1

Information similar to the following is displayed: (The UUID of **/dev/xvdb1** is displayed.)

```
[root@ecs-b656 test]# blkid /dev/xvdb1
/dev/xvdb1: UUID="1851e23f-1c57-40ab-86bb-5fc5fc606ffa" TYPE="ext4"
```

Step 2 Run the following command to open the **fstab** file using the vi editor:

vi /etc/fstab

Step 3 Press **i** to enter the editing mode.

Step 4 Move the cursor to the end of the file and press **Enter**. Then add the following information:

UUID=xxx attachment directory file system defaults 0 2

Assuming that the file system is **ext4** and the attachment directory is **/mnt/sdc**.
UUID=1851e23f-1c57-40ab-86bb-5fc5fc606ffa /mnt/sdc ext4 defaults 0 2

NOTICE

After automatic attachment upon instance start is configured, comment out or delete the line in the **fstab** file before detaching the disk. Otherwise, you may fail to access the OS after the disk is detached.

Step 5 Press **Esc**, enter **:wq**, and press **Enter**.

The system saves the configuration and exits the vi editor.

----End

3.6.4 Initializing a Linux Data Disk (parted)

A data disk attached to an instance or created together with an instance can be used by the instance only after being initialized. This section uses an instance running CentOS 7.0 (64-bit) as an example, and uses the parted partition tool to set up partitions for the data disk. Initialization operations vary with operating systems.

Prerequisites

- For details about how to log in to an ECS, see [Operation Help Center > Compute > Elastic Cloud Server > User Guide > Logging In to an ECS](#).
- A disk has been attached to an instance and has not been initialized.

Mounting Partitions to a Disk

The following example shows how to create a partition on a new data disk that has been attached to an instance. The partition will be created using parted and GPT is the default partition style. Furthermore, the partition will be formatted using the ext4 file system, mounted on the **/mnt/sdc** directory, and set to be automatically mounted upon a system start.

Step 1 Run the following command to view information about the added data disk:

lsblk

Information similar to the following is displayed:

NOTE

If you do not log in to the ECS and run the **umount** command but directly detach the **/dev/xvdb** or **/dev/vdb** EVS disk on the management console, the disk name in the ECS may encounter a release delay. When you attach the disk to the server again, the mount point displayed on the management console may be inconsistent with that in the server. For example, device name **/dev/sdb** or **/dev/vdb** is selected for attachment, but **/dev/xvdc** or **/dev/vdc** may be displayed as the disk name in the OS. This issue does not adversely affect services.

```
[root@ecs-centos-70 linux]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
xvda 202:0 0 40G 0 disk
├─xvda1 202:1 0 4G 0 part [SWAP]
└─xvda2 202:2 0 36G 0 part /
xvdb 202:16 0 10G 0 disk
```

The command output indicates that the server contains two disks. **/dev/xvda** is the system disk and **/dev/xvdb** is the new data disk.

Step 2 Run the following command to enter parted to partition the added data disk:

parted *Added data disk*

In this example, **/dev/xvdb** is the newly added data disk.

parted /dev/xvdb

Information similar to the following is displayed:

```
[root@ecs-centos-70 linux]# parted /dev/xvdb
GNU Parted 3.1
```

```
Using /dev/xvdb
Welcome to GNU Parted! Type 'help' to view a list of commands.
```

Step 3 Enter **p** and press **Enter** to view the current disk partition style.

Information similar to the following is displayed:

```
(parted) p
Error: /dev/xvdb: unrecognised disk label
Model: Xen Virtual Block Device (xvd)
Disk /dev/xvdb: 10.7GB
Sector size (logical/physical): 512B/512B
Partition Table: unknown
Disk Flags:
```

In the command output, the **Partition Table** value is **unknown**, indicating that the disk partition style is unknown.

NOTE

The capacity displayed here is inconsistent with the capacity of the EVS disk applied for on ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios). The reason is as follows: The capacity of EVS disks is calculated using the unit of GiB (Gibibyte), while the capacity unit in Linux OS is GB (Gigabyte). The GiB is calculated in binary mode, and the GB is calculated in decimal format. 1 GiB = 1,073,741,824 Bytes and 1 GB = 1,000,000,000 Bytes.

Step 4 Run the following command to set the disk partition style:

mklabel *Disk partition style*

The disk partition styles include MBR and GPT. For example, run the following command to set the partition style to GPT:

mklabel gpt

NOTICE

If you change the disk partition style after the disk has been used, the original data on the disk will be cleared. Therefore, select a proper disk partition style when initializing the disk.

Step 5 Enter **p** and press **Enter** to view the disk partition style.

Information similar to the following is displayed:

```
(parted) mklabel gpt
(parted) p
Model: Xen Virtual Block Device (xvd)
Disk /dev/xvdb: 20971520s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start  End  Size  File system  Name  Flags
```

Step 6 Enter **unit s** and press **Enter** to set the measurement unit of the disk to sector numbers.

Step 7 Enter **mkpart opt 2048s 100%** and press **Enter**.

In the command, **opt** is the name of the new partition, **2048s** indicates the start of the partition, and **100%** indicates the end of the partition. You can plan the number and capacity of disk partitions based on service requirements.

Information similar to the following is displayed:

```
(parted) mkpart opt 2048s 100%
Warning: The resulting partition is not properly aligned for best performance.
Ignore/Cancel? Cancel
```

If the preceding warning message is displayed, enter **Cancel** to stop the partitioning. Then, find the first sector with the best disk performance and use that value to partition the disk.

Step 8 Enter **p** and press **Enter** to view the details about the created partition.

Information similar to the following is displayed:

```
(parted) p
Model: Xen Virtual Block Device (xvd)
Disk /dev/xvdb: 20971520s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number Start End      Size      File system Name Flags
 1      2048s 20969471s 20967424s                opt
```

Details about the **/dev/xvdb1** partition are displayed.

Step 9 Enter **q** and press **Enter** to exit parted.

Step 10 Run the following command to view the disk partition information:

lsblk

Information similar to the following is displayed:

```
[root@ecs-centos-70 linux]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
xvda 202:0  0  40G  0 disk
├─xvda1 202:1  0   4G  0 part [SWAP]
└─xvda2 202:2  0  36G  0 part /
xvdb 202:16  0 100G  0 disk
└─xvdb1 202:17  0 100G  0 part
```

In the command output, **/dev/xvdb1** is the partition you created.

Step 11 Run the following command to set the format for the file system of the newly created partition:

NOTICE

The partition sizes supported by file systems vary. Therefore, you are advised to choose an appropriate file system based on your service requirements.

mkfs -t *File system format* **/dev/xvdb1**

For example, run the following command to set the **ext4** file system for the **/dev/xvdb1** partition:

mkfs -t ext4 /dev/xvdb1

Information similar to the following is displayed:

```
[root@ecs-centos-70 linux]# mkfs -t ext4 /dev/xvdb1
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
655360 inodes, 2620928 blocks
131046 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2151677925
80 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632

Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

The formatting takes a period of time. Observe the system running status, and do not exit.

Step 12 Run the following command to create a mount point:

mkdir *Mount point*

For example, run the following command to create the **/mnt/sdc** mount point:

mkdir **/mnt/sdc**

Step 13 Run the following command to mount the new partition to the mount point created in [Step 12](#):

mount **/dev/xvdb1** *Mount point*

For example, run the following command to mount the newly created partition on **/mnt/sdc**:

mount **/dev/xvdb1** **/mnt/sdc**

Step 14 Run the following command to view the mount result:

df -TH

Information similar to the following is displayed:

```
[root@ecs-centos-70 linux]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/xvda2      xfs       39G   4.0G   35G   11% /
devtmpfs        devtmpfs  946M    0  946M    0% /dev
tmpfs           tmpfs     954M    0  954M    0% /dev/shm
tmpfs           tmpfs     954M   9.1M  945M    1% /run
tmpfs           tmpfs     954M    0  954M    0% /sys/fs/cgroup
/dev/xvdb1      ext4      11G   38M  101G    1% /mnt/sdc
```

The newly created **/dev/xvdb1** is mounted on **/mnt/sdc**.

----End

Setting Automatic Disk Attachment at a System Start

If you require a disk to be automatically attached to an instance when the instance is started, enable automatic disk attachment upon an instance start by

referring to operations provided in this section. When enabling automatic disk attachment, you cannot directly specify **/dev/xvdb1** in **/etc/fstab**. This is because the sequence codes of the instance may change during an instance stop or start process. You are advised to use the universally unique identifier (UUID) in **/etc/fstab** to automatically attach the disk at a system start.

NOTE

The UUID of a disk is a character string that uniquely identifies a storage device in a Linux system.

Step 1 Run the following command to query the partition UUID:

blkid *Disk partition*

For example, run the following command to query the UUID of **/dev/xvdb1**:

blkid /dev/xvdb1

Information similar to the following is displayed: (The UUID of **/dev/xvdb1** is displayed.)

```
[root@ecs-b656 test]# blkid /dev/xvdb1
/dev/xvdb1: UUID="1851e23f-1c57-40ab-86bb-5fc5fc606ffa" TYPE="ext4"
```

Step 2 Run the following command to open the **fstab** file using the vi editor:

vi /etc/fstab

Step 3 Press **i** to enter the editing mode.

Step 4 Move the cursor to the end of the file and press **Enter**. Then add the following information:

UUID=xxx attachment directory file system defaults 0 2

Assuming that the file system is **ext4** and the attachment directory is **/mnt/sdc**.
UUID=1851e23f-1c57-40ab-86bb-5fc5fc606ffa /mnt/sdc ext4 defaults 0 2

NOTICE

After automatic attachment upon instance start is configured, comment out or delete the line in the **fstab** file before detaching the disk. Otherwise, you may fail to access the OS after the disk is detached.

Step 5 Press **Esc**, enter **:wq**, and press **Enter**.

The system saves the configuration and exits the vi editor.

----End

3.7 Expanding EVS Disk Capacity

3.7.1 Overview

You can expand EVS disk capacity if the disk capacity becomes insufficient. For system disks, you can only expand the capacity of existing disks. For data disks,

you can either expand the capacity of existing disks or add data disks by referring to [add new disks](#) and attach them to instances.

After expanding the capacity of a disk, you need to create partitions for the new capacity or create partitions to replace the original disk partitions.

- **Creating partitions for the new capacity**
When using this method, you need to create partitions, without unmounting any existing partitions. This operation does not interrupt ongoing services and has minor impacts on services. This method is recommended for system disks or disks carrying services that cannot be interrupted. If the MBR partition style is used, the disk capacity must be less than 2 TB and the number of partitions does not exceed the upper limit after the expansion.
- **Creating partitions to replace existing ones**
If the MBR partition style is used and the number of disk partitions has reached the upper limit, new partitions cannot be created. In this situation, you need to unmount existing partitions first and then create new ones to replace them. This operation does not delete data in existing partitions, but services must be suspended during the operation, affecting ongoing services.

[Table 3-17](#) describes the post-expansion operations.

Table 3-17 Post-expansion operations for an EVS disk

OS	Disk Capacity After Expansion	Current Disk Partition Style	Post-Expansion Operation
Linux	≤ 2 TB	MBR	Use the fdisk or parted tool to create partitions for the new capacity. Operations After Expanding the Capacity of a Disk in Linux (Adding Partitions Using fdisk) Operations After Expanding the Capacity of a Disk in Linux (Adding Partitions Using parted)
			Use the fdisk or parted tool to create new partitions to replace existing ones. This operation interrupts ongoing user services and is not recommended after system disk capacity expansion. Operations After Expanding the Capacity of a Disk in Linux (Replacing the Original Partitions Using fdisk) Operations After Expanding the Capacity of a Disk in Linux (Replacing the Original Partitions Using parted)

OS	Disk Capacity After Expansion	Current Disk Partition Style	Post-Expansion Operation
		GPT	Use the parted tool to add partitions for the new capacity or create partitions to replace the original partitions. Operations After Expanding the Capacity of a Disk in Linux (Adding Partitions Using parted) Operations After Expanding the Capacity of a Disk in Linux (Replacing the Original Partitions Using parted)
	> 2 TB	MBR	Use the parted tool to change the partition style from MBR to GPT. However, this operation will clear disk data.
		GPT	Use the parted tool to add partitions for the new capacity or create partitions to replace the original partitions. Operations After Expanding the Capacity of a Disk in Linux (Adding Partitions Using parted) Operations After Expanding the Capacity of a Disk in Linux (Replacing the Original Partitions Using parted)
Windows	-	MBR	Allocate the additional disk space to existing partitions.
		GPT	Operations After Expanding the Capacity of a Disk in Windows

3.7.2 Expanding Disk Capacity Online

This section describes how to expand the capacity of an EVS disk attached to an instance.

Restrictions

- When you expand the capacity of a disk online, the instance to which the disk is attached must be in the **Running** or **Stopped** state.
- Shared EVS disks do not support online capacity expansion, that is, the capacity of a shared EVS disk can be expanded only when the disk is in the **Available** state.
- The capacity of a disk configured with the DR service (CSHA/CSDR/VHA) cannot be expanded.

- When the backend storage is Huawei SAN storage (OceanStor V3/V5 series or OceanStor Dorado V3 series/6.x) or heterogeneous storage, if the EVS disk has snapshots, capacity expansion is not supported. When the backend storage is Huawei Distributed Block Storage, capacity expansion can be performed for an EVS disk with snapshots.
- If backend storage of the disk is heterogeneous storage, online capacity expansion is not supported while offline capacity expansion is supported.
- Capacity expansion cannot be performed when the disk status is **Reserved** or **Maintenance**.
- For the OSs that support online expansion, see [Table 3-18](#).

Table 3-18 OSs

OS	Version
CentOS	7.3 64-bit
	7.2 64-bit
	6.8 64-bit
	6.7 64-bit
	6.5 64-bit
Debian	8.6.0 64-bit
	8.5.0 64-bit
Fedora	25 64-bit
	24 64-bit
SUSE	SUSE Linux Enterprise Server 12 SP2 (64-bit)
	SUSE Linux Enterprise Server 12 SP1 (64-bit)
	SUSE Linux Enterprise Server 11 SP4 (64-bit)
	SUSE Linux Enterprise Server 12 (64-bit)
OpenSUSE	42.2 64-bit
	42.1 64-bit
Oracle Linux Server release	7.3 64-bit
	7.2 64-bit
	6.8 64-bit
	6.7 64-bit
Ubuntu Server	16.04 64-bit
	14.04 64-bit
	14.04.4 64-bit

OS	Version
Windows	Windows Server 2008 R2 Enterprise 64-bit
	Windows Server 2012 R2 Standard 64-bit
	Windows Server 2016 R2 Standard 64-bit
Red Hat Enterprise Linux (RHEL)	7.3 64-bit
	6.8 64-bit

Procedure

Step 1 [Log in to EVS Console as a VDC administrator or VDC operator](#).

Step 2 In the EVS disk list, locate the row that contains the target disk, click **More** in the **Operation** column, and choose **Expand Capacity**.

The **Expand Capacity** page is displayed.

NOTE

To view information about the instance to which the EVS disk is attached, click the instance name in the **Attaching Information** column.

Step 3 Set **Added Capacity (GB)** as prompted and click **Next**.

Step 4 On the **Resource Details** page, confirm the EVS disk specifications.

- If you do not need to modify the specifications, click **Apply Now** to start the EVS disk capacity expansion.
- If you need to modify the specifications, click **Previous** to modify parameters.

After the expansion is submitted, the disk list page is displayed.

Step 5 If the EVS disk whose capacity is to be expanded requires approval, contact the administrator for approval. Otherwise, skip this step.

On the **Elastic Volume Service** page, view the capacity of the EVS disk. If the disk capacity has increased, the expansion is successful.

----End

Follow-up Procedure

After you have expanded the capacity of an EVS disk, perform follow-up operations for the additional capacity.

NOTE

If the instance to which the EVS disk is attached is stopped during capacity expansion, power on the instance before capacity expansion.

- For Windows, see [Operations After Expanding Disk Capacity in Windows](#).
- For Linux, see [Table 3-17](#).

3.7.3 Expanding Disk Capacity Offline

This section describes how to expand the capacity of an EVS disk not attached to any instance.

Restrictions

- The capacity of a disk configured with the DR service (CSHA/CSDR/VHA) cannot be expanded.
- When the backend storage is Huawei SAN storage (OceanStor V3/V5 series or OceanStor Dorado V3 series/6.x) or heterogeneous storage, if the EVS disk has snapshots, capacity expansion is not supported. When the backend storage is Huawei Distributed Block Storage, capacity expansion can be performed for an EVS disk with snapshots.
- Capacity expansion cannot be performed when the disk status is **Reserved** or **Maintenance**.

Procedure

Step 1 [Log in to EVS Console as a VDC administrator or VDC operator](#).

Step 2 In the EVS disk list, locate the row that contains the target disk, click **More** in the **Operation** column, and choose **Expand Capacity**.

The **Expand Capacity** page is displayed.

NOTE

To view information about the instance to which the EVS disk is attached, click the instance name in the **Attaching Information** column.

Step 3 Set **Added Capacity (GB)** as prompted and click **Next**.

Step 4 On the **Resource Details** page, confirm the EVS disk specifications.

- If you do not need to modify the specifications, click **Apply Now** to start the EVS disk capacity expansion.
- If you need to modify the specifications, click **Previous** to modify parameters.

After the expansion is submitted, the disk list page is displayed.

Step 5 If the EVS disk whose capacity is to be expanded requires approval, contact the administrator for approval. Otherwise, skip this step.

On the **Elastic Volume Service** page, view the capacity of the EVS disk. If the disk capacity has increased, the expansion is successful.

----End

Follow-up Procedure

After you have expanded the capacity of an EVS disk, perform follow-up operations on the additional volume.

1. Attach the expanded EVS disk to an instance. For details, see section [Attaching an EVS Disk](#).

2. Operations after attaching:
 - For Windows, see [Operations After Expanding Disk Capacity in Windows](#).
 - For Linux, see [Table 1 Post-expansion operations for an EVS disk](#).

3.7.4 Operations After Expanding Disk Capacity in Windows

This section describes how to perform follow-up operations for the additional disk space after you have expanded the capacity of an EVS disk and attached the EVS disk to an instance. The method for allocating the extended space of an EVS disk varies depending on the in-use server OS. This section uses Windows Server 2008 R2 Enterprise as the example OS to describe how to perform post-expansion operations for an EVS disk. For other Windows OSs, see the corresponding OS product documents.

Prerequisites

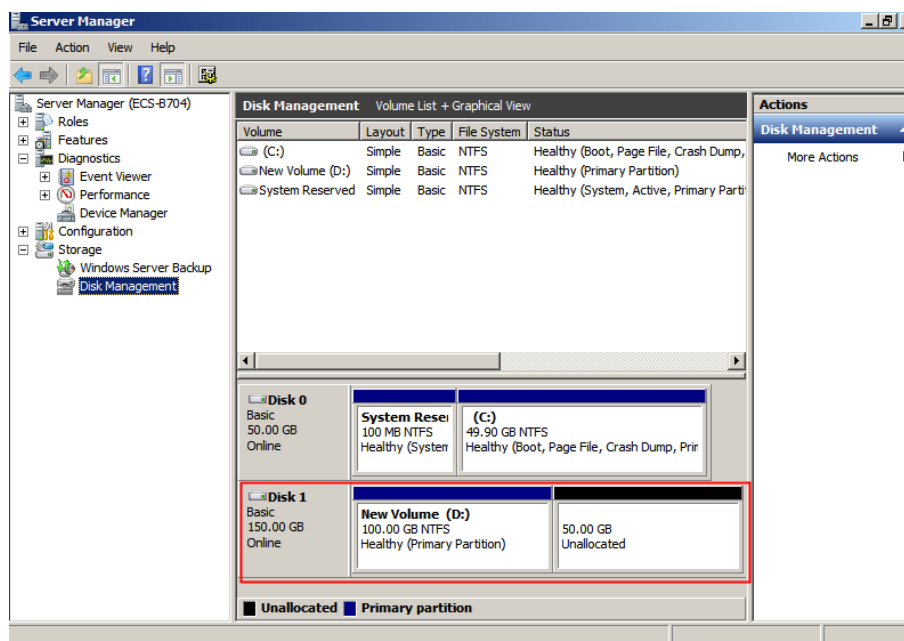
- For details about how to log in to an ECS, see [Operation Help Center > Compute > Elastic Cloud Server > User Guide > Logging In to an ECS](#).
- A disk has been attached to an instance and has not been initialized.

Context

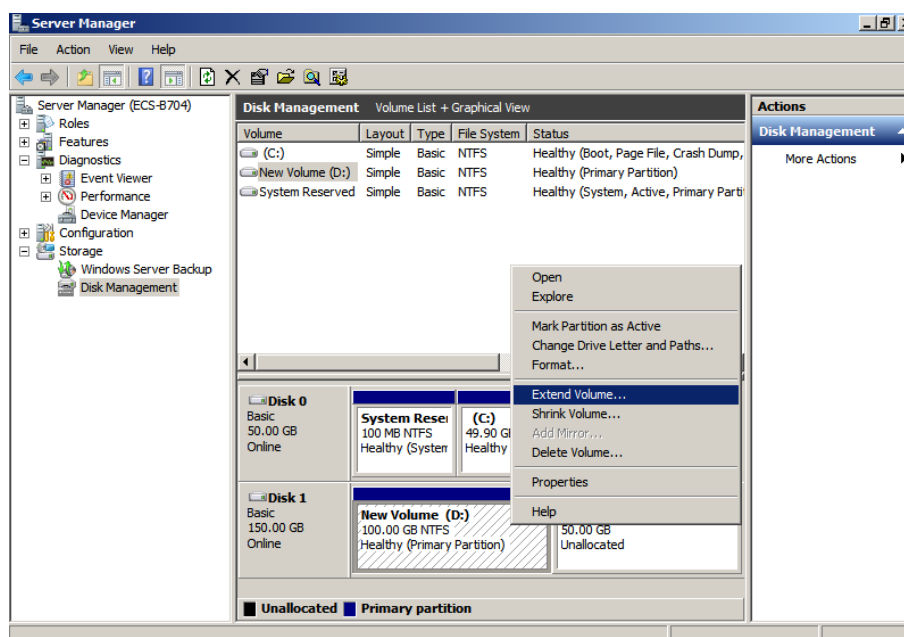
After the capacity expansion has succeeded, the new EVS disk space needs to be allocated to an existing partition or a new partition. This section describes how to allocate new EVS disk space to an existing partition.

Procedure

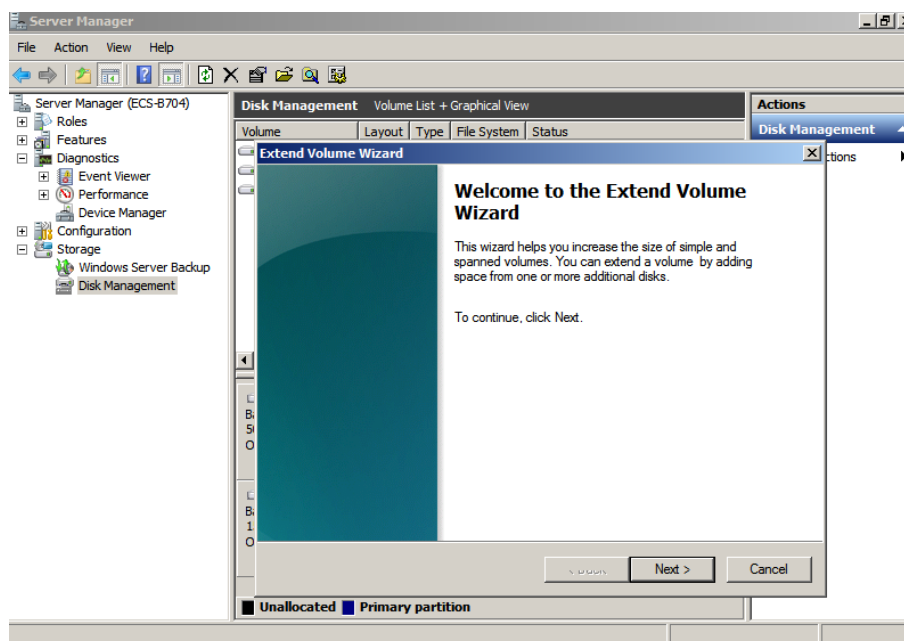
- Step 1** In desktop, right-click **Computer** and choose **Manage** from the shortcut menu.
The **Server Manager** window is displayed.
- Step 2** In the navigation tree, choose **Storage > Disk Management**.
The **Disk Management** page is displayed.
- Step 3** On the **Disk Management** page, select the disk and partition that needs to be extended. The current partition size and unallocated disk space are displayed.



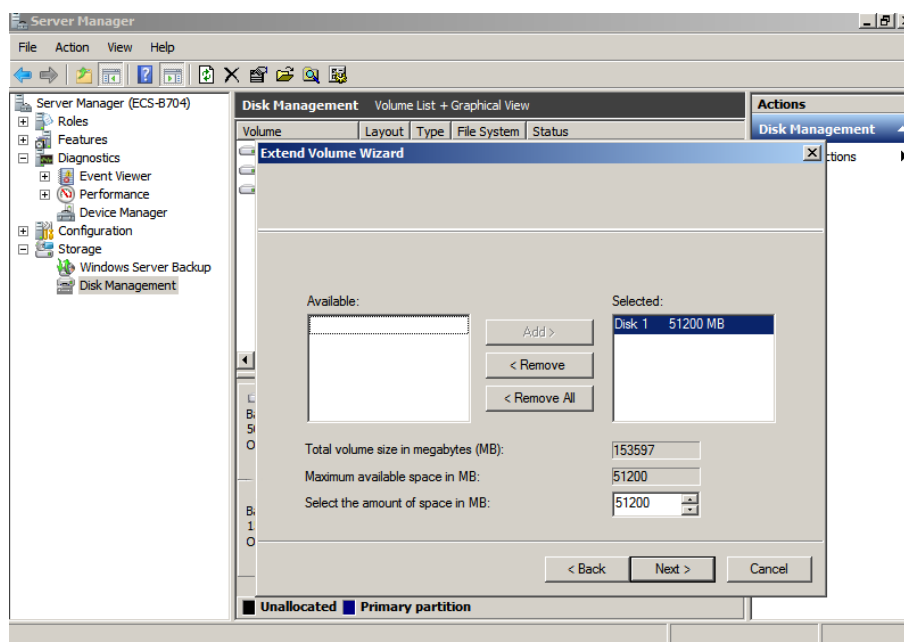
Step 4 Right-click the selected disk and choose **Extend Volume**.



Step 5 On the displayed **Extend Volume Wizard** page, click **Next**.

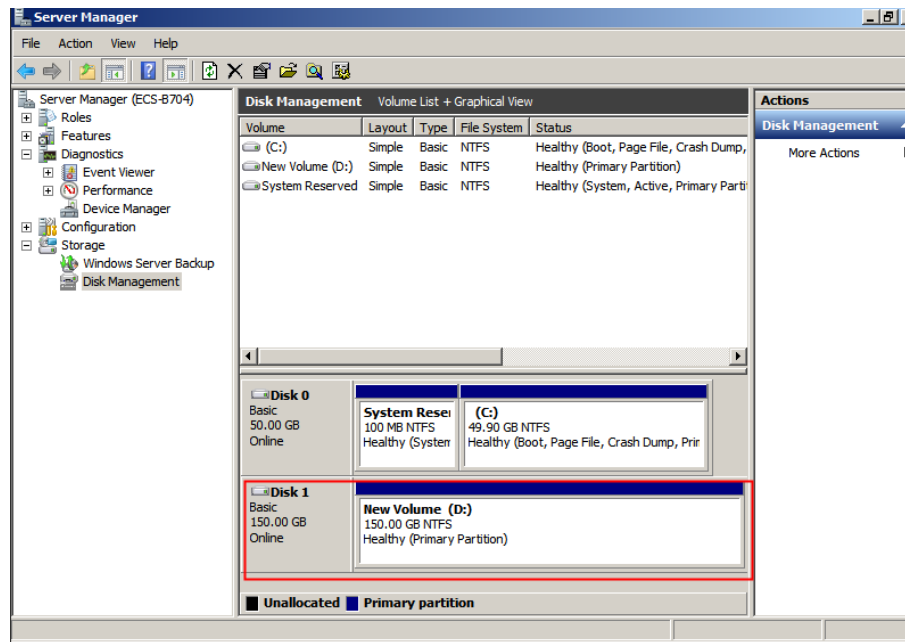


Step 6 In the text box to the right of **Select the amount of space in MB**, enter the extended capacity and click **Next**.



Step 7 Click **Finish** to complete the wizard.

After the expansion is successful, the disk capacity is greater than the original capacity.



----End

3.7.5 Operations After Expanding Disk Capacity in Linux (Adding Partitions Using fdisk)

After the capacity of an EVS disk is expanded and the disk is attached to an instance, the additional space needs to be allocated to a new partition. This section uses CentOS 7.0 (64-bit) as an example to describe how to use fdisk to create partitions for a disk after capacity expansion.

Prerequisites

- For details about how to log in to an ECS, see **Operation Help Center > Compute > Elastic Cloud Server > User Guide > Logging In to an ECS**.
- A disk has been attached to an instance and has not been initialized.

Procedure

The steps below describe how to allocate the additional space to a new partition and mount the partition on **/opt**. The disk is partitioned using MBR, the disk capacity is less than 2 TB, and the number of partitions is lower than the upper limit. You can use either the fdisk or parted tool to partition the disk space. This section uses the fdisk tool as an example.

Step 1 Run the following command to query and view the disk information:

fdisk -l

Information similar to the following is displayed: (**/dev/xvda** is the system disk.)

```
[root@ecs-bab9 test]# fdisk -l
```

```
Disk /dev/xvda: 64.4 GB, 64424509440 bytes, 125829120 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
```

I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x000cc4ad

Device	Boot	Start	End	Blocks	Id	System
/dev/xvda1	*	2048	2050047	1024000	83	Linux
/dev/xvda2		2050048	22530047	10240000	83	Linux
/dev/xvda3		22530048	24578047	1024000	83	Linux
/dev/xvda4		24578048	83886079	29654016	5	Extended
/dev/xvda5		24580096	26628095	1024000	82	Linux swap / Solaris

NOTE

The capacity displayed here is inconsistent with the capacity of the EVS disk applied for on ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios). The reason is as follows: The capacity of EVS disks is calculated using the unit of GiB (Gibibyte), while the capacity unit in Linux OS is GB (Gigabyte). The GiB is calculated in binary mode, and the GB is calculated in decimal format. 1 GiB = 1,073,741,824 bytes and 1 GB = 1,000,000,000 bytes.

Step 2 Run the following command to enter the fdisk (/dev/xvda is used as an example):

fdisk /dev/xvda

Information similar to the following is displayed:

```
[root@ecs-bab9 test]# fdisk /dev/xvda
Welcome to fdisk (util-linux 2.23.2).
```

Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.

Command (m for help):

Step 3 Enter **n** and press **Enter** to create a partition.

Because the system disk has five existing partitions, the system automatically creates the sixth one.

Information similar to the following is displayed:

```
Command (m for help): n
All primary partitions are in use
Adding logical partition 6
First sector (26630144-83886079, default 26630144):
```

Step 4 Enter the new partition's first sector number, for example the default value, and press **Enter**.

The first sector number must be greater than the last sector numbers of existing partitions.

Information similar to the following is displayed:

```
First sector (26630144-83886079, default 26630144):
Using default value 26630144
Last sector, +sectors or +size{K,M,G} (26630144-83886079, default 83886079):
```

Step 5 Enter the new partition's last sector number and press **Enter**.

The default last sector number is used as an example.

Information similar to the following is displayed:

```
Last sector, +sectors or +size{K,M,G} (26630144-83886079, default 83886079):
Using default value 83886079
Partition 6 of type Linux and of size 27.3 GiB is set
```

Command (m for help):

Step 6 Enter **p** and press **Enter** to view the created partition.

Information similar to the following is displayed:

```
Disk /dev/xvda: 64.4 GB, 64424509440 bytes, 125829120 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x000cc4ad

   Device Boot      Start         End      Blocks   Id  System
/dev/xvda1 *        2048       2050047       1024000   83  Linux
/dev/xvda2          2050048       22530047       10240000   83  Linux
/dev/xvda3          22530048       24578047        1024000   83  Linux
/dev/xvda4          24578048       83886079       29654016    5  Extended
/dev/xvda5          24580096       26628095        1024000   82  Linux swap / Solaris
/dev/xvda6          26630144       83886079       28627968   83  Linux
```

Command (m for help):

Step 7 Enter **w** and press **Enter**.

The partition result is written into the partition table, and the partitioning is complete.

Information similar to the following is displayed:

```
Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.

WARNING: Re-reading the partition table failed with error 16: Device or resource busy.
The kernel still uses the old table. The new table will be used at
the next reboot or after you run partprobe(8) or kpartx(8)
Syncing disks.
```

Step 8 Run the following command to synchronize the new partition table to the OS:

partprobe

Step 9 Run the following command to set the format for the file system of the new partition:

In this example, the ext4 file system is used.

mkfs -t ext4 /dev/xvda6

Information similar to the following is displayed:

```
[root@ecs-bab9 test]# mkfs -t ext4 /dev/xvda6
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
1790544 inodes, 7156992 blocks
357849 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2155872256
219 block groups
32768 blocks per group, 32768 fragments per group
8176 inodes per group
```

Superblock backups stored on blocks:

32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
4096000

Allocating group tables: done

Writing inode tables: done

Creating journal (32768 blocks): done

Writing superblocks and filesystem accounting information: done

The formatting takes a while, and you need to observe the system running status. Once **done** is displayed in the command output, the formatting is complete.

Step 10 Run the following command to mount the new partition to the space-demanding directory, for example **/opt**:

mount /dev/xvda6 /opt

Information similar to the following is displayed:

```
[root@ecs-bab9 test]# mount /dev/xvda6 /opt
[root@ecs-bab9 test]#
```

NOTE

If the new partition is mounted to a directory that is not empty, the subdirectories and files in the directory will be hidden. In this situation, you are advised to mount the new partition to an empty directory or a new directory. If the new partition must be mounted to a directory that is not empty, move the subdirectories and files in this directory to another directory temporarily. After the partition is mounted, move the subdirectories and files back.

Step 11 Run the following command to view the mount result:

df -TH

Information similar to the following is displayed:

```
[root@ecs-bab9 test]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/xvda2      xfs       11G   7.4G  3.2G  71% /
devtmpfs       devtmpfs  4.1G   0  4.1G   0% /dev
tmpfs          tmpfs     4.1G  82k  4.1G   1% /dev/shm
tmpfs          tmpfs     4.1G  9.2M  4.1G   1% /run
tmpfs          tmpfs     4.1G   0  4.1G   0% /sys/fs/cgroup
/dev/xvda3      xfs       1.1G   39M  1.1G   4% /home
/dev/xvda1      xfs       1.1G  131M  915M  13% /boot
/dev/xvda6      ext4      29G   47M  28G   1% /opt
```

----End

Setting Automatic Disk Mounting at a System Start

If you require a disk to be automatically attached to an instance when the instance is started, enable automatic disk attachment upon an instance start by referring to operations provided in this section. When enabling automatic disk attachment, you cannot directly specify **/dev/xvdb1** in **/etc/fstab**. This is because the sequence codes of the instance may change during an instance stop or start process. You are advised to use the universally unique identifier (UUID) in **/etc/fstab** to automatically attach the disk at a system start.

NOTE

The UUID of a disk is a character string that uniquely identifies a storage device in a Linux system.

Step 1 Run the following command to query the partition UUID:

blkid *Disk partition*

For example, run the following command to query the UUID of **/dev/xvdb1**:

blkid /dev/xvdb1

Information similar to the following is displayed: (The UUID of **/dev/xvdb1** is displayed.)

```
[root@ecs-b656 test]# blkid /dev/xvdb1
/dev/xvdb1: UUID="1851e23f-1c57-40ab-86bb-5fc5fc606ffa" TYPE="ext4"
```

Step 2 Run the following command to open the **fstab** file using the vi editor:

vi /etc/fstab

Step 3 Press **i** to enter the editing mode.

Step 4 Move the cursor to the end of the file and press **Enter**. Then add the following information:

UUID=xxx attachment directory file system defaults 0 2

Assuming that the file system is **ext4** and the attachment directory is **/mnt/sdc**.

```
UUID=1851e23f-1c57-40ab-86bb-5fc5fc606ffa /mnt/sdc ext4 defaults 0 2
```

NOTICE

After automatic attachment upon instance start is configured, comment out or delete the line in the **fstab** file before detaching the disk. Otherwise, you may fail to access the OS after the disk is detached.

Step 5 Press **Esc**, enter **:wq**, and press **Enter**.

The system saves the configuration and exits the vi editor.

----End

3.7.6 Operations After Expanding Disk Capacity in Linux (Adding Partitions Using parted)

This section describes how to perform follow-up operations for the additional disk space after you have expanded the capacity of an EVS disk and attached the EVS disk to an instance. This section uses EulerOS 2.0 (64-bit) to describe how to allocate the additional disk space to a partition using parted.

Prerequisites

- For details about how to log in to an ECS, see **Operation Help Center > Compute > Elastic Cloud Server > User Guide > Logging In to an ECS**.
- A disk has been attached to an instance and has not been initialized.

Procedure

The steps below describe how to allocate the additional space to a new partition and mount the partition on **/opt**. The disk is partitioned using MBR, the disk

capacity is less than 2 TB, and the number of partitions is lower than the upper limit. You can use either the `fdisk` or `parted` tool to partition the disk space. This section uses the `parted` tool as an example.

Step 1 Run the following command to view the disk partition information:

lsblk

If the following information is displayed, the total capacity of the current system disk **dev/xvda** is 80 GB. 40 GB of the disk has been allocated to partitions, and the remaining 40 GB is additional disk space and have not been allocated to any partition.

```
[root@ecs-1120 linux]# lsblk
NAME        MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
xvda  202:0    0  80G  0 disk
├─xvda1 202:1    0  40G  0 part /
├─xvdb  202:16   0  250G  0 disk
│   └─xvdb1 202:17   0  100G  0 part
│       └─xvdb2 202:18   0   50G  0 part
└─xvdc  202:32   0   40G  0 disk
    ├─xvdc1 202:33   0    8G  0 part
    └─xvdc2 202:34   0   32G  0 part
```

Step 2 Run the following command to create a partition for the additional system disk space:

parted *System disk*

parted **/dev/xvda**

The following information is displayed:

```
[root@ecs-1120 linux]# parted /dev/xvda
GNU Parted 3.1
Using /dev/xvda
Welcome to GNU Parted! Type 'help' to view a list of commands.
```

Step 3 Enter **unit s** and press **Enter** to set the measurement unit of the disk to sector numbers.

Step 4 Enter **p** and press **Enter** to view the current disk partition style.

Partition Table specifies the partition style for existing disks. **msdos** indicates that the disk partition style is MBR, and **gpt** indicates that the disk partition style is GPT.

The following information is displayed:

```
(parted) unit s
(parted) p
Model: Xen Virtual Block Device (xvd)
Disk /dev/xvda: 167772160s
Sector size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:

Number Start End      Size      Type     File system  Flags
 1      2048s 83886079s 83884032s primary ext4
```

Step 5 Enter **mkpart** and press **Enter** to create a partition.

Step 6 Enter **p** and press **Enter** to create a primary partition. Creating a primary partition is used as an example.

Set the file system format to **ext4**.

Set the first sector and last sector to **83886080** and **167772159** for the **dev/xvda2** partition, respectively.

Set the parameters as required.

The following information is displayed:

```
(parted) mkpart
Partition type? primary/extended? p
File system type? [ext2]? ext4
Start? 83886080
End? 167772159
```

NOTE

The file system type may fail to be configured in this step. You can reconfigure the file system format according to [Step 9](#).

Step 7 Enter **p** and press **Enter** to view the created partition.

Information similar to the following is displayed:

```
(parted) p
Model: Xen Virtual Block Device (xvd)
Disk /dev/xvda: 167772160s
Sector size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:

Number Start      End          Size         Type      File system  Flags
 1      2048s    83886079s   83884032s   primary  ext4
 2     83886080s 167772159s  83886080s   primary
```

Step 8 Enter **q** and press **Enter** to exit parted.

Step 9 Run the following command to set the format for the file system of the new partition:

In this example, the ext4 file system is used.

mkfs -t ext4 /dev/xvda2

Information similar to the following is displayed:

Wait for the formatting. If **done** is displayed in the command output, the formatting has been complete.

```
[[root@ecs-1120 linux]# mkfs -t ext4 /dev/xvda2
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
2621440 inodes, 10485760 blocks
524288 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2157969408
320 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
?32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
?4096000, 7962624

Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

Step 10 Run the following command to mount the new partition to the space-demanding directory, for example **/opt**:

```
mount /dev/xvda2 /opt
```

 **NOTE**

If the new partition is mounted to a directory that is not empty, the subdirectories and files in the directory will be hidden. In this situation, you are advised to mount the new partition to an empty directory or a new directory. If the new partition must be mounted to a directory that is not empty, move the subdirectories and files in this directory to another directory temporarily. After the partition is mounted, move the subdirectories and files back.

Step 11 Run the following command to view the mount result:

```
df -TH
```

Information similar to the following is displayed:

```
[root@ecs-1120 linux]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/xvda1      ext4      43G   8.3G   33G   21% /
devtmpfs        devtmpfs  885M    0 885M    0% /dev
tmpfs           tmpfs     894M    0 894M    0% /dev/shm
tmpfs           tmpfs     894M   18M 877M    2% /run
tmpfs           tmpfs     894M    0 894M    0% /sys/fs/cgroup
tmpfs           tmpfs     179M    0 179M    0% /run/user/2000
tmpfs           tmpfs     179M    0 179M    0% /run/user/0
tmpfs           tmpfs     179M    0 179M    0% /run/user/1001
/dev/xvda2      ext4      43G    51M   40G    1% /opt
```

----End

Setting Automatic Disk Mounting at a System Start

If you require a disk to be automatically attached to an instance when the instance is started, enable automatic disk attachment upon an instance start by referring to operations provided in this section. When enabling automatic disk attachment, you cannot directly specify **/dev/xvdb1** in **/etc/fstab**. This is because the sequence codes of the instance may change during an instance stop or start process. You are advised to use the universally unique identifier (UUID) in **/etc/fstab** to automatically attach the disk at a system start.

 **NOTE**

The UUID of a disk is a character string that uniquely identifies a storage device in a Linux system.

Step 1 Run the following command to query the partition UUID:

```
blkid Disk partition
```

For example, run the following command to query the UUID of **/dev/xvdb1**:

```
blkid /dev/xvdb1
```

Information similar to the following is displayed: (The UUID of **/dev/xvdb1** is displayed.)

```
[root@ecs-b656 test]# blkid /dev/xvdb1
/dev/xvdb1: UUID="1851e23f-1c57-40ab-86bb-5fc5fc606ffa" TYPE="ext4"
```

Step 2 Run the following command to open the **fstab** file using the vi editor:

```
vi /etc/fstab
```

Step 3 Press **i** to enter the editing mode.

Step 4 Move the cursor to the end of the file and press **Enter**. Then add the following information:

```
UUID=xxx attachment directory file system defaults 0 2
```

Assuming that the file system is **ext4** and the attachment directory is **/mnt/sdc**.

```
UUID=1851e23f-1c57-40ab-86bb-5fc5fc606ffa /mnt/sdc ext4 defaults 0 2
```

NOTICE

After automatic attachment upon instance start is configured, comment out or delete the line in the **fstab** file before detaching the disk. Otherwise, you may fail to access the OS after the disk is detached.

Step 5 Press **Esc**, enter **:wq**, and press **Enter**.

The system saves the configuration and exits the vi editor.

----End

3.7.7 Operations After Expanding Disk Capacity in Linux (Replacing Original Partitions Using fdisk)

This section describes how to perform follow-up operations for the additional disk space after you have expanded the capacity of an EVS disk and attached the EVS disk to an instance.

NOTICE

For details, contact the OS vendor of the target VM. This section uses EulerOS 2.0 (64-bit) to describe how to allocate the additional disk space to a partition using fdisk.

Prerequisites

- For details about how to log in to an ECS, see **Operation Help Center > Compute > Elastic Cloud Server > User Guide > Logging In to an ECS**.
- A disk has been attached to an instance and has not been initialized.

Procedure

The following example shows how to create new partitions to replace existing ones on a disk that has been attached to an instance. In this example, the disk partition is **/dev/xvdb1**, the partition style is MBR, and the mounting directory is **/mnt/sdc**. After capacity expansion, the disk capacity is less than 2 TB. You can use either the fdisk or parted tool to partition the disk space. This section uses the fdisk tool as an example.

Step 1 Run the following command to view the disk partition information:

fdisk -l

Information similar to the following is displayed: (In the command output, the server contains two disks. **/dev/xvda** is the system disk, and **/dev/xvdb** is the data disk.)

```
[root@ecs-b656 test]# fdisk -l
```

```

Disk /dev/xvda: 42.9 GB, 42949672960 bytes, 83886080 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x000cc4ad

```

Device	Boot	Start	End	Blocks	Id	System
/dev/xvda1	*	2048	2050047	1024000	83	Linux
/dev/xvda2		2050048	22530047	10240000	83	Linux
/dev/xvda3		22530048	24578047	1024000	83	Linux
/dev/xvda4		24578048	83886079	29654016	5	Extended
/dev/xvda5		24580096	26628095	1024000	82	Linux swap / Solaris

```

Disk /dev/xvdb: 24.7 GB, 24696061952 bytes, 48234496 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0xb00005bd

```

Device	Boot	Start	End	Blocks	Id	System
/dev/xvdb1		2048	20971519	10484736	83	Linux

In the command output, parameter **Disk label type** indicates the disk partition style. Value **dos** indicates the MBR partition style, and value **gpt** indicates the GPT partition style.

View the **/dev/xvdb** capacity and check whether the additional space is included.

NOTE

The capacity displayed here is inconsistent with the capacity of the EVS disk applied for on ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios). The reason is as follows: The capacity of EVS disks is calculated using the unit of GiB (Gibibyte), while the capacity unit in Linux OS is GB (Gigabyte). The GiB is calculated in binary mode, and the GB is calculated in decimal format. 1 GiB = 1,073,741,824 Bytes and 1 GB = 1,000,000,000 Bytes.

- If the additional space is displayed, go to [Step 2](#).
- If the new capacity is not displayed in the command output, run the following command to update the disk capacity:

```
echo 1 > /sys/class/scsi_device/%d:%d:%d:%d/device/rescan &
```

In the command, **%d:%d:%d:%d** indicates a folder in the **/sys/class/scsi_device/** directory and can be obtained using **ll /sys/class/scsi_device/**.

Information similar to the following is displayed: (**2:0:0:0** indicates the folder to be obtained.)

```

cs-xen-02:/sys/class/scsi_device # ll /sys/class/scsi_device/
total 0
lrwxrwxrwx 1 root root 0 Sep 26 11:37 2:0:0:0-> ../../devices/xen/vscsi-2064/host2/target2:0:0/2:0:0:0/scsi_device/2:0:0:0

```

Example command:

```
echo 1 > /sys/class/scsi_device/2:0:0:0/device/rescan &
```

After the update is complete, run the **fdisk -l** command to view the disk partition information.

Step 2 Run the following command to unmount the disk:

```
umount /mnt/sdc
```

Step 3 Run the following command and enter **d** and then **1** to delete existing partition **/dev/xvdb1**. The partition to be deleted varies between scenarios.

```
fdisk /dev/xvdb
```

The command output is as follows:

```
[root@ecs-b656 test]# fdisk /dev/xvdb
Welcome to fdisk (util-linux 2.23.2).

Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.

Command (m for help): d
Selected partition 1
Partition 1 is deleted

Command (m for help):
```

NOTE

Deleting partitions will not cause data losses on the data disk.

Step 4 Enter **n** and press **Enter** to create a partition.

Entering **n** creates a partition.

There are two types of disk partitions:

- Choosing **p** creates a primary partition.
- Choosing **e** creates an extended partition.

```
Command (m for help): n
Partition type:
  p   primary (0 primary, 0 extended, 4 free)
  e   extended
```

Step 5 Ensure that the entered partition type is the same as that of the partition to be replaced. In this example, creating a primary partition is used. Therefore, enter **p** and press **Enter** to create a primary partition.

Information similar to the following is displayed: (**Partition number** indicates the serial number of the primary partition. The value can be **1** to **4**.)

```
Select (default p): p
Partition number (1-4, default 1):
```

Step 6 Ensure that entered partition number is the same as that of the partition to be replaced. In this example, the partition number **1** is used. Therefore, enter **1** and press **Enter**.

Information similar to the following is displayed: (**First sector** indicates the first sector number. The value can be **2048** to **20971519**, and the default value is **2048**.)

```
Partition number (1-4, default 1): 1
First sector (2048-20971519, default 2048):
```

Step 7 Press Enter.

The default start sector number is used as an example.

Information similar to the following is displayed: (**Last sector** indicates the last sector number. The value can be **2048** to **20971519**, and the default value is **20971519**.)

```
First sector (2048-20971519, default 2048):  
Using default value 2048  
Last sector, +sectors or +size{K,M,G} (2048-20971519, default 20971519):
```

Step 8 Press Enter.

The default last sector number is used as an example.

Information similar to the following is displayed, indicating that the partition is successfully created.

```
Last sector, +sectors or +size{K,M,G} (2048-20971519, default 20971519):  
Using default value 20971519  
Partition 1 of type Linux and of size 10 GiB is set  
Command (m for help):
```

Step 9 Enter p and press Enter to view the details about the created partition.

Information similar to the following is displayed: (Details about the **/dev/xvdb1** partition are displayed.)

```
Command (m for help): p  
  
Disk /dev/xvdb: 10.7 GB, 10737418240 bytes, 20971520 sectors  
Units = sectors of 1 * 512 = 512 bytes  
Sector size (logical/physical): 512 bytes / 512 bytes  
I/O size (minimum/optimal): 512 bytes / 512 bytes  
Disk label type: dos  
Disk identifier: 0xb00005bd
```

Device	Boot	Start	End	Blocks	Id	System
/dev/xvdb1		2048	20971519	10484736	83	Linux

```
Command (m for help):
```

Step 10 Enter w and press Enter to write the changes into the partition table.

Information similar to the following is displayed: (The partition is successfully created.)

```
Command (m for help): w  
The partition table has been altered!  
  
Calling ioctl() to re-read partition table.  
Syncing disks.
```

If the following error is displayed when you write partition results to the partition table, the new partition table will be updated upon the next OS restart.

```
Command (m for help): w  
The partition table has been altered!
```

```
Calling ioctl() to re-read partition table.
```

```
WARNING: Re-reading the partition table failed with error 16: Device or resource busy.  
The kernel still uses the old table, The new table will be used at the next reboot or after you run  
partprobe(8) or kpartx(8)  
Syncing disks.
```

Step 11 Run the following commands to check and adjust the size of the file system on `/dev/xvdb1`:

e2fsck -f /dev/xvdb1

Information similar to the following is displayed:

```
[root@ecs-b656 test]# e2fsck -f /dev/xvdb1
e2fsck 1.42.9 (28-Dec-2013)
Pass 1: Checking inodes, blocks, and sizes
Pass 2: Checking directory structure
Pass 3: Checking directory connectivity
Pass 4: Checking reference counts
Pass 5: Checking group summary information
/dev/xvdb1: 11/655360 files (0.0% non-contiguous), 83137/2621184 blocks
```

resize2fs /dev/xvdb1

Information similar to the following is displayed:

```
[root@ecs-b656 test]# resize2fs /dev/xvdb1
resize2fs 1.42.9 (28-Dec-2013)
Resizing the filesystem on /dev/xvdb1 to 6029056 (4k) blocks.
The filesystem on /dev/xvdb1 is now 6029056 blocks long.
```

----End

3.7.8 Operations After Expanding Disk Capacity in Linux (Replacing Original Partitions Using parted)

This section describes how to perform follow-up operations for the additional disk space after you have expanded the capacity of an EVS disk and attached the EVS disk to an instance.

NOTICE

For details, contact the OS vendor of the target VM. This section uses EulerOS 2.0 (64-bit) to describe how to allocate the additional disk space to a partition using parted.

Prerequisites

- For details about how to log in to an ECS, see **Operation Help Center > Compute > Elastic Cloud Server > User Guide > Logging In to an ECS**.
- A disk has been attached to an instance and has not been initialized.

Procedure

The following example shows how to create a partition to replace the `/dev/xvdc1` partition mounted on `/mnt/sdc`. `/dev/xvdc1` is the only partition of the `/dev/xvdc` disk attached to an instance and uses the GPT partition style. There are two disks attached to the instance. During the partition creation, services will be interrupted.

NOTICE

After the disk capacity has been expanded, the additional space is added to the end of the disk. When the disk has multiple partitions, only the partition at the end of the disk can be expanded.

Step 1 Run the following command to view the disk partition information:

lsblk

Information similar to the following is displayed:

```
[root@ecs-1120 sdc]# lsblk
NAME        MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
xvda        202:0    0   80G  0 disk
├─xvda1     202:1    0   40G  0 part /
├─xvda2     202:2    0   40G  0 part /opt
xvdb        202:16   0  350G  0 disk
├─xvdb1     202:17   0  100G  0 part
├─xvdb2     202:18   0  200G  0 part
xvdc        202:32   0   60G  0 disk
└─xvdc1     202:33   0   10G  0 part /mnt/sdc
```

Indicates that the total capacity of the current data disk **/dev/xvdc** is 60 GB and the allocated partition capacity is 10 GB. The last partition is **/dev/xvdc1**, which is attached to the **/mnt/sdc** directory.

View the **/dev/xvdc** capacity and check whether the additional space is included.

- If the additional space is included, go to [Step 2](#).
- If the new capacity is not displayed in the command output, run the following command to update the disk capacity:

echo 1 > /sys/class/scsi_device/%d:%d:%d:%d/device/rescan &

In the command, **%d:%d:%d:%d** indicates a folder in the **/sys/class/scsi_device/** directory and can be obtained using **ll /sys/class/scsi_device/**.

Information similar to the following is displayed: (**2:0:0:0** indicates the folder to be obtained.)

```
cs-xen-02:/sys/class/scsi_device # ll /sys/class/scsi_device/
total 0
lrwxrwxrwx 1 root root 0 Sep 26 11:37 2:0:0:0-> ../../devices/xen/vscsi-2064/host2/target2:0:0/2:0:0:0/scsi_device/2:0:0:0
```

Example command:

echo 1 > /sys/class/scsi_device/2:0:0:0/device/rescan &

After the update is complete, run the **fdisk -l** command to view the disk partition information.

Step 2 Run the following command to unmount the disk partition:

umount /mnt/sdc

Step 3 Run the following command to view the unmount result:

lsblk

Information similar to the following is displayed:

```
[root@ecs-1120 linux]# umount /mnt/sdc
[root@ecs-1120 linux]# lsblk
NAME        MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
xvda        202:0    0   80G  0 disk
├─xvda1     202:1    0   40G  0 part /
```

```
└─xvda2 202:2 0 40G 0 part /opt
xvdb 202:16 0 350G 0 disk
└─xvdb1 202:17 0 100G 0 part
└─xvdb2 202:18 0 200G 0 part
xvdc 202:32 0 60G 0 disk
└─xvdc1 202:33 0 10G 0 part
```

Step 4 Run the following command to enter parted to allocate the additional space of the data disk to a partition:

parted *Data disk*

In this example, **/dev/xvdc** is the data disk.

parted /dev/xvdc

Information similar to the following is displayed:

```
[root@ecs-1120 linux]# parted /dev/xvdc
GNU Parted 3.1
Using /dev/xvdc
Welcome to GNU Parted! Type 'help' to view a list of commands.
```

Step 5 Enter **unit s** and press **Enter** to set the measurement unit of the disk to sector numbers.

Step 6 Enter **p** and press **Enter** to view the current disk partition style.

Information similar to the following is displayed:

```
(parted) unit s
(parted) p
Error: The backup GPT table is not at the end of the disk, as it should be.
This might mean that another operating system believes the disk is smaller.
Fix, by moving the backup to the end (and removing the old backup)?
Fix/Ignore/Cancel? Fix
Warning: Not all of the space available to /dev/xvdb appears to be used,
you can fix the GPT to use all of the space (an extra 104857600 blocks)
or continue with the current setting?
Fix/Ignore? Fix
Model: Xen Virtual Block Device (xvd)
Disk /dev/xvdc: 125829120s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number Start      End          Size         File system  Name  Flags
 1      2048s      20969471s   20967424s    ext4         opt
```

Partition Table specifies the partition style for existing disks. **msdos** indicates that the disk partition style is MBR, and **gpt** indicates that the disk partition style is GPT.

If the preceding information is displayed, enter **Fix** to rectify the disk exception. Then take note of the first and last sectors of the **/dev/xvdc1** partition. These values will be used during the partition recreation. In this example, the partition's first sector is **2048**, and its last sector is **20969471**.

The **/dev/xvdc1** partition number is **1**. Therefore, enter **rm 1** and press **Enter** to delete the partition.

Step 7 Enter **p** and press **Enter** to check whether the **/dev/xvdc1** partition has been deleted.

```
(parted) rm 1
(parted) p
Model: Xen Virtual Block Device (xvd)
Disk /dev/xvdc: 125829120s
```

```
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start   End     Size    File system  Name  Flags
```

Step 8 Enter **mkpart opt 2048s 100%** and press **Enter** to create a partition.

opt indicates the name of the created partition, **2048s** indicates the first sector recorded in [Step 6](#), and **100%** indicates that all capacity of the disk is allocated to one partition, which must be greater than or equal to the last sector recorded in [Step 6](#).

Information similar to the following is displayed:

```
(parted) mkpart opt 2048s 100%
Warning: You requested a partition from 2048s to 125829199s (sectors 2048..125829199).
The closest location we can manage is 2048s to 125829036s (sectors 2048..125829036).
Is this still acceptable to you?
Yes/No? Yes
```

Enter **Yes** as prompted to set the last sector.

If the following warning message is displayed, enter **Cancel** to stop the partitioning. Then, find the first sector with the best disk performance and use that value to partition the disk. The warning message will not be displayed if the first sector with the best disk performance has been entered. In this example, **2048s** is one of such first sectors. Therefore, the system does not display the warning message.

```
Warning: The resulting partition is not properly aligned for best performance.
Ignore/Cancel? Cancel
```

NOTE

Data will be lost if the following operations are performed:

- Select a first sector which is inconsistent with that of the original partition.
- Select a last sector which is smaller than that of the original partition.

Step 9 Enter **p** and press **Enter** to check whether the **/dev/xvdc1** partition has been recreated.

```
(parted) p
Model: Xen Virtual Block Device (xvd)
Disk /dev/xvdb: 125829120s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start   End     Size    File system  Name  Flags
1       2048s   125829086s 125827039s ext4      opt
```

The **/dev/xvdc1** partition has been recreated.

Step 10 Enter **q** and press **Enter** to exit parted.

Step 11 Run the following command to view the disk partition information after the partition expansion:

lsblk

If information similar to the following is displayed, the total capacity of **/dev/xvdc** is 60 GB. The new 50 GB of the capacity is allocated to the **/dev/xvdc1** partition, and the partition is mounted on the **/mnt/sdc** directory. Skip [Step 12](#), [Step 14](#), and [Step 15](#).

```
[root@ecs-1120 sdc]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
xvda 202:0 0 80G 0 disk
├─xvda1 202:1 0 40G 0 part /
└─xvda2 202:2 0 40G 0 part /opt
xvdb 202:16 0 350G 0 disk
├─xvdb1 202:17 0 100G 0 part
└─xvdb2 202:18 0 200G 0 part
xvdc 202:32 0 60G 0 disk
└─xvdc1 202:33 0 60G 0 part /mnt/sdc
```

- Step 12** Run the following command to check the correctness of the file system on **/dev/xvdc1**:

e2fsck -f /dev/xvdc1

Information similar to the following is displayed:

```
[root@ecs-1120 linux]# e2fsck -f /dev/xvdb2
e2fsck 1.42.9 (28-Dec-2013)
Pass 1: Checking inodes, blocks, and sizes
Pass 2: Checking directory structure
Pass 3: Checking directory connectivity
Pass 4: Checking reference counts
Pass 5: Checking group summary information
/dev/xvdc1: 11/655360 files (0.0% non-contiguous), 83137/2620928 blocks
```

- Step 13** Run the following command to expand the size of the file system on **/dev/xvdc1**:

resize2fs /dev/xvdc1

Information similar to the following is displayed:

```
[root@ecs-1120 linux]# resize2fs /dev/xvdc1
resize2fs 1.42.9 (28-Dec-2013)
Resizing the filesystem on /dev/xvdc1 to 15728379 (4k) blocks.
The filesystem on /dev/xvdc1 is now 15728379 blocks long.
```

- Step 14** Run the following command to view the disk partition information after the partition expansion:

lsblk

Information similar to the following is displayed:

```
[root@ecs-1120 linux]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
xvda 202:0 0 80G 0 disk
├─xvda1 202:1 0 40G 0 part /
└─xvda2 202:2 0 40G 0 part /opt
xvdb 202:16 0 350G 0 disk
├─xvdb1 202:17 0 100G 0 part
└─xvdb2 202:18 0 200G 0 part
xvdc 202:32 0 60G 0 disk
└─xvdc1 202:33 0 60G 0 part
```

In the command output, the total capacity of the **/dev/xvdc** disk is 60 GB, in which the additional 50 GB has been allocated to the **dev/xvdc1** partition.

- Step 15** Run the following command to mount the created partition to the **/mnt/sdc** directory:

mount /dev/xvdc1 /mnt/sdc

- Step 16** Run the following command to view the attachment result of **/dev/xvdc1**:

df -TH

Information similar to the following is displayed:

```
[root@ecs-1120 linux]# mount /dev/xvdc1 /mnt/sdc
[root@ecs-1120 linux]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/xvda1      ext4      43G   8.3G   33G   21% /
devtmpfs        devtmpfs  885M    0  885M    0% /dev
tmpfs           tmpfs     894M    0  894M    0% /dev/shm
tmpfs           tmpfs     894M   18M   877M    2% /run
tmpfs           tmpfs     894M    0  894M    0% /sys/fs/cgroup
tmpfs           tmpfs     179M    0  179M    0% /run/user/2000
tmpfs           tmpfs     179M    0  179M    0% /run/user/0
tmpfs           tmpfs     179M    0  179M    0% /run/user/1001
/dev/xvda2      ext4      43G   51M   40G    1% /opt
/dev/xvdc1      ext4      64G   55M   60G    1% /mnt/sdc
```

The **/dev/xvdc1** partition has been mounted on the **/mnt/sdc** directory.

----End

3.7.9 Adding a Data Disk

You can expand disk capacity by adding data disks. This section describes how to add data disks to expand the available capacity of an instance.

Procedure

- Step 1** Apply for a data disk as required by services. For details, see section [Applying for an EVS Disk](#).
- Step 2** Attach the data disk to the instance requiring capacity expansion. For details, see [Attaching an EVS Disk](#).
- Step 3** After the attachment is complete, initialize the data disk by referring to section [Initializing a Data Disk](#).

----End

3.8 Releasing an EVS Disk

3.8.1 Detaching an EVS Disk

This section describes how to detach an EVS disk from an ECS.

The following describes how to detach a data disk. The operations for detaching a system disk are the same.

You can detach a data disk in the following ways:

- ECS console: To detach multiple data disks from one ECS, perform related operations on the ECS console.
- EVS console: To detach a shared disk from multiple ECSs, perform related operations on the EVS console. In other scenarios, you can perform related operations on either console.

Prerequisites

If automatic attachment upon instance start has been configured for the Linux operating system by referring to [Setting Automatic Disk Attachment Upon](#)

Instance Start, comment out or delete the added content in the `/etc/fstab` file before detaching the disk. Otherwise, you may fail to access the operating system after the disk is detached.

Restrictions

- ECSs of the KVM virtualization type support online data disk detachment, namely, you can detach a data disk from an ECS in **Running** state when the following conditions are met:
 - The ECS is running steadily. You are not advised to detach a disk immediately after the ECS is started. Otherwise, the operation may fail.
 - For a VBD disk, **Disk Device Type** should be set to **scsi** or **virtio** when the operating system of the ECS is registered with Service OM.
 - The operating system of the ECS supports online disk detaching. For details, see the latest SIA compatibility document. Search for **Hot-plugging out disks** in the document to check whether the operating system supports online disk detaching.

NOTE

Obtain *FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Private Cloud)*:

- x86
 - Carrier users: [Click here](#). Search for **FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Private Cloud)**.
 - Enterprise users: [Click here](#). Search for **FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Private Cloud)**.
- Arm
 - Carrier users: [Click here](#). Search for **FusionSphere SIA Huawei Guest OS Compatibility Guide (ARM)**.
 - Enterprise users: [Click here](#). Search for **FusionSphere SIA Huawei Guest OS Compatibility Guide (ARM)**.
- ECSs of the Xen virtualization type support online data disk detachment. However, ECSs of the VMware virtualization type do not support online data disk detachment.
- System disks cannot be detached online.
- Before detaching a disk online from an instance running Windows, log in to the instance to perform the offline operation and confirm that UVP VMTools has been installed on the ECS and is running properly. At the same time, ensure that this disk is not being read and written. Otherwise, the disk will fail to be detached.
- Before detaching a disk online from an instance running Linux, log in to the instance, run the **umount** command to cancel the relationship between the disk and the file system, and confirm that the disk is not being read and written. Otherwise, the disk will fail to be detached.

ECS Console


Step 1 Log in to ManageOne as a VDC administrator or VDC operator using a browser.

URL in non-B2B scenarios: <https://Domain name of ManageOne Operation Portal>, for example, <https://console.demo.com>.

URL in B2B scenarios: **https://Domain name of ManageOne Tenant Portal**, for example, **https://tenant.demo.com**.

You can log in using a password or a USB key.

- Login using a password: Enter the username and password.
The password is that of the VDC administrator or VDC operator.
- Login using a USB key: Insert a USB key with preset user certificates, select the required device and certificate, and enter a PIN.

Step 2 Click  in the upper left corner, select a region and resource set, and choose **Storage > Elastic Volume Service**. The EVS console is displayed.

Step 3 In the ECS list, locate the ECS from which you want to detach the disk.

Step 4 Click the name of the ECS.

The page providing details about the ECS is displayed.


Step 5 On the **EVS** tab, select the data disk to be detached and click **Detach**.

Step 6 Click **OK**.

Step 7 After the detaching is complete, the EVS disk will no longer be displayed on the **EVS** tab of the ECS.

----End

Operations on EVS Console

Step 1 Click  in the upper left corner, select a region and resource set, and choose **Storage > Elastic Volume Service**. The EVS console is displayed.

Step 2 Click **Detach** in the row where the EVS disk to be detached is located.

If you are attempting to detach a shared disk, select the instance where the shared disk resides. You can select multiple instances.

Step 3 Click **OK**.

- If the disk being detached is a non-shared disk, the detachment is successful after the EVS disk state changes to **Available**.
- If a shared disk is detached, **Status** of the EVS disk becomes **Available** only after it is detached from all instances.

----End

Follow-up Procedure

After the EVS disk is detached, you can delete it by referring to [delete the disk](#) to release space.

3.8.2 Deleting an EVS Disk

If an EVS disk is no longer needed, you can soft delete it or delete it permanently. The soft deleted EVS disk is moved to the recycle bin and can be restore. A permanently deleted EVS disk cannot be restored.

The following describes how to delete a data disk. The operations for deleting a system disk are the same.

Restrictions

- If a disk has been attached to an instance, the disk cannot be deleted.
- If a disk has snapshots, the disk can be deleted only when the snapshot status is **Available** or **Error**.
- You can delete a disk only when the disk status is **Available**, **Error**, **Restoration failed**, or **Rollback failed**, and no VM snapshot has been created for the ECS where the disk resides.
- Disks configured with the DR service (CSDR/CSHA/VHA) cannot be deleted.
- If an EVS disk has a snapshot, the EVS disk can be soft deleted only when the snapshot is in the **Available** or **Error** state.
- When an EVS disk is permanently deleted, all snapshots of the EVS disk are also deleted.
- A shared disk to be deleted must have been detached from all instances.
- If the ECS to which the EVS disk belongs has not been created, the EVS disk cannot be deleted.
- Local disks can be used as data disks or system disks for an ECS. When a local disk is used as the system disk or a data disk, its life cycle starts and ends with the ECS, and cannot be manually detached or deleted.

Procedure

Step 1 Log in to EVS Console as a VDC administrator or VDC operator. For details, see [Logging In to the EVS Console as a VDC Administrator or VDC Operator](#).

Step 2 On the **Elastic Volume Service** page, locate the row that contains the target data disk, click **More** in the **Operation** column, and choose **Delete**.

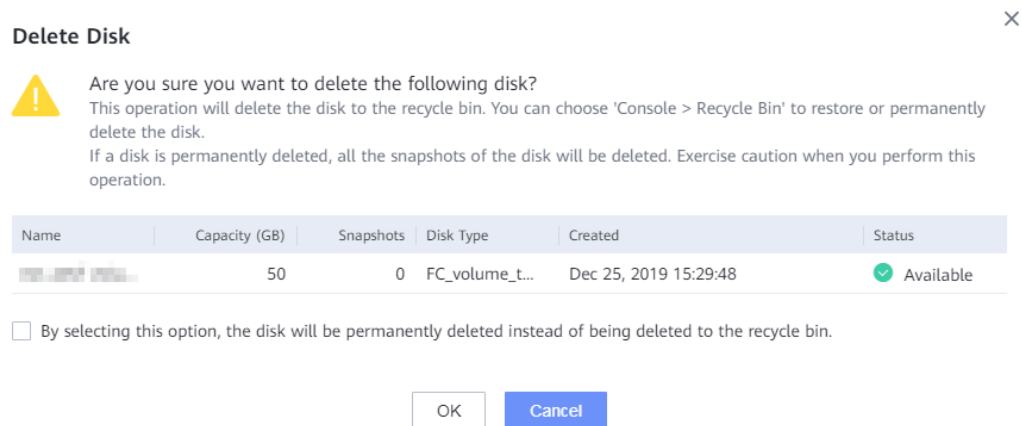
To delete multiple data disks at a time, select ☐ in front of the data disks on the **Elastic Volume Service** page.

Step 3 Delete the data disk.

- If you want to delete an EVS disk permanently, select **By selecting this option, the disk will be permanently deleted instead of being deleted to the recycle bin.** and click **OK**.

When an EVS disk is permanently deleted, all snapshots of the EVS disk are also deleted and cannot be restored.

- If you want to delete a data disk but require that the deleted data disk can be restored later, click **OK** to soft delete the data disk.





Step 4 If deleting a data disk requires approval, contact the administrator for approval. Otherwise, skip this step.

- After a data disk is permanently deleted, it is no longer displayed in the EVS disk list.
- After a data disk is soft deleted, the status of the data disk changes to **Soft deleted**.

----End

Follow-up Procedure

A deleted EVS disk is stored in the recycle bin and still occupies storage space.

- To restore data on a disk that has been **Soft deleted**, click  in the upper left corner of the page and choose **Recycle Bin**. On the **Recycle Bin** page, select the target EVS disk and click **Restore**.
- To permanently delete an EVS disk that has been **Soft deleted**, click  in the upper left corner of the page and choose **Recycle Bin**. On the **Recycle Bin** page, select the target EVS disk and click **Delete Permanently**.

3.9 Managing Disk Types

3.9.1 Changing a Disk Type

When the read and write performance of the storage device where the upper-layer service resides no longer suits the service, you can change the disk type to alter the type of the storage device to change the read and write performance, meeting the requirements of varying service performance of the instance. For example: Thick LUNs have been configured. It is found later that the actually used capacity is small. Therefore, thick LUNs need to be changed to thin LUNs. SAS disks have been purchased but later services require new high-performance SSDs. In this case, services need to be migrated from SAS disks to SSDs.

Restrictions

- Changing the disk type is supported when the backend storage is OceanStor V3/V5, OceanStor Dorado V3/6.x, or Huawei Distributed Block Storage.
- If the backend storage is OceanStor V3/V5 or OceanStor Dorado V3/6.x, the disk type can be changed between different storage pools in the same storage system. If the backend storage is Huawei Distributed Block Storage, the disk type can be changed only in the same storage pool.
- If the backend storage is Huawei Distributed Block Storage, the disk type can be changed only by modifying the QoS attribute.
- The administrator needs to import the SmartMigration license on the device in advance if the backend storage is OceanStor V3/V5 or Dorado V3.
- When changing the disk type, if the backend storage is OceanStor Dorado 6.x, the administrator needs to check whether the SmartMigration license has been imported to the device in advance. (The basic software package of OceanStor Dorado 6.x contains the SmartMigration license.)
- You can change the type of the EVS disk only in the **Available** or **In-use** state.
- If a disk has snapshots or is configured with the backup service (VBS/CSBS) or the DR service (CSDR/CSHA/VHA), the disk type cannot be changed.
- If backend storage of the disk is heterogeneous storage, the disk type cannot be changed.

NOTICE

If you change the disk type of an EVS disk that has been attached to an instance, the service performance of the instance will be affected.

Procedure

Step 1 Log in to the EVS console as a VDC administrator or VDC operator. For details, see [Logging In to the EVS Console as a VDC Administrator or VDC Operator](#).

Step 2 On the **Elastic Volume Service** page, locate the row that contains the target EVS disk, click **More** in the **Operation** column, and choose **Change Disk Type**.

<input type="checkbox"/>	Name	Status	Share...	Capacity...	Disk Type	Disk properties	Created	Expires	AZ	Attaching Inf...	Operation
<input type="checkbox"/>	ecs-d6ee-volume-0000	In-use	No	1	business_type...	System disk	Dec 25, 2019 19:30:56	Never	ecs-d6ee	ECS	Attach Detach More
<input type="checkbox"/>	ecs-f35f-volume-0001	In-use	Yes	60	business_type...	Data disk	Dec 25, 2019 19:11:06	Never	ecs-f35f	ECS	Expand Capacity Apply for Disk Apply for Snapshot Delete
<input type="checkbox"/>	ecs-f35f-volume-0000	In-use	No	1	business_type...	System disk	Dec 25, 2019 19:10:33	Never	ecs-f35f	ECS	Extend
<input type="checkbox"/>	ecs-7f9e-volume-0001	Available	No	60	business_type...	Data disk	Dec 25, 2019 17:26:05	Never	--	--	Change Disk Type
<input type="checkbox"/>	ecs-7f9e-volume-0000	In-use	No	1	business_type...	System disk	Dec 25, 2019 17:25:30	Never	ecs-7f9e	ECS	

Step 3 Select the target disk type, and click **Next**.

NOTE

When the backend storage device of the source disk type is Huawei Distributed Block Storage, at least one of the backend storage devices of the selected target disk type should be Huawei Distributed Block Storage. Otherwise, the disk type cannot be changed.

Step 4 Confirm the application information and click **Apply Now**.

Step 5 If changing the disk type of this EVS disk requires approval, contact the administrator for approval. Otherwise, skip this step.

On the **Elastic Volume Service** page, view the disk type of the EVS disk. If the disk type of the EVS disk has changed as expected, changing the disk type is successful.

 **NOTE**

If the EVS disk is still in **Retyping** state for a long time, contact the administrator to rectify the fault by referring to "Storage Services" > "Elastic Volume Service (EVS for ECS)" > "FAQs" > "What Can I Do If an EVS Disk Whose Disk Type Is Changed Remains in the Retyping State for a Long Time?" in *HUAWEI CLOUD Stack 8.1.0 Resource Provisioning Guide*.

----End

3.10 Managing Snapshots

3.10.1 Applying for a Snapshot

A snapshot can capture the data and status of a disk at a certain time point. If a service change or application software upgrade is required, you can create a snapshot for the disk in advance. If a fault occurs during the change or upgrade, you can use the snapshot to quickly restore disk data, ensuring service continuity and security. You can also use snapshots for routine backup of disk data.

Restrictions

- If backend storage is one of OceanStor V3/V5 or OceanStor Dorado V3 series, it is necessary for the administrator to import the HyperSnap license on the device in advance.
- Snapshots can be created only for disks in the **Available** or **In-use** state.
- A snapshot name cannot be the same as that of the prefix of the temporary snapshot created by the backup service (VBS/CSBS), the DR service (CSDR/CSHA/VHA), or the VM snapshot.
- Snapshots created using the EVS console consume the capacity quota instead of quantity quota of EVS disks.
- Temporary snapshots created by the backup service (VBS/CSBS) or the DR service (CSDR/CSHA/VHA) do not consume EVS disk quotas. Snapshots created using the VM snapshot function do not consume EVS disk quotas.
- Snapshots created using the EVS console, temporary snapshots created by DR and backup services, and snapshots created using the VM snapshot function consume backend storage capacity. If a large number of snapshots are created, contact the administrator to set the thin provisioning ratio of backend storage to a large value, preventing EVS disk provisioning failures caused by excessive snapshots.
- If backend storage of the disk is heterogeneous storage, snapshots can be created.
- In the VRM or VMware scenario, no snapshots can be created for shared EVS disks.

- If an EVS disk is created from data storage of the VIMS type in the VRM scenario and the EVS disk has been attached to an ECS, a snapshot can be created for the EVS disk only when the ECS is in the stopped state.
- In the VMware scenario, no snapshots can be created for an EVS disk that has been attached to an ECS.
- No snapshots can be created for disks that have expired.
- No snapshots can be created for disks that have been soft deleted.
- Snapshots cannot be created when the disk status is **Reserved** or **Maintenance**.
- If a task for creating a snapshot fails, the task is automatically deleted.

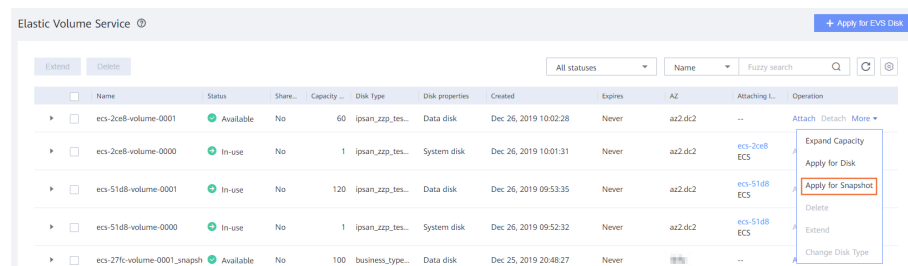
Procedure

Step 1 Log in to the EVS console as a VDC administrator or VDC operator. For details, see [Logging In to the EVS Console as a VDC Administrator or VDC Operator](#).

Step 2 Use one of the following methods to display the **Apply for Snapshot** page.

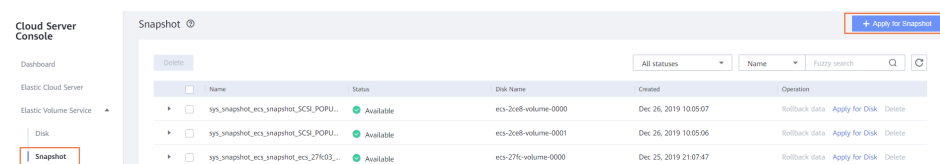
- Method 1

In the EVS disk list, locate the row that contains the target disk, click **More** in the **Operation** column, and choose **Apply for Snapshot**.



- Method 2

In the navigation pane on the left, choose **Elastic Volume Service > Snapshot**. Then, click **Apply for Snapshot**.



Step 3 Configure the basic information about the snapshot, as shown in [Table 3-19](#).

Table 3-19 Parameters for creating a snapshot

Parameter	Description	Example Value
Name	The value contains only digits, letters, underscores (_), and hyphens (-) and cannot exceed 63 characters. NOTE The snapshot name cannot be the same as the name prefix of any temporary snapshot created by the backup service (VBS or CSBS) or the DR service (CSDR, CSHA, or VHA) or the name prefix of any VM snapshot. The prefixes include autobk_snapshot , manualbk_snapshot , and sys_snapshot .	snapshot-01
EVS Disk	<ul style="list-style-type: none">When you create a snapshot for the specified disk, the disk has been determined.When you create a snapshot on the snapshot page, click Select, select the target disk, and click OK.	volume-01
Description	Description of the created snapshot. The value contains a maximum of 63 bytes.	-

Step 4 Click **Next**.

Step 5 If you do not need to modify the specifications, click **Apply Now** to start the snapshot application.

If you need to modify the specifications, click **Back** to modify parameter settings.

Step 6 If creating a snapshot requires approval, contact the administrator for approval. If no, skip this step.

Return to the **Snapshot** page. If the snapshot status is **Available**, the snapshot is created successfully.

----End

3.10.2 Rolling Back an EVS Disk Using a Snapshot

If the data of an EVS disk is lost or damaged, or if you want to roll back the disk to a certain time point and you have created a snapshot for the disk, you can use the snapshot to roll back the disk to the time when the snapshot was created.

Restrictions

- Temporary snapshots created by the backup service (VBS/CSBS) or the DR service (CSDR/CSHA/ VHA) cannot be rolled back.
- Snapshots created for disks having any DR service (CSDR/CSHA/VHA) configured cannot be rolled back.
- Snapshots created using the VM snapshot function cannot be used for EVS disk rollback.

- After an EVS disk without VM snapshots is attached to a VM with VM snapshots, the EVS disk will be detached when the VM is rolled back using a VM snapshot.
- If backend storage of the disk is heterogeneous storage (excluding XSKY), EVS disk rollback from a snapshot is not supported.
- A snapshot can be used to roll back its source EVS disk, and cannot be used to roll back any other EVS disk.
- A rollback can be performed only when the snapshot status is **Available** and the status of the snapshot source disk is **Available** (that is, the snapshot is not attached to any instance) or **Rollback failed**.
- When the source disk of a snapshot is in the recycle bin, EVS disk rollback from the snapshot is not supported.
- The rollback operation is irreversible. After the rollback is complete, all data generated from the time when the snapshot was created to the time when the rollback operation is performed will be lost and cannot be restored. Exercise caution when performing this operation.

Procedure

- Step 1** Log in to the EVS console as a VDC administrator or VDC operator. For details, see [Logging In to the EVS Console as a VDC Administrator or VDC Operator](#).
- Step 2** In the navigation pane on the left, choose **Elastic Volume Service > Snapshot**.
The **Snapshot** page is displayed.
- Step 3** In the snapshot list, locate the row that contains the target snapshot and click **Rollback data** in the **Operation** column.
- Step 4** After confirming that the rollback information is correct, click **OK** to perform the rollback.

Return to the **Snapshot** page. After the snapshot status changes from **Rolling back** to **Available**, the data rollback is successful.

----End

3.10.3 Deleting a Snapshot

Snapshots occupy the disk capacity quota. Therefore, if some snapshots are not required, you are advised to delete them in a timely manner to release space.

Context

- Users are allowed to delete a temporary snapshot created by the backup service (VBS/CSBS). After the snapshot is deleted, if users want to back up the EVS disk corresponding to the snapshot, full backup is performed for the first time.
- Temporary snapshots created by the DR service (CSDR/CSHA/VHA) cannot be deleted.
- A snapshot created using the VM snapshot function cannot be deleted, and the name of the snapshot cannot be changed.

- You can delete a snapshot only when its state is **Available**, **Deletion failed**, or **Error**.

NOTICE

The deletion operation is irreversible. Exercise caution when performing this operation.

Procedure

- Step 1** Log in to the EVS console as a VDC administrator or VDC operator. For details, see [Logging In to the EVS Console as a VDC Administrator or VDC Operator](#).
- Step 2** In the navigation pane on the left, choose **Elastic Volume Service > Snapshot**.
- Step 3** Delete snapshots.
- Deleting one snapshot
 - a. On the **Snapshot** page, locate the row that contains the target snapshot and click **Delete** in the **Operation** column.
 - b. In the displayed dialog box, confirm the information and click **OK**.
 - Deleting multiple snapshots at a time
 - a. On the **Snapshot** page, select multiple snapshots by selecting ☐ next to them.
 - b. Click **Delete** in the upper part of the page. In the displayed dialog box, click **OK**.
- Step 4** If deleting a snapshot requires approval, contact the administrator for approval. If no, skip this step.

----End

3.11 Managing Backups

3.11.1 Creating a Backup

To prevent data loss caused by misoperations or system faults, you can create a backup for an EVS disk to ensure data correctness and security.

You can create backups for EVS disks in the following ways:

- EVS console: To create backups for a single EVS disk, perform related operations on the EVS console.
- VBS console: To create backups for multiple EVS disks in batches, perform related operations on the VBS console.

Restrictions

- Only disks in the **Available** or **In-use** state can be backed up.

- If backend storage of the disk is heterogeneous storage, backups cannot be created.

EVS Console

Before creating a backup on the EVS console, ensure that the VBS service has been deployed in the environment and enabled by the administrator. Otherwise, this function is unavailable. For details about how to enable the VBS service, see "Storage Services" > "Elastic Volume Service (EVS for ECS)" > "FAQs" > "How Do I Enable the VBS?" in *HUAWEI CLOUD Stack 8.1.0 Resource Provisioning Guide*.

Step 1 Log in to the EVS console as a VDC administrator or VDC operator. For details, see [Logging In to the EVS Console as a VDC Administrator or VDC Operator](#).

Step 2 On the **Elastic Volume Service** page, locate the row that contains the target EVS disk, click **More** in the **Operation** column, and then choose **Create Backup**.

Step 3 On the **Create VBS Backup** page, configure parameters as prompted.

For details, see [Operation Help Center > DR & Backup > User Guide > Creating a Periodic Backup Task > Creating a Backup Task](#).

----End

VBS Console

For details about how to create a backup task on the VBS page, see [Operation Help Center > DR & Backup > User Guide > Creating a Periodic Backup Task > Creating a Backup Task](#).

Follow-up Procedure

After an EVS disk is backed up successfully, a VBS backup of the disk is generated. You can view the details about the backup on [Operation Help Center > DR & Backup > User Guide > Managing Backups and Replicas > Viewing EVS Disk Backups and Replicas](#). You can also view backup tree information on the **Elastic Volume Service** page.

3.11.2 Using a Backup to Restore Data

If the data on the current EVS disk is lost or damaged and you have created a backup for the disk, you can use the backup to restore the data on the disk to the time when the backup was created.

- You can use a backup to restore data to the original EVS disk. For details, see [Operation Help Center > DR & Backup > User Guide > Restoring a Disk Using a Backup > Restoring Backup Data to the Source EVS Disk](#).
- If the current EVS disk is damaged, you can use the backup to restore the data to another EVS disk. For details, see [Operation Help Center > DR & Backup > User Guide > Restoring a Disk Using Backup > Restoring Backup Data to Another EVS Disk](#).

3.12 Managing EVS Disks

3.12.1 Extending the EVS Disk Validity Period

This section describes how to extend the validity period of an EVS disk. You can extend the validity period of multiple EVS disks in batches. When an EVS disk expires, you can only delete the disk or extending the disk validity period. If your EVS disk is about to expire, you are advised to extend the validity period as soon as possible.

Restrictions

- If an EVS disk is created with an instance, the validity period of the EVS disk is unlimited.
- If the validity period of an EVS disk is unlimited, the validity period cannot be extended.
- You can extend the EVS disk validity period only when the disk is in the **available** or **In-use** status.
- If an EVS disk has expired, its snapshot cannot be used to roll back the EVS disk or create an EVS disk. To continue using this EVS disk, extend its validity period.
- When an EVS disk expires, its data will not be deleted. You can continue using this EVS disk after extending its validity period.

Procedure

Step 1 Log in to the EVS console as a VDC administrator or VDC operator. For details, see [Logging In to the EVS Console as a VDC Administrator or VDC Operator](#).

Step 2 On the **Elastic Volume Service** page, click  to select the target EVS disk.

You can also choose **More > Extend** to extend the validity period of a single EVS disk.

Step 3 Click **Extend**.

Step 4 In the **Configuration** area, select **Extension Period** and enter the description of the EVS disk. Click **Next**.

The extension period can be set to **Unlimited**, **1 year**, or **Custom**.

Step 5 On the **Resource Details** page, confirm the extension information.

- After confirming that the information is correct, click **Apply Now** to start extending the EVS disk validity period.
- If you need to modify the specifications, click **Back** to modify parameter settings.

Step 6 If this application requires approval, contact the administrator for approval. If no, skip this step.

On the **Elastic Volume Service** page, view the expiration time of the EVS disk. If the expiration time of the EVS disk changes as expected, validity period extending has succeeded.

----End

3.12.2 Viewing Monitoring Data

If an EVS disk has been attached to an instance which is not of the VMware virtualization type and in use, you can refer to this section to view the monitoring data of the EVS disk, such as the read and write rates.

Procedure

- Step 1** Log in to the EVS console as a VDC administrator or VDC operator. For details, see [Logging In to the EVS Console as a VDC Administrator or VDC Operator](#).
 - Step 2** Find the EVS disk whose monitoring data you want to view and click ► next to the disk. The details about the disk will be displayed.
 - Step 3** Click **View Monitoring Data** to view the read and write rates of the EVS disk.
- End

3.12.3 Changing the VDC Quota

When a ManageOne operation administrator creates a tenant and configures the VDC service, the operation administrator needs to set not only the resource pool for each region but also the resource quota for the tenant and the resource quota for the VDC of this level. The EVS disk quota consists of the total capacity and total quantity for a tenant and the total capacity and total quantity for the VDC of this level, which you can configure. When the EVS disk quota in the VDC is insufficient, contact the administrator to change the quota. For details, see .

3.13 FAQs

3.13.1 Can I Attach an EVS Disk to Multiple Instances?

A non-shared EVS disk can be attached to only one instance.

A shared EVS disk can be attached to a maximum of 16 instances by default.

3.13.2 How Can I Attach an EVS Disk to an ECS?

For details, see section [Attaching an EVS Disk](#).

3.13.3 How Many States Does an EVS Disk Have?

Disks are managed using the EVS during the entire process from disk creation to release, ensuring optimal user experience of applications or sites hosted on them.

[Figure 3-30](#) shows the switching between different states of an EVS disk, and [Table 3-20](#) describes the meaning and supported operations of each state.

Figure 3-30 EVS disk status

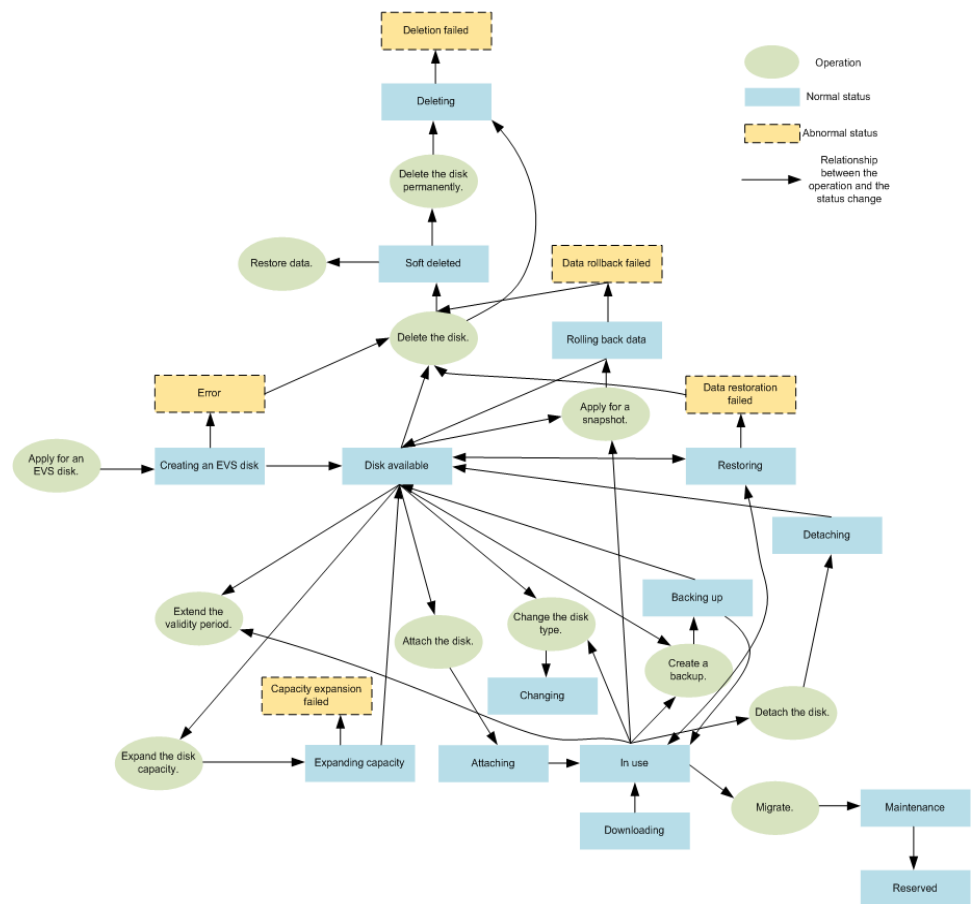


Table 3-20 EVS disk status description

EVS Disk Status	Status Description	Supported Operation
In-use	The EVS disk is attached to an instance and is in use.	<ul style="list-style-type: none"> Detach Expand capacity Create snapshot Create backup Extend validity period of EVS disks (except EVS disks whose expiration time is Unlimited) Change disk type
Downloadin g	Data is being downloaded from an image to the EVS disk. An EVS disk is in this status when you are creating an ECS.	-

EVS Disk Status	Status Description	Supported Operation
Available	The EVS disk is created and has not been attached to any server.	<ul style="list-style-type: none">• Attach• Expand capacity• Create snapshot• Create backup• Delete• Extend validity period of EVS disks (except EVS disks whose expiration time is Unlimited)• Change disk type
Creating	The EVS disk is being created.	-
Attaching	The EVS disk is being attached to an instance.	-
Detaching	The EVS disk is being detached from an instance.	-
Deleting	The EVS disk is being deleted.	-
Extending	The capacity of the EVS disk is being expanded.	-
Rolling back	The EVS disk data is being rolled back using the snapshot.	-
Backing up	The EVS disk is being backed up.	-
Restoring	The EVS disk data is being restored using the backup.	-
Retyping	The EVS disk type is being changed.	-
Soft deleted	The EVS disk has been soft deleted and moved to the recycle bin.	<ul style="list-style-type: none">• Restore• Permanently delete
Error	An error occurs when you are creating an EVS disk. You can delete the EVS disk and create it again.	Delete
Deletion failed	An error occurs when you are deleting an EVS disk. Contact the administrator.	None

EVS Disk Status	Status Description	Supported Operation
Expansion failed	An error occurs when you are expanding the capacity of an EVS disk. The administrator will contact you and help you handle this error. Do not perform any operations on the disk before the administrator contact you. If you require that the error be handled as soon as possible, contact the administrator.	-
Rollback failed	An error occurs during an EVS disk rollback.	Delete
Restoration failed	An error occurs during EVS disk restoration. The administrator will contact you and help you handle this error. Do not perform any operations on the disk before the administrator contact you. If you require that the error be handled as soon as possible, contact the administrator.	Delete
Maintenance	The EVS disk is being migrated.	-
Reserved	Status of the source disk after the migration task is complete	-

3.13.4 Does an EVS Disk or Snapshot Generate Metering Information in All States?

EVS disks and snapshots do not generate metering information in the following states:

- EVS disk state: Creating, Error, Restoration failed, Downloading, Expansion failed, Deleting, Deletion failed, Rollingback, Rollback failed, or Retyping
- Snapshot state: Creating, Error, Deleting, or Deletion failed

3.13.5 Can I Change the EVS Disk Capacity?

EVS disk capacity can be expanded but cannot not be reduced at present.

For details about how to expand the EVS disk capacity, see [Expanding EVS Disk Capacity](#).

3.13.6 What Can I Do If an EVS Disk Whose Disk Type Is Changed Remains in the Retyping State for a Long Time?

This section describes how to handle the problem that an EVS disk whose disk type is changed remains in the **Retyping** state for a long time.

Procedure

Step 1 Check whether the migration task is complete on the storage device.

1. Log in to the OceanStor DeviceManager as the **admin** user.
See the default password of the target account on the "Type A (Portal)" sheet in *HUAWEI CLOUD Stack 8.1.0 Account List*.
2. Choose **Provisioning > Resource Performance Tuning > SmartMigration** to display the migration task page.
3. Use the source LUN ID to search for the migration task.
You can use the ID of the EVS disk whose disk type is changed to obtain the corresponding source LUN ID. For details, see [Related Operations](#).
 - If the running status of the current LUN is **Migrating**, the task for changing the disk type is not complete. Wait until the task is complete.
 - If the running status of the current LUN is **Migrated** and the status of the EVS disk on the EVS console is **Retyping**, go to step [Step 2](#).

Step 2 Tick ☐ before the source LUN, and click **Delete** to delete the completed migration task.

Step 3 Reset the statuses of the EVS disk and the migration task.

1. Log in to the FusionSphere OpenStack controller node and import environment variables. For details, see [Step 1](#) to [Step 3](#) in [Related Operations](#).
2. Run the following commands to reset the statuses of the EVS disk and the migration task:

```
cinder reset-state --state Status of the EVS disk before changing the disk type volume_id
```

```
cinder reset-state --reset-migration-status error volume_id
```

In the command, *Status of the EVS disk before changing the disk type* can be **in-use** or **available**. If an EVS disk has been attached, its status is **in-use**. If an EVS disk has not been attached, its status is **available**.

volume-id is the EVS disk ID. You can view it in the EVS disk details on the EVS console.

Step 4 Change the disk type again.

----End

Related Operations

To use an EVS disk ID to obtain the corresponding LUN ID, perform the following operations:

Step 1 Use PuTTY to log in to the FusionSphere OpenStack controller node.

Account: **fsp**

See the default password of the account for logging in to the target node on the "Type A (Background)" sheet in *HUAWEI CLOUD Stack 8.1.0 Account List*.

Step 2 Run the **su root** command and enter the password of user **root** to switch to user **root**.

See the default password of the account for logging in to the target node on the "Type A (Background)" sheet in *HUAWEI CLOUD Stack 8.1.0 Account List*.

Step 3 Run the following command to import the environment variables:

source set_env

During the import, enter **1**, select the environment variable **openstack environment variable (keystone v3)**, and enter the password of account **dc_admin**.

See the default password of the target account on the "Type B (FusionSphere OpenStack)" sheet in *HUAWEI CLOUD Stack 8.1.0 Account List*.

Step 4 Run the following command to query the LUN ID:

cinder show volume-id

In the command, *volume-id* is the EVS disk ID. You can view it in the EVS disk details on the EVS console.

```
EBF8BBAC-5D50-63B5-E811-0A3EC4703446:/opt # cinder show 617143ce-a425-4413-8675-1fba7238c012
+-----+-----+
| Property | Value |
+-----+-----+
| attachments | [] |
| availability_zone | az1.dc1 |
| bootable | False |
| consistencygroup_id | None |
| created_at | 2018-08-31T10:01:32.344815 |
| description | None |
| encrypted | False |
| id | 617143ce-a425-4413-8675-1fba7238c012 |
| metadata | {'StorageType': 'Dorado V5-V500R007C10', 'lun_wwn': '63ce82410069ebad7064f24a000021bc'} |
| migration_status | None |
| multiattach | False |
| name | ww'jecho vat2fj994j 3hs9um2ks4 #xzwx |
| os-vol-host-attr:host | cinder-kvm016@su_test002#su_test002 |
| os-vol-mig-status-attr:migstat | None |
| os-vol-mig-status-attr:name_id | None |
| os-vol-pro-location-attr:provider_location | 8636 |
| os-vol-tenant-attr:tenant_id | 688e90536de0458b9538449e17be389a |
| os-volume-replication:driver_data | [{"lun_id": "8636", "sn": "2102351PX5102000005"}] |
| os-volume-replication:extended_status | None |
| replication_status | disabled |
| shareable | False |
| size | 1 |
| snapshot_id | None |
| source_vol_id | None |
| status | available |
| updated_at | 2018-08-31T10:01:34.292718 |
| user_id | ca71b3431da84f52957aed0559c99ae1 |
| volume_type | su_test002 |
+-----+-----+
```

----End

3.13.7 How Can I Test the Performance of a Windows EVS Disk?

This section describes how to test the performance of an EVS disk on an instance that runs on the Windows OS.

Prerequisites

Before testing the performance, log in to the Iometer official website to download and install the Iometer performance test tool.

Precautions

- In the EVS disk performance test, if the start cylinder number is not 4-KB aligned, the EVS disk performance will be greatly affected. Ensure that the start cylinder number is 4-KB aligned before you start the test.
- Do not perform the fio test on system disks to avoid damaging important system files.
- You are advised to perform the fio test on idle disks that do not store important data and to create a file system after the test is complete.
- Do not perform the test on data disks. Otherwise, disk data may be damaged if the metadata of the underlying file system is damaged.

Procedure

The method for testing an EVS disk varies depending on the in-use server OS. This section uses Windows 7 Professional 64-bit as the example OS to describe how to test the EVS disk performance.

Step 1 Log in to an instance.

Step 2 Press **Win+R** to open the **Run** window. Enter **msinfo32** and click **OK**.

The system information window is displayed.

Step 3 Choose **Component > Storage > Disk**. View the start offset of the partition in the right pane, and check whether the start cylinder number of the disk partition is 4-KB aligned.

- If 4096 can be divided by the start cylinder number, the partition is 4-KB aligned. Go to [Step 4](#) to start the test.
- If 4096 cannot be divided by the start cylinder number, the partition is not 4-KB aligned. If you need to continue to perform the test, ensure the 4-KB alignment for the partition.

NOTICE

Data losses will occur when you delete the partition and select another start cylinder number for the 4-KB alignment. Exercise caution when performing such an operation.

Step 4 Use Iometer to test the EVS disk performance. For details, see the Iometer product document.

When the disk IOPS and throughput are tested, the parameter configurations for Iometer and fio are the same. For details, see [Table 3-21](#).

----End

3.13.8 How Can I Test the Performance of a Linux EVS Disk?

This section describes how to test the performance of an EVS disk on an instance that runs on the Linux OS.

Prerequisites

fiio is a tool for testing disk performance. It is used to perform pressure tests on hardware. You are advised to use the I/O engine of Libaio to perform the test. You can log in to the official website to download and install fio and Libaio.

Precautions

- When you use an old version Linux OS, for example CentOS 6.5, and run **fdisk** to create partitions, the default start cylinder number is not 4-KB aligned, which will greatly affect the test performance. For that reason, if such an OS is used, you are advised to select a new start cylinder number that is 4-KB aligned when creating partitions.
- Do not perform the fio test on system disks to avoid damaging important system files.
- You are advised to perform the fio test on idle disks that do not store important data and to create a file system after the test is complete.
- Do not perform the test on data disks. Otherwise, disk data may be damaged if the metadata of the underlying file system is damaged.
- When testing disk performance, you are advised to directly test raw disks. When testing the performance of the file system, you are advised to specify a specific file for the test.

Procedure

The method for testing an EVS disk varies depending on the in-use server OS. This section uses CentOS 7.2 (64-bit) as the example OS to describe how to test the EVS disk performance.

Step 1 Log in to an instance and switch to the **root** user.

Step 2 Before you start the test, run the following command to check whether the start cylinder number is 4-KB aligned:

fdisk -lu

Information similar to the following is displayed:

```
[root@ecs-centos sdc]# fdisk -lu
```

```
Disk /dev/xvda: 10.7 GB, 10737418240 bytes, 20971520 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x7db77aa5
```

Device	Boot	Start	End	Blocks	Id	System
/dev/xvda1	*	2048	20968919	10483436	83	Linux

```
Disk /dev/xvdb: 10.7 GB, 10737418240 bytes, 20971520 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
```

```
Disk /dev/xvdc: 53.7 GB, 53687091200 bytes, 104857600 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
```

Disk label type: dos
Disk identifier: 0x3cf3265c

Device	Boot	Start	End	Blocks	Id	System
/dev/xvdc1		2048	41943039	20970496	83	Linux

- If 8 can be divided by the start cylinder number, the partition is 4-KB aligned. Go to [Step 3](#) to start the test.
- If 8 cannot be divided by the start cylinder number, the partition is not 4-KB aligned. If you need to continue to perform the test, select another start cylinder number to ensure the 4-KB alignment for the partition.

NOTICE

Data losses will occur when you delete the partition and select another start cylinder number for the 4-KB alignment. Exercise caution when performing such an operation.

Step 3 Run commands and use fio to test the EVS disk performance:

- To test random write IOPS, run **fio -direct=1 -iodepth=128 -rw=randwrite -ioengine=libaio -bs=4k -size=10G -numjobs=1 -runtime=600 -group_reporting -filename=/dev/[device] -name=Rand_Write_IOPS_Test**.
- To test random read IOPS, run **fio -direct=1 -iodepth=128 -rw=randread -ioengine=libaio -bs=4k -size=10G -numjobs=1 -runtime=600 -group_reporting -filename=/dev/[device] -name=Rand_Read_IOPS_Test**.
- To test write throughput, run **fio -direct=1 -iodepth=32 -rw=write -ioengine=libaio -bs=1024k -size=10G -numjobs=1 -runtime=600 -group_reporting -filename=/dev/[device] -name=Write_BandWidth_Test**.
- To test read throughput, run **fio -direct=1 -iodepth=32 -rw=read -ioengine=libaio -bs=1024k -size=10G -numjobs=1 -runtime=600 -group_reporting -filename=/dev/[device] -name=Read_BandWidth_Test**.

[Table 3-21](#) describes the parameters.

Table 3-21 Parameters for testing the disk performance

Parameter	Description
direct	Indicates whether to use the direct I/O. The options are as follows: <ul style="list-style-type: none">– If the value is 0, it indicates that the buffered I/O is used.– If the value is 1, it indicates that the direct I/O is used.
iodepth	Defines the depth of the I/O queue during the test. The default value is 1. This queue depth refers to the queue depth of each thread. For example, when multiple threads are tested, the parameter defines the queue depth of each thread. Total number of concurrent I/Os of fio = iodepth x numjobs

Parameter	Description
rw	Defines the test read/write policy. The values are as follows: <ul style="list-style-type: none">- Random read: randread- Random write: randwrite- Sequential read: read- Sequential write: write- Mixed random read/write: randrw
ioengine	Defines how the fio delivers I/O requests. There are synchronous I/O requests and asynchronous I/O requests. <ul style="list-style-type: none">- Only one synchronous I/O request can be delivered at a time. The request is returned only after the kernel processing completes. In this way, the I/O queue depth for a single thread is always less than 1. However, multi-thread concurrent processing can be used. Generally, 16 to 32 threads are used to work at the same time to stuff the I/O queue depth.- Usually, a batch of asynchronous I/O requests is submitted at a time by using the libaio method. After a batch of requests is processed, the number of interaction times is reduced, which is more efficient.
bs	Defines the I/O block size. The unit is k, K, m, or M. The default size is 4 KB.
size	Defines the amount of data processed by the test I/Os. If parameters, such as runtime , are not specified, the test ends until fio has processed all the specified data amount. The value can be a number with a unit or percentage. A number with a unit indicates the read/write data amount, for example size=10G , indicating a 10-GB read/write data amount. A percentage indicates the ratio of read/write data amount to the capacity of total files, for example size=20% , indicating the read/write data amount takes 20% of the total file space.
numjobs	Defines the number of concurrent threads.
runtime	Defines the test time. If this parameter is not specified, the test ends until the specified data amount is processed by the block size defined using parameter size .
group_reporting	Defines the test result display mode. The parameter value displays the statistics on a single thread instead of that on all jobs.
filename	Defines the name of the test file or device.
name	Defines the test task name.

----End

3.13.9 Logging In to the EVS Console as a VDC Administrator or VDC Operator

Procedure


Step 1 Log in to ManageOne as a VDC administrator or VDC operator using a browser.

URL in non-B2B scenarios: **https://Domain name of ManageOne Operation Portal**, for example, **https://console.demo.com**.

URL in B2B scenarios: **https://Domain name of ManageOne Tenant Portal**, for example, **https://tenant.demo.com**.

You can log in using a password or a USB key.

- Login using a password: Enter the username and password.
The password is that of the VDC administrator or VDC operator.
- Login using a USB key: Insert a USB key with preset user certificates, select the required device and certificate, and enter a PIN.

Step 2 Click  in the upper left corner, select a region and resource set, and choose **Storage > Elastic Volume Service**. The EVS console is displayed.

----End

4 User Guide (for BMS)

- [4.1 Introduction](#)
- [4.2 Related Concepts](#)
- [4.3 Operation Process](#)
- [4.4 Applying for an EVS Disk](#)
- [4.5 Attaching an EVS Disk](#)
- [4.6 Initializing a Data Disk](#)
- [4.7 Expanding EVS Disk Capacity](#)
- [4.8 Releasing an EVS Disk](#)
- [4.9 Managing Disk Types](#)
- [4.10 Managing Snapshots](#)
- [4.11 Managing Backups](#)
- [4.12 Managing EVS Disks](#)
- [4.13 FAQs](#)

4.1 Introduction

4.1.1 What Is Elastic Volume Service?

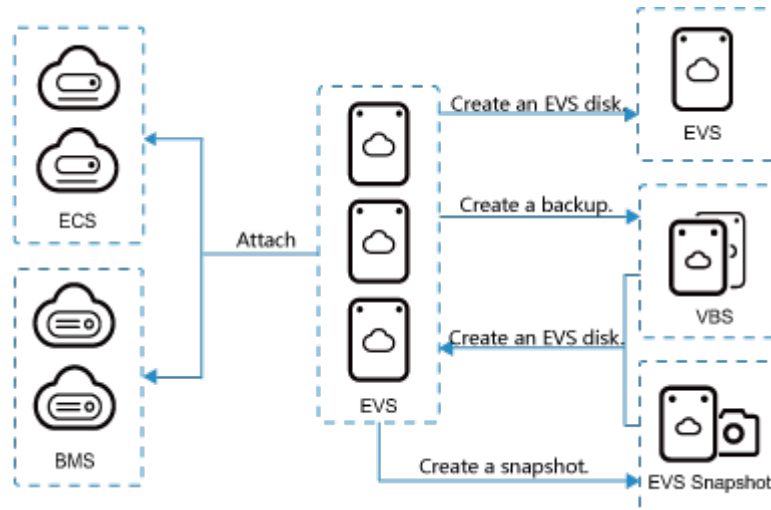
Definition

Elastic Volume Service (EVS) is a virtual block storage service, which provides block storage space for Elastic Cloud Servers (ECSs) and Bare Metal Servers (BMSs). You can create EVS disks on the console and attach them to ECSs. The method for using EVS disks is the same as that for using disks on physical servers. EVS disks have higher data reliability and I/O throughput and are easier to use. EVS disks are suitable for file systems, databases, or system software or applications that require block storage devices. [Figure 4-1](#) shows how to use an EVS disk.

In this document, an EVS disk is also referred to as a disk.

In this document, instances refer to the ECSs or BMSs that you apply for.

Figure 4-1 EVS functions



Functions

Various types of persistent EVS disks are available. You can choose disk types based on your needs and store files and build databases on EVS disks. EVS has the following major features:

- Elastic attaching and detaching
An EVS disk is like a raw, unformatted, external block device that you can attach to a single instance. Disks will not be affected by the running time of instances. After attaching a disk to an instance, you can use the disk as if you were using a physical disk. You can also detach a disk from an instance and attach the disk to another instance.
- Various disk types
A disk type represents backend storage devices used by a group of disks. You can divide disk types of EVS disks based on backend storage types to meet different performance requirements of services. When the read and write performance of the storage device no longer suits your services, you can change the disk type to alter the type of the storage media where the disk resides to change the read and write performance, meeting the requirements of the instance for higher storage service performance.
- Scalability
A single disk has a maximum capacity of 64 TB. You can configure storage capacity and expand the capacity on demand to deal with your service data increase.
- Snapshot
You can back up your data by taking a snapshot of your disk data at a specific point in time to prevent data loss caused by data tampering or mis-deletion and ensure a quick rollback in the event of a service fault. You can also create disks from snapshots and attach them to other instances to provide data resources for a variety of services, such as data mining, report query, and

development and test. Snapshots can be used to protect original data or create disks for rapidly deploying other services, meeting diversified service data requirements of enterprises.

- Shared disk

Multiple instances can access (read and write) a shared disk at the same time, meeting the requirements of key enterprises that require cluster deployment and high availability (HA).

Comparison Between EVS and SFS

Table 4-1 compares EVS and SFS.

Table 4-1 Comparison between EVS and SFS

Dimension	EVS	SFS
Usage	Provides persistent block storage for compute services such as ECS and BMS. EVS disks feature high availability, high reliability, and low latency. You can format, create file systems on, and persistently store data on EVS disks.	Provides ECSs with a high-performance shared file system that supports on-demand elastic scaling. The file system complies with the standard file protocol and delivers scalable performance, supporting massive amount of data and bandwidth-demanding applications.
Data access mode	Data access is limited within the internal network of a data center.	Data access is limited within the internal network of a data center.
Sharing mode	Supports EVS disk sharing. A shared EVS disk can be attached to a maximum of 16 ECSs in the cluster management system.	Supports data sharing. A file system can be mounted to a maximum of 256 ECSs.
Storage capacity	The maximum capacity of a single disk is 64 TB.	The maximum capacity of a single file is 240 TB, and the file system capacity can be scaled to the PB level.
Backend storage	Supports Huawei SAN storage, Huawei Distributed Block Storage, and heterogeneous storage.	OceanStor 9000 and OceanStor Dorado 6.x

Dimension	EVS	SFS
Recommended scenarios	Scenarios such as database, enterprise office applications, and development and testing.	Scenarios such as media processing and file sharing.

4.1.2 Advantages

- Varying specifications
EVS disks of different performance levels are provided. You can choose and configure EVS disks of appropriate performance levels to meet your service requirements.
- Scalable
EVS disks provide ultra-large block storage and a single EVS disk has a maximum capacity of 64 TB. You can expand the EVS disk capacity on running ECSs to meet your increasing service requirements.
 - On-demand expansion
You can expand the capacity of EVS disks based on your needs, with at least 1 GB added at a time.
 - Linear performance improvement
You can expand the capacity of EVS disks on running ECSs to implement linear performance improvement, thereby meeting your service requirements.
- Secure and reliable
Distributed storage is adopted, and data is stored in multiple identical copies, ensuring zero data loss. Data durability reaches 99.9999999%.
- Backup and restoration
Functions, such as EVS disk backup and EVS disk snapshot, are supported to prevent incorrect data caused by application exceptions or attacks.
 - EVS disk backup
This function enables the system to create EVS disk backups. The backups can be used to roll back EVS disks, maximizing user data accuracy and security and ensuring service availability.
 - EVS disk snapshot
This function enables the system to create snapshots for EVS disks. A snapshot can be used to roll back an EVS disk to the state when the snapshot is created, maximizing data accuracy and security and ensuring service availability.

4.1.3 Application Scenarios

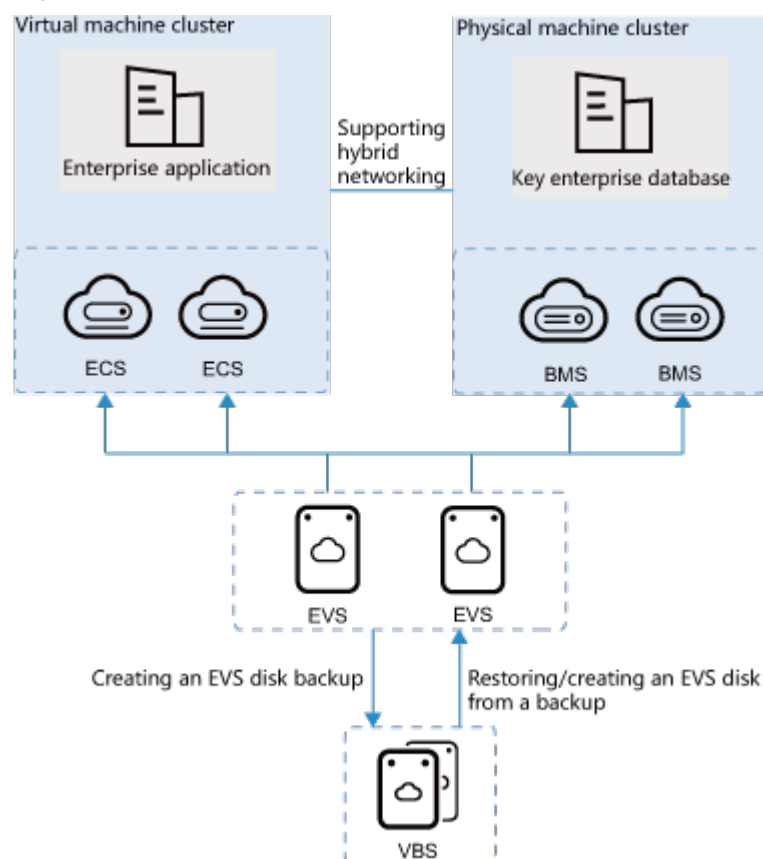
You can configure and select disk types with different service levels based on your application requirements for flexible deployment.

Relational Database

The service core database needs to support massive access at traffic peaks, and requires disks with persistent and stable high performance and low latency. You can use the disk type with ultra-high performance to implement a combination of excellent performance and superior reliability, meeting the high requirements for low latency and high I/O performance in data-intensive scenarios, such as NoSQL and relational databases. [Figure 4-2](#) shows the architecture in these scenarios. Disks with ultra-high performance service levels can meet the following performance requirements:

- The latency is shorter than 1 ms.
- The performance ranges from 2000 IOPS/TB to 20,000 IOPS/TB.
- Typical configurations: Enterprise storage Dorado 5000 V3 is selected for backend storage, 25 1 TB, 2 TB, or 4 TB SSDs are configured for every dual controllers, and RAID 6 is configured. Deduplication and compression functions are enabled, and a maximum of four controllers and 50 disks (30 TB, 60 TB, or 120 TB) are configured for a single system.

Figure 4-2 Architecture in the relational database scenario



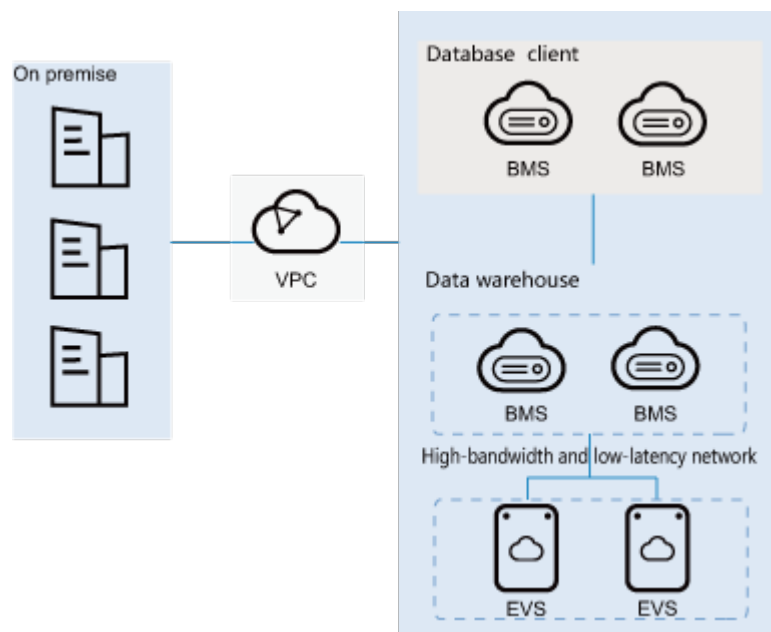
Data Warehouse

In scenarios with intensive data reads, deploy data warehouses, and it is recommended that you use the disk type with high performance to meet the application requirements for low latency, high read and write speed, and large

throughput. **Figure 4-3** shows the architecture in these scenarios. Disks with high performance service levels can meet the following performance requirements:

- The delay ranges from 1 ms to 3 ms.
- The performance ranges from 500 IOPS/TB to 4000 IOPS/TB.
- Typical configuration 1: OceanStor 6800 V5 is selected for backend storage, 50 1.92 TB, 3.84 TB, or 7.68 TB SSDs are configured for every dual-controller, and RAID 5 is configured. A maximum of eight controllers and 200 disks (300 TB, 600 TB, or 1200 TB) are configured for a single system.
- Typical configuration 2: Huawei Distributed Block Storage is selected for backend storage. RH2288H V5 servers are used. Twelve 4 TB, 6 TB, 8 TB, or 10 TB SATA disks are configured. Three-duplicate mode is adopted. One 1.6 TB or 3.2 TB SSD is configured. The total available space on each node is about 15.2 TB, 22.8 TB, 30.4 TB, or 38 TB.

Figure 4-3 Architecture in the data warehouse scenario



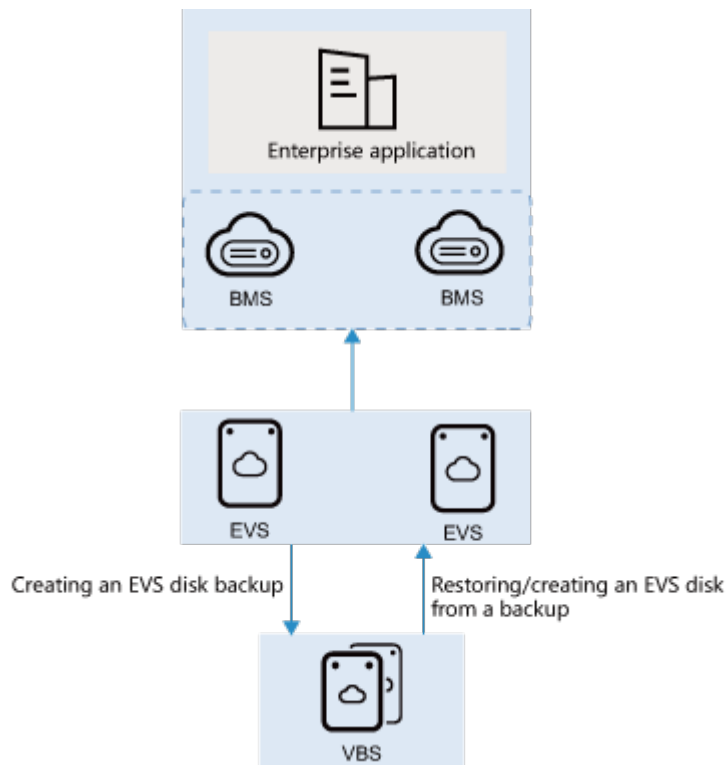
Enterprise Application System

Mission-critical applications of enterprises are deployed in these scenarios. These scenarios, such as common databases, application VMs, and middleware VMs, require relatively low performance but rich enterprise-class features. It is recommended that you use the disk type with medium performance. **Figure 4-4** shows the architecture in these scenarios. Disks with medium performance service levels can meet the following performance requirements:

- The delay ranges from 3 ms to 10 ms.
- The performance ranges from 250 IOPS/TB to 1000 IOPS/TB.
- Typical configuration 1: OceanStor 5500 V5 is selected for backend storage. Fewer than 250 disks are configured for every dual-controller, including ten 1.92 TB, 3.84 TB, or 7.68 TB SSDs and fewer than 240 600 GB, 1.2 TB, or 1.8 TB SAS disks. RAID 5 is configured. A single system supports a maximum of six controllers and 750 disks (360 TB, 720 TB, or 1116 TB).

- Typical configuration 2: Huawei Distributed Block Storage is selected for backend storage. 5288 V3 servers are used. Thirty-six 2 TB, 4 TB, 6 TB, or 8 TB SATA disks are configured. Three-duplicate mode is adopted. Two 1.6 TB or 3.2 TB SSDs are configured. The total available space on each node is about 22.8 TB, 45.6 TB, 68.4 TB, or 91.2 TB.

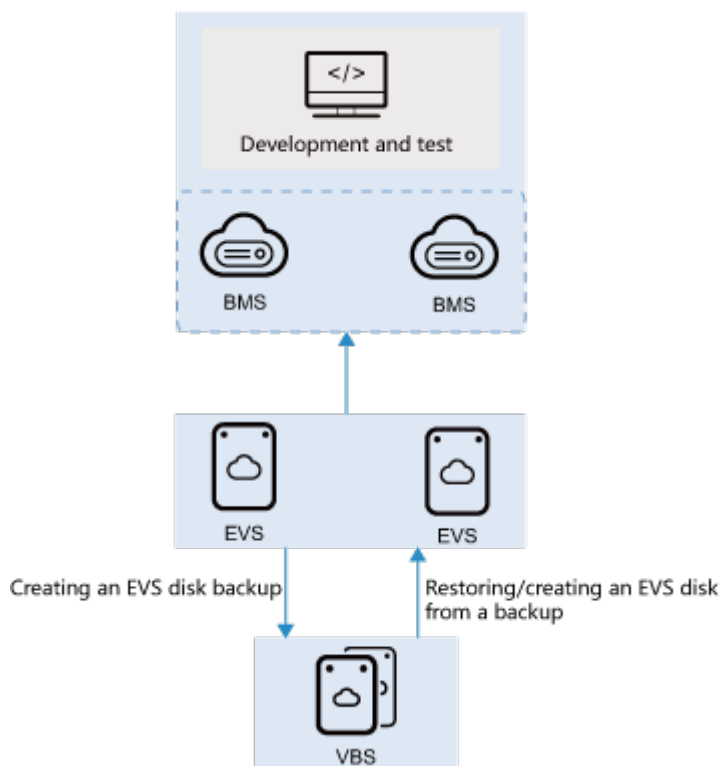
Figure 4-4 Architecture in the enterprise application system scenario



Development and Test

In these scenarios, development and test applications are deployed. It is recommended that you use the disk type with common performance to meet the requirements of development, test, deployment, and O&M. [Figure 4-5](#) shows the architecture in these scenarios. Disks with common performance service levels can meet the following performance requirements:

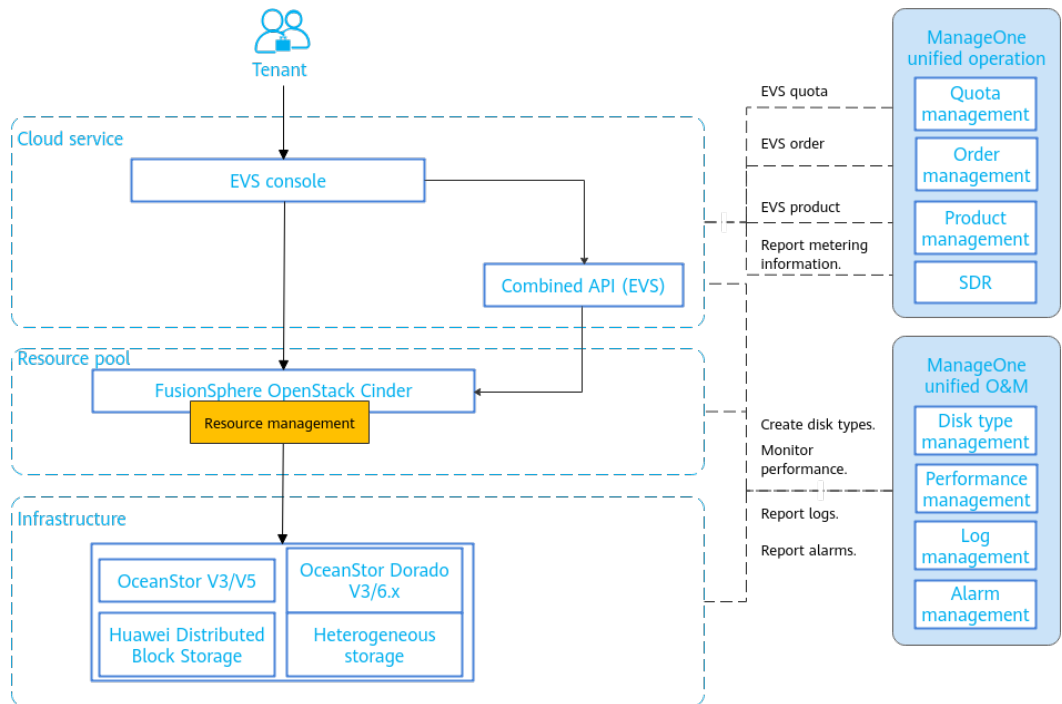
- The delay ranges from 10 ms to 20 ms.
- The performance ranges from 5 IOPS/TB to 25 IOPS/TB.
- Typical configuration: OceanStor 5300 V5 is selected for backend storage. Fewer than 396 disks (2 TB/4 TB/6 TB/8 TB/10 TB NL-SAS disks) are configured for every dual controllers. RAID 6 is configured. A single system supports a maximum of two controllers (612 TB/1224 TB/1840 TB/2460 TB/3060 TB).

Figure 4-5 Architecture in the development and test scenario

4.1.4 Implementation Principles

Architecture

EVS includes components such as the EVS console, EVS service API, FusionSphere OpenStack Cinder, and storage device. [Figure 4-6](#) shows the logical architecture of EVS.

Figure 4-6 Logical architecture of EVS**Table 4-2** EVS component description

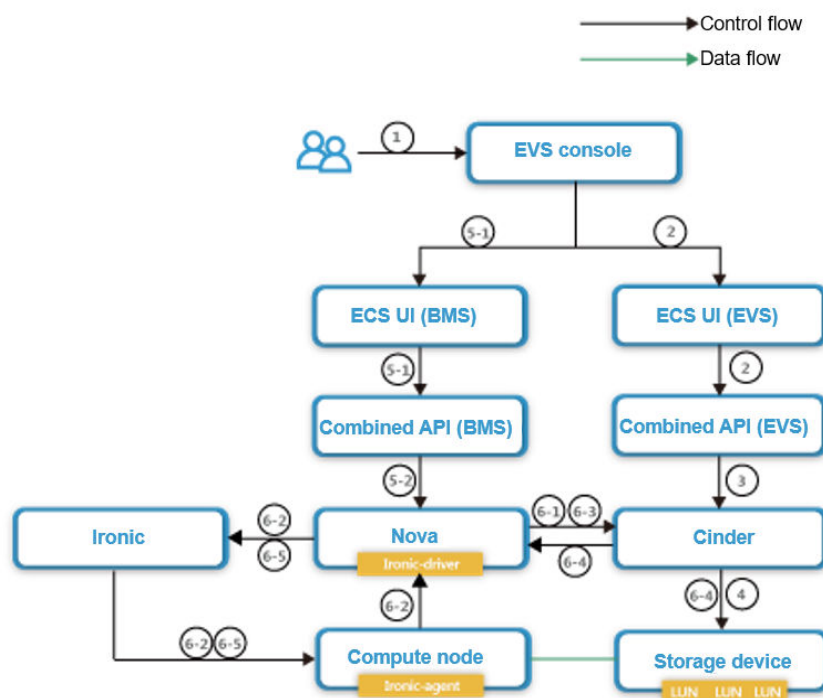
Component Name	Description
EVS console	The EVS console provides tenants with an entry to EVS. Tenants can apply for EVS disks on the console.
Combined API (EVS)	The EVS service API encapsulates or combines the logic based on the native Cinder interface to implement some EVS functions. The EVS service API can be invoked by the EVS console or tenants.
FusionSphere OpenStack Cinder	FusionSphere OpenStack Cinder provides persistent block storage to manage block storage resources. It is mainly used to create disk types in EVS. Disks are created on the storage device and attached to ECSs or BMSs.
Infrastructure	Infrastructure refers to the physical storage device that provides block storage based on physical resources. The following storage devices can function as the backend storage of EVS: Huawei SAN storage (OceanStor V3/V5 and OceanStor Dorado V3/6.x), Huawei Distributed Block Storage, and heterogeneous storage (such as HP 3PAR 8000 series).
ManageOne unified operation	ManageOne unified operation provides quota management, order management, product management, and resource metering and charging for EVS.

Component Name	Description
ManageOne unified O&M	ManageOne unified O&M provides disk type management, performance monitoring, logging, and alarm reporting for EVS.

Workflow

Figure 4-7 shows the workflow for EVS to provision EVS disks and attach EVS disks to BMSs.

Figure 4-7 EVS workflow



1. A VDC administrator or VDC operator applies for storage resources on the EVS console.
2. The EVS console sends the request to Combined API (EVS) through ECS UI (EVS).
3. Combined API distributes the request to Cinder.
4. Cinder creates volumes in the storage pool based on the policy for applying for storage resources. Cinder includes the following components:
 - Cinder API: receives external requests.
 - Cinder Scheduler: selects a proper backend storage server and specifies the storage server where the created volume resides.
 - Cinder Volume: connects to various storage device drivers and delivers requests to specific storage devices.
5. The VDC administrator or VDC operator attaches the applied storage resources to BMSs through the EVS console.

- a. The EVS console sends the request to Combined API (BMS) through ECS UI (BMS).
 - b. Combined API distributes the request to Nova.
6. Nova instructs Cinder to attach EVS disks.
 - a. Nova obtains EVS disk information and instructs Cinder to reserve EVS disks.
 - b. Nova uses the Ironic driver and ironic-agent to obtain information about the initiator of the physical machine.
 - c. Nova transmits initiator information to Cinder.
 - d. Cinder instructs the storage array to map the initiator and target and returns the Nova target information.
 - e. Nova completes the attachment task.

4.1.5 Related Services

Figure 4-8 shows the dependencies between EVS and other cloud services. **Table 4-3** provides more details.

Figure 4-8 Relationship between the EVS service and other cloud services

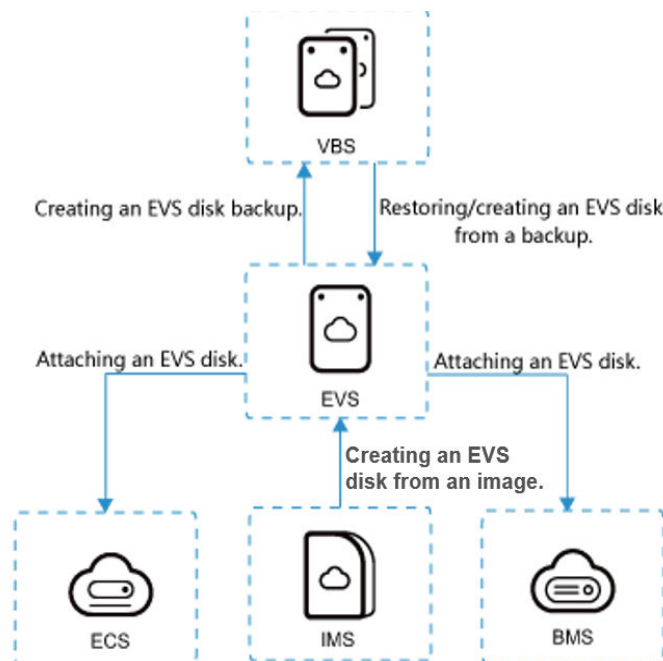


Table 4-3 Dependencies between EVS and other cloud services

Service Name	Description
ECS	You can attach EVS disks to ECSs to provide scalable block storage.
BMS	You can attach SCSI EVS disks to BMSs to provide scalable block storage.

Service Name	Description
VBS	Volume Backup Service (VBS) can be used to create backups for EVS disks. EVS disk data can be restored using the backups. Backups can be used to create EVS disks.
IMS	EVS can be used to create data disks from data disk images and system disks from system disk images. Image Management Service (IMS) can be used to create data disk images or system disk images.

4.1.6 Key Metrics

[Table 4-4](#) lists the key metrics of EVS.

Table 4-4 Key metrics of the EVS service

Item	Metric
Maximum number of EVS disks that you can obtain at a time	This metric is related to the EVS disk quota. <ul style="list-style-type: none">If the number of EVS disks in the quota is greater than 100, a maximum number of 100 EVS disks can be applied for each time.If the number of EVS disks in a quota is less than 100, the maximum number of EVS disks that can be applied for each time is equal to the quota quantity.
Maximum number of instances to which a shared disk can be attached simultaneously	16 If Huawei SAN storage is used as backend storage and the storage version is earlier than V300R006C50, a shared disk can be attached to fewer than eight instances simultaneously.
Maximum number of snapshots that can be created for an EVS disk	32 (recommended) This item is related to the backend storage type. The maximum number of snapshots that can be created varies with the backend storage type. For details, see the product documentation of the corresponding backend storage.

4.1.7 Restrictions

Before using EVS, learn the restrictions described in [Table 4-5](#).

Table 4-5 Restrictions on EVS

Item	Restrictions
Backend storage	<ul style="list-style-type: none">Supported Huawei storage devices include enterprise storage OceanStor V3/V5 and OceanStor Dorado V3, and Huawei Distributed Block Storage. You can visit HUAWEI CLOUD Stack Information Center to query the specific storage models and versions.For details about the supported heterogeneous storage types and versions, visit HUAWEI CLOUD Stack Information Center.When OceanStor Dorado V3/6.x or heterogeneous storage is used as the backend storage, AZs with the virtualization capability of Xen cannot be accessed but AZs with the virtualization capability of KVM or IroniC can be accessed.An AZ can contain multiple Huawei SAN storage devices, such as OceanStor V3/V5 and OceanStor Dorado V3/6.x. Huawei Distributed Block Storage, Huawei SAN storage, and heterogeneous storage cannot be used in one AZ.Only one set of Huawei Distributed Block Storage can be deployed in an AZ.FC SAN and IP SAN protocols cannot be used in the same AZ at the same time.It is recommended that a disk type contains only one type of backend storage to ensure that backend storage has the same performance.If the backend storage is FusionStorage 8.0.0 or later, deduplication and compression are enabled by default. Therefore, the provisioned disks have the deduplication and compression functions.If the backend storage is FusionStorage 8.0.0 or later and the type of the access storage pool is self-encrypting, the provisioned disks support data encryption.If the backend storage is OceanStor V500R007C60SPC300 or later, or OceanStor Dorado 6.0.0 or later, set Max. Sessions per User to 0 on DeviceManager, indicating that the number of sessions is not limited.If the backend storage is OceanStor Dorado 6.x, manually disable the recycle bin function. For details, see "Configure" > "Basic Storage Service Configuration Guide for Block" > "Managing Basic Storage Services" > "Managing LUNs" > "Managing the Recycle Bin" > "Configuring the Recycle Bin" in OceanStor Dorado 8000 and Dorado 18000 Product Documentation.

Item	Restrictions
Applying for an EVS disk	<ul style="list-style-type: none">• The maximum capacity of a single disk is 64 TB.• When the backend storage type is heterogeneous: Blank EVS disks can be created and EVS disks can be created from an existing disk or a snapshot. EVS disks cannot be created from a backup and the disk type cannot be VBD or shared.• EVS disks cannot be used as system disks on BMSs.• Shared disks can be used as data disks and cannot be used as system disks.• When you use an existing disk to create a disk, the restrictions are as follows:<ul style="list-style-type: none">– If backend storage is one of OceanStor V3/V5 series and you use an existing disk to create a disk, it is necessary for the administrator to import the license for HyperCopy in advance on the device side.– If the backend storage type is OceanStor Dorado V3, the version must be Dorado V300R001C21 or later.– If the selected service has preset disk capacity and disk type, you can choose only a disk whose capacity is less than or equal to the preset disk capacity as the source disk, and the disk type of the disk must be the same as the preset disk type.– The source disk and the disk to be created must be in the same AZ.– When you use an existing disk to create a disk, the disk capacity can be configured but must be greater than or equal to that of the source disk. The disk type and disk mode cannot be changed, which are the same as those of the source disk.• When creating a disk using a snapshot, if the backend storage type is OceanStor V3/V5 series, the administrator needs to import the license of the HyperCopy feature on the device in advance.• Snapshots in one AZ cannot be used to create disks in another AZ.
Attaching an EVS disk	<ul style="list-style-type: none">• Only SCSI EVS disks can be attached to BMSs.• Regardless if a shared EVS disk or non-shared EVS disk is attached to an instance, the EVS disk and the instance must be in the same AZ.• An EVS disk cannot be attached to an instance that has expired.• An EVS disk cannot be attached to an instance that has been soft deleted.• An EVS disk cannot be attached to an instance that has been stopped.• If the ECS to which an EVS disk belongs has not been created, the EVS disk cannot be attached to another ECS.

Item	Restrictions
Creating a snapshot	<ul style="list-style-type: none">• If backend storage is one of OceanStor V3/V5 or OceanStor Dorado V3 series, it is necessary for the administrator to import the license for HyperSnap in advance on the device side.• Snapshots can be created only for disks in the Available or In-use state.• A snapshot name cannot be the same as the prefix of the temporary snapshot created by the backup service, such as Volume Backup Service (VBS) and Cloud Server Backup Service (CSBS), or the disaster recovery service, such as Cloud Server Disaster Recovery (CSDR), Cloud Server High Availability (CSHA), and VHA.• Snapshots created using the EVS console consume the capacity quota instead of quantity quota of EVS disks.• Snapshots created using the EVS console and temporary snapshots created by the DR and backup service (VBS, CSBS, CSDR, CSHA, or VHA) consume backend storage capacity. If a large number of snapshots are created, contact the administrator to set the thin provisioning ratio of backend storage to a large value, preventing EVS disk provisioning failures caused by excessive snapshots.• Temporary snapshots created by the backup service (VBS or CSBS) or the disaster recovery service (CSDR, CSHA, or VHA) do not consume EVS disk quotas.• No snapshots can be created for disks that have expired.• No snapshots can be created for disks that have been soft deleted.• If a task for creating a snapshot fails, the task is automatically deleted.• If backend storage of the disk is heterogeneous storage, snapshots can be created.
Rolling back a disk from a snapshot	<ul style="list-style-type: none">• A temporary snapshot created by the backup service (VBS or CSBS) or the disaster recovery service (CSDR, CSHA, or VHA) cannot be used to roll back the EVS disk.• Snapshots created for disks having any DR service (CSDR/CSHA/VHA) configured cannot be rolled back.• A snapshot can be used to roll back its source EVS disk, and cannot be used to roll back any other EVS disk.• When the source disk of a snapshot is in the recycle bin, EVS disk rollback from the snapshot is not supported.• If backend storage of the disk is heterogeneous storage (excluding XSKY), EVS disk rollback from a snapshot is not supported.
Creating a backup	<ul style="list-style-type: none">• Only disks in the Available or In-use state can be backed up.• If backend storage of the disk is heterogeneous storage, backups cannot be created.


Item	Restrictions
Expanding the capacity of an EVS disk	<ul style="list-style-type: none">• When you expand the capacity of a disk online, the instance to which the disk is attached must be in the Running or Stopped state.• Shared EVS disks do not support online capacity expansion, that is, the capacity of a shared EVS disk can be expanded only when the disk is in the Available state.• The capacity of a disk configured with the disaster recovery service (CSHA, CSDR, or VHA) cannot be expanded.• When the backend storage is Huawei SAN storage (OceanStor V3/V5 series and OceanStor Dorado V3/6.x series) or heterogeneous storage, if the EVS disk has snapshots, capacity expansion is not supported. When the backend storage is Huawei Distributed Block Storage, capacity expansion can be performed for an EVS disk with snapshots.• If backend storage of the disk is heterogeneous storage, online capacity expansion is not supported while offline capacity expansion is supported.
Changing the disk type	<ul style="list-style-type: none">• Changing the disk type is supported when the backend storage is OceanStor V3/V5, OceanStor Dorado V3, or Huawei Distributed Block Storage.• If the backend storage is OceanStor V3/V5 or OceanStor Dorado V3, the disk type can be changed between different storage pools in the same storage system. If the backend storage is Huawei Distributed Block Storage, the disk type can be changed only in the same storage pool.• If the backend storage is Huawei Distributed Block Storage, the disk type can be changed only by modifying the QoS attribute.• The administrator needs to import the SmartMigration license on the device in advance if the backend storage is OceanStor V3/V5 or OceanStor Dorado V3.• When changing the disk type, if the backend storage is OceanStor Dorado 6.x, the administrator needs to check whether the SmartMigration license has been imported to the device in advance. (The basic software package of OceanStor Dorado 6.x contains the SmartMigration license.)• You can change the type of the EVS disk only in the Available or In-use state.• If a disk has snapshots or is configured with the backup service (VBS or CSBS) or the disaster recovery service (CSDR, CSHA, or VHA), the disk type cannot be changed.• If backend storage of a disk is heterogeneous storage, the disk type cannot be changed.

Item	Restrictions
Extending the validity period of an EVS disk	<ul style="list-style-type: none">• If an EVS disk is created with an instance, the validity period of the EVS disk is unlimited.• If the validity period of an EVS disk is unlimited, the validity period cannot be extended.• You can extend the EVS disk validity period only when the disk is in the available or In-use status.• If an EVS disk has expired, its snapshot cannot be used to roll back the EVS disk or create an EVS disk. To continue using this EVS disk, extend its validity period.• When an EVS disk expires, its data will not be deleted. You can continue using this EVS disk after extending its validity period.
Detaching an EVS disk	<ul style="list-style-type: none">• Data disks can be detached online, that is, data disks can be detached from BMSs in the running state.• Before detaching a disk online from an instance running Windows, log in to the instance to perform the offline operation and confirm that the disk is not being read and written. Otherwise, the disk will fail to be detached.• Before detaching a disk online from an instance running Linux, log in to the instance, run the umount command to cancel the relationship between the disk and the file system, and confirm that the disk is not being read and written. Otherwise, the disk will fail to be detached.
Deleting an EVS disk	<ul style="list-style-type: none">• A disk that has been attached to an instance cannot be deleted.• If a disk has been configured with the disaster recovery service (CSDR, CSHA, or VHA), the disk cannot be deleted.• If an EVS disk has a snapshot, the EVS disk can be soft deleted only when the snapshot is in the Available or Error state.• When an EVS disk is permanently deleted, all snapshots of the EVS disk are also deleted.• A shared disk to be deleted must have been detached from all instances.• If the ECS to which an EVS disk belongs has not been created, the EVS disk cannot be deleted.
Deleting a snapshot	<ul style="list-style-type: none">• Users are allowed to delete a temporary snapshot created by the backup service (VBS or CSBS). After the snapshot is deleted, if users want to back up the EVS disk corresponding to the snapshot, full backup is performed for the first time.• Temporary snapshots created by the disaster recovery service (CSDR, CSHA, or VHA) cannot be deleted.• You can delete a snapshot only when its state is Available, Deletion failed, or Error.

Item	Restrictions
Creating and Associating a QoS	<ul style="list-style-type: none">• The QoS function is supported only in KVM and BMS virtualization scenarios.• IOPS and bandwidth upper limits can be set only when the backend storage is OceanStor V3/V5, OceanStor Dorado V3/6.x, or Huawei Distributed Block Storage.• The I/O priority can be set only when the backend storage is OceanStor V3/V5 or OceanStor Dorado V3/6.x.• A QoS policy cannot be associated with a disk type with disks provisioned.• One disk type can be associated with only one QoS policy. One QoS policy can be associated with multiple disk types.• Before creating a QoS policy, if the backend storage is Huawei SAN storage, check on OceanStor DeviceManager that the SmartQoS license has been activated.
Disk Migration	<ul style="list-style-type: none">• Advanced migration applies to Huawei SAN storage (OceanStor V3/V5 and OceanStor Dorado V3/6.x) and does not apply to Huawei Distributed Block Storage and heterogeneous storage. The source storage and target storage must be Huawei SAN storage and must meet the version requirements.• During migration, the source storage and target storage must be in the same AZ.• Only unattached disks can be migrated.• Disks with snapshots cannot be migrated.• Shared disks can be migrated.• Before the migration, check on OceanStor DeviceManager that the SmartMigration license has been activated in the target storage.• After the migration is complete, the disk has all features of the target disk type.• No more than three sets of source storage devices can be migrated to one set of target storage device. It is recommended that one set of source storage device be migrated to one target storage device.• During the migration, do not perform other operations on the disk.• The remaining capacity of the storage pool to which the disk to be migrated must be greater than 1% of the total capacity of the storage pool.• During disk migration, if a resource tag has been set for the source disk type, a resource tag must be set for the target disk type. Otherwise, disk migration is not supported.

4.1.8 Accessing and Using the Cloud Service

Two methods are available:

- Web UI
Log in to ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios) as a tenant, click  in the upper left corner of the page, select a region, and select the cloud service.
- API
Use this mode if you need to integrate the cloud service into a third-party system for secondary development. For details, see API reference of the service on **Operation Help Center**.

4.2 Related Concepts

4.2.1 Device Type

Definition

Device types of EVS disks are divided based on whether advanced SCSI commands are supported. The device type can be Virtual Block Device (VBD) or Small Computer System Interface (SCSI).

- VBD: EVS disks of this type support only basic SCSI read and write commands. They are usually used in common scenarios, for example, OA, tests, Linux clusters such as RHCS.
- SCSI: EVS disks of this type support transparent SCSI command transmission and allow the ECS operating system to directly access the underlying storage media. SCSI EVS disks support advanced SCSI commands (such as SCSI-3 persistent pre-lock) in addition to basic SCSI read and write commands. They can be used in cluster scenarios where data security is ensured by using the SCSI lock mechanism, such as the Windows MSCS cluster.

NOTE

For details about BMS OSs supported and BMS software required by SCSI EVS disks, see [Usage requirements on SCSI EVS disks](#).

Usage requirements on SCSI EVS disks

Currently, only SCSI EVS disks can be attached to BMSs. VBD EVS disks attached to BMSs will be used as SCSI EVS disks by default.

The BMS OS is preinstalled with the driver required for using SCSI EVS disks, and you do not need to install the driver.

4.2.2 Disk Type

Definition

A disk type can be selected during disk creation. A disk type represents backend storage devices used by a group of disks. You can divide disk types of EVS disks

based on backend storage types to meet different performance requirements of services.

Based on performance differences of backend storage used by disks, typical disk types and their application scenarios are as follows:

- Common performance: EVS disks of this type are suitable for scenarios that require large capacity, medium-level read and write speed, and relative fewer transactions, such as the scenario for deploying development and test applications.
- Medium performance: EVS disks of this type are suitable for scenarios that require common performance but rich enterprise-class features. They can be used in common databases, application VMs, and middleware VMs.
- High performance: EVS disks of this type are suitable for scenarios that require high performance, fast read and write speed, and large throughput, such as data warehouses.
- Ultra-high performance: EVS disks of this type are suitable for data-intensive scenarios that require very high I/O performance, such as NoSQL and relational databases.

Changing a Disk Type

When the read and write performance of the storage device where the upper-layer service resides no longer suits the service, you can change the disk type to alter the type of the storage device to change the read and write performance, meeting the requirements of varying service performance of the instance. Examples are as follows:

- When your service requires a higher read and write performance, you can migrate your service from disks created on low-speed storage media to disks created on high-speed storage media to improve the read and write performance.
- If the priority of the performance of a service degrades, you can migrate your service to disks created on low-performance storage media. This helps release storage resources for high-performance disks for other services.

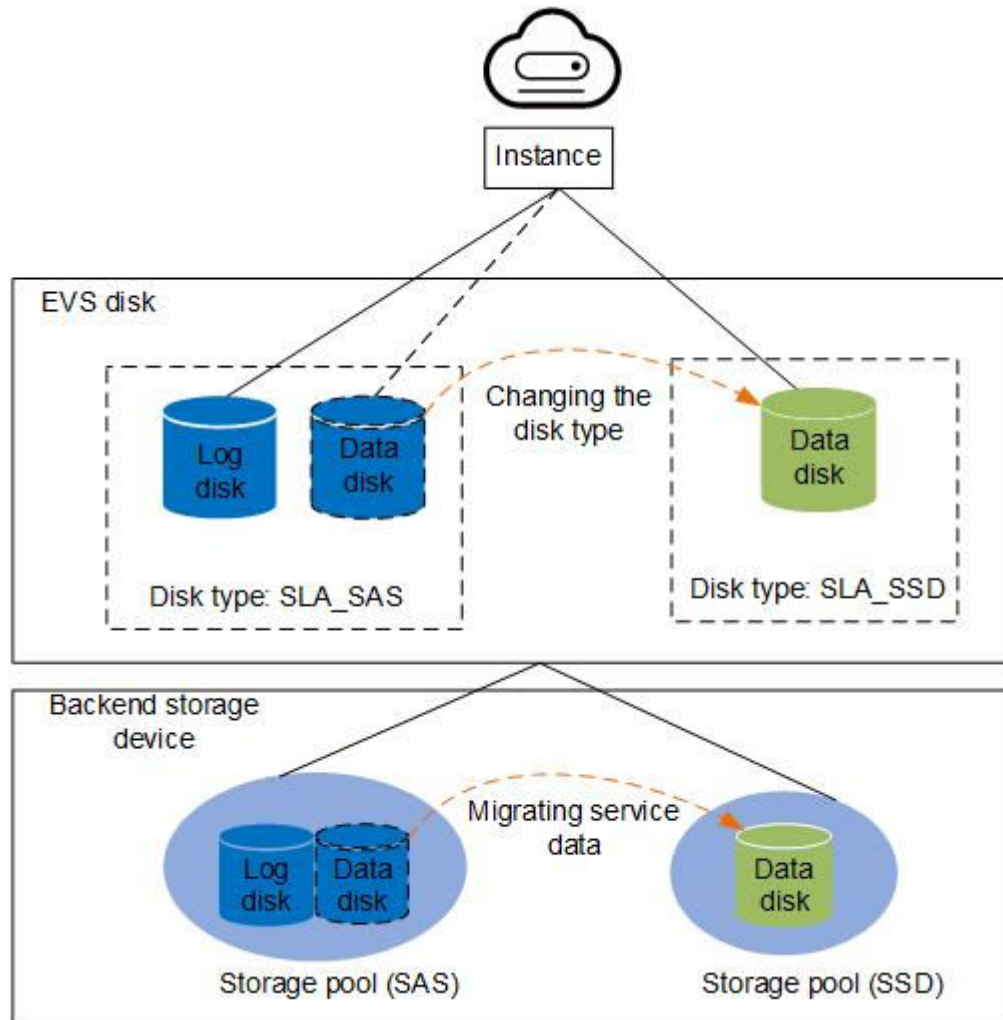
You can change the disk type of an in-use disk (a disk that has been attached to an instance). You can also detach a disk from the instance, and then change the disk type of the disk.

If you change the disk type of an in-use EVS disk, the service of the source EVS disk on the instance will be migrated to the destination EVS disk without interrupting host services. After service migration, the destination EVS disk replaces the source EVS disk to run the service, without any adverse impact on customer experience. However, when you change the disk type of an in-use EVS disk, the performance of the instance is adversely affected to some extent.

Figure 4-9 shows the implementation principle of changing a disk type. In the following figure, two disks are attached to an instance. One of the disks serves as a log disk, and the other serves as a data disk. The original disk type of the two disks is SLA_SAS. Because the service has a higher performance requirement on the data disk, the disk type of the data disk is changed from SLA_SAS to SLA_SSD, seamlessly migrating service data to a disk of the target disk type. The backend storage device performs service data migration. After service data migration, the

system automatically attaches the destination disk to the instance, without service interruption. In addition, the source disk will be deleted to release storage resources for other services.

Figure 4-9 Implementation principle of changing a disk type



4.2.3 EVS Disk Data Redundancy Protection

The backend storage of EVS disks supports multiple storage types, including Huawei SAN, Huawei Distributed Block Storage, and heterogeneous storage. A storage array provides data redundancy protection. When a physical device is faulty, data on the faulty device can be automatically restored and data can still be accessed. For example: Huawei SAN storage uses the RAID technology and allows one or two disks to fail at the same time. Huawei Distributed Block Storage uses the three-copy technology to create two identical copies for each piece of data. The data and copies are stored on different storage nodes. The Erasure Coding (EC) technology is added to OceanStor Pacific series 8.1.1. When the storage type is OceanStor Pacific series 8.1.1, EC is used as the default data redundancy mode. This section uses the EC technology as an example to describe EVS disk data redundancy protection. For more information about storage device redundancy protection, see the product documentation of the corresponding storage device.

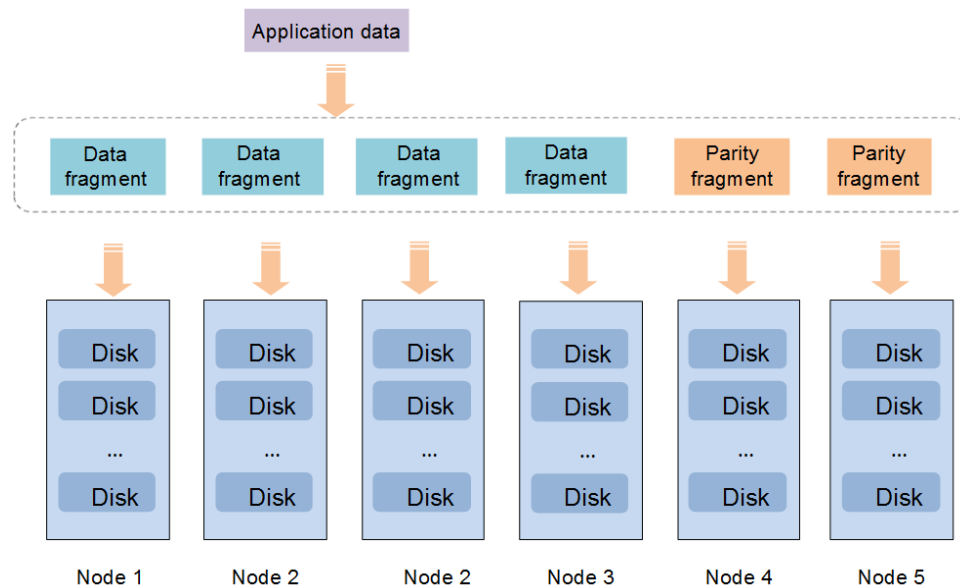
EC Principle

EC protection levels are expressed in N+M mode. N indicates the number of data fragments (data strips) and M indicates the number of parity fragments (parity strips). Huawei Distributed Block Storage supports N+2, N+3, and N+4 protection levels and allows a maximum of four storage nodes or cabinets to fail at the same time. The minimum node quantity and corresponding disk utilization rate vary by protection level. For details, see "General Information" > "Product Description" > "Functions and Features" > "Key Features" > "Block Service" > "EC and Multi-Copy" in *OceanStor Pacific Series 8.1.1 Product Documentation*.

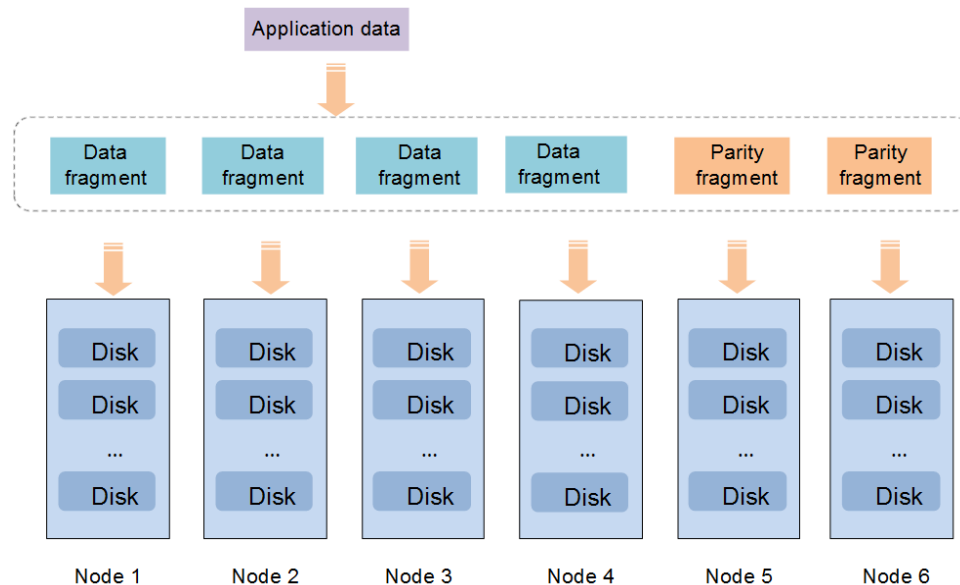
The following uses the N+2 level (N is 4) as an example to describe the EC implementation process.

- **Figure 4-10** shows how data protection is implemented when the number of storage nodes is greater than or equal to N+2. In this example, two storage nodes or cabinets can be faulty at the same time.
- **Figure 4-11** shows how data protection is implemented when the number of storage nodes is greater than or equal to (N+2)/2 and less than N+2. In this example, one storage node or cabinet is allowed to fail at the same time, and two disks or storage nodes are allowed to fail at the same time on a single storage node.

Figure 4-10 Huawei Distributed Block Storage EC technology (N+M configuration mode)



1. Divide written data into four data fragments.
2. Calculate and generate two parity fragments based on the four data fragments.
3. Write all the data fragments and parity fragments to the six storage nodes in the redundancy ratio mode.

Figure 4-11 Huawei Distributed Block Storage EC technology (N+M:1 configuration mode)

1. Divide written data into four data fragments.
2. Calculate and generate two parity fragments based on the four data fragments.
3. Write all the data fragments and parity fragments to the three storage nodes in the redundancy ratio mode.

Huawei Distributed Block Storage EC data fragmentation improves data read and write performance and ensures high disk utilization. If no more than M disks in a Huawei Distributed Block Storage cluster are faulty, the system can quickly restore lost data by reconstructing data among nodes. This ensures high data reliability and service availability.

4.2.4 Shared Disk

In the traditional cluster architecture, multiple computing nodes need to access the same data. This ensures that when a single or multiple computing nodes are faulty, the HA cluster can continue providing services, which means that a faulty component will not cause service interruption. Therefore, important data files need to be stored on shared block system, and shared block storage is centrally managed using the cluster file system. When multiple frontend computing nodes access data, the data will be the same on the multiple computing nodes.

The shared disk is designed for the core service HA architecture of enterprise customers. The shared disk is suitable for scenarios that require shared block storage access in the share-everything architecture. The scenarios include the HA Oracle RAC database architecture for government, enterprise, and finance customers and the HA server cluster architecture.

Definition

Shared EVS disks are block storage devices that support concurrent read/write operations of multiple ECSs/BMSs. Shared EVS disks feature multiple attachments,

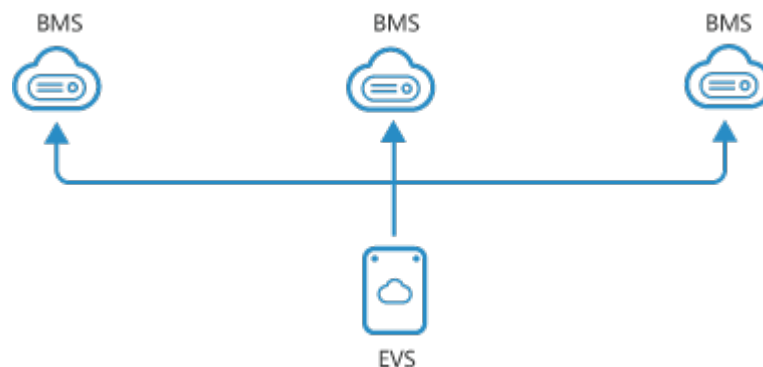
high-concurrency, high-performance, and high-reliability. A shared EVS disk can be attached to a maximum of 16 ECSs/BMSs. A non-shared EVS disk can be attached to only one ECS/BMS. This document uses BMS as an example, as shown in [Figure 4-12](#).

Currently, shared EVS disks can be used as data disks only and cannot be used as system disks.

You can use the EVS console to create VBD shared EVS disks or SCSI shared EVS disks. However, only SCSI EVS disks can be attached to BMSs. Therefore, you can attach only SCSI shared EVS disks to BMSs.

You can use the BMS console to create VBD shared EVS disks (default EVS disks) together with BMSs, and attach the VBD shared EVS disks to BMSs as data disks. VBD EVS disks attached to BMSs will be used as SCSI EVS disks by default.

Figure 4-12 Shared EVS disk

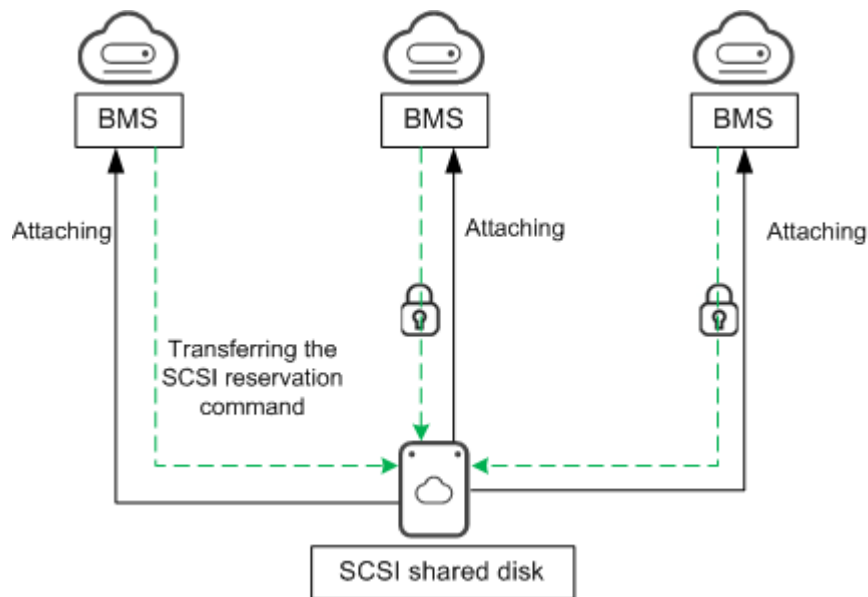


SCSI Reservation

SCSI shared EVS disks support SCSI reservation. If SCSI reservation is required for your applications, create SCSI shared EVS disks.

SCSI reservation is the basic mechanism for multiple hosts to use disks. In a shared storage environment, multiple service hosts may access a disk simultaneously. If multiple hosts perform the write operation on the disk at the same time, the disk does not know data from which host will be written first. To prevent this problem that may cause data damage, SCSI reservation is introduced.

[Figure 4-13](#) shows how SCSI reservation is implemented. When a SCSI shared disk is attached to multiple BMSs, if one of the BMSs sends a SCSI reservation command to the SCSI shared disk, the SCSI shared disk is locked for the other BMSs. In this case, the other BMSs cannot write data into the SCSI shared disk.

Figure 4-13 SCSI reservation implementation mechanism

Precautions for Using the Shared EVS Disk

A shared EVS disk is essentially the disk that can be attached to multiple instances 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 and write rules, such as the read and write sequence and meaning, between these servers are not defined, data read and write interference between servers or other unpredictable errors may occur.

Shared EVS disks provide block storage devices whose data can be randomly read or written and allows shared access. Shared EVS disks do not provide the cluster file system. You need to install the cluster file system to manage shared EVS disks.

If a shared EVS disk is attached to multiple instances but is managed using a common file system, disk space allocation conflict will occur and data files will be inconsistent. The details are as follows:

- Disk space allocation conflict

Suppose that a shared EVS disk is attached to multiple instances. When a process on instance A writes files into the shared EVS disk, it checks the file system and available disk space. After files are written into the shared EVS disk, instance A will change its own space allocation records, but will not change the space allocation records on the other instances. Therefore, when instance B attempts to write files to the shared EVS disk, it may allocate disk space addresses that have been allocated by instance A, resulting in disk space allocation conflict.

- Inconsistent data files

Suppose instance A reads data and records it in the cache. When another process on instance A accesses the same data, the process will read the data directly from the cache. If instance B changes the data, instance A will not know and will read the data from the cache. As a result, service data will be inconsistent on instance A and instance B.

Therefore, the proper method for using shared EVS disks is to use a cluster file system to centrally manage the block devices. The cluster file system can be Oracle RAC, Windows WSFC cluster, Linux RHCS cluster, Veritas VCS cluster, or CFS cluster application.

4.2.5 EVS Disk Snapshot

Definition

EVS disk snapshot is an important data recovery method that records the status of EVS disk data at a specific point in time. The snapshot created for an EVS disk at a certain point in time is independent from the life cycle of the EVS disk. The snapshot can be used to roll back and restore data of the EVS disk at the time when the snapshot was taken.

You can create an EVS disk from a snapshot. The created EVS disk contains the data of the snapshot, and is a precise copy of the source EVS disk. An EVS disk created from a snapshot does not need to be partitioned or formatted, and no file system needs to be created. When the EVS disk is attached to an instance, the EVS disk can read and write data. Therefore, the snapshot is an important way of sharing and migrating data.

Comparison Between EVS Snapshot and VBS

Both EVS snapshot and backup can protect EVS data. [Table 4-6](#) describes the differences between them.

Table 4-6 Comparison between EVS snapshot and VBS

Item	EVS Disk Snapshot	VBS
DR methods	A snapshot records the status of an EVS disk at a specific point in time. You can roll back and restore data of the EVS disk at the time when the snapshot was taken. A snapshot is not an actual disk data copy. If the disk is physically damaged, data cannot be restored using the snapshot rollback function. In this case, backup can be used.	A backup is a copy of EVS disk data at a certain point in time Backup is implemented based on the snapshot technology and snapshot comparison technology. A disk is backed up in full backup mode only when they are initially backed up. After the full backup, the disk is backed up in incremental backup mode.
Storage mode	Snapshot data is stored with disk data, consuming the capacity quota instead of quantity quota of EVS disks. The storage space consumed by snapshots is small. Therefore, snapshots can be created faster than backups.	The backup data is stored in the backup storage. It consumes no quota of EVS disks.

Item	EVS Disk Snapshot	VBS
Periodic snapshot/ backup	Periodic snapshots are not supported. Snapshots have to be created manually.	Periodic backup is supported. You can create a backup policy to automatically back up disks based on the backup policy.
DR range	Snapshots are region-specific. Snapshots, and the corresponding source disks belong to the same AZ. When you create an EVS disk using a snapshot, the new EVS disk and the snapshot are in the same AZ.	EVS disks and backups can be in different AZs. Backup copies can be used to restore EVS disk data across AZs.
Service restoration	You can use the snapshot to roll back the disk to the time when the snapshot was created, or create an EVS disk using the snapshot to retrieve the data at the time when the snapshot was created.	You can restore the backup to the original disk or another disk, or create an EVS disk from the backup to retrieve the data at the backup time point.
Deployment mode	The snapshot function can be used after EVS is deployed.	The advanced cloud service VBS must be deployed.

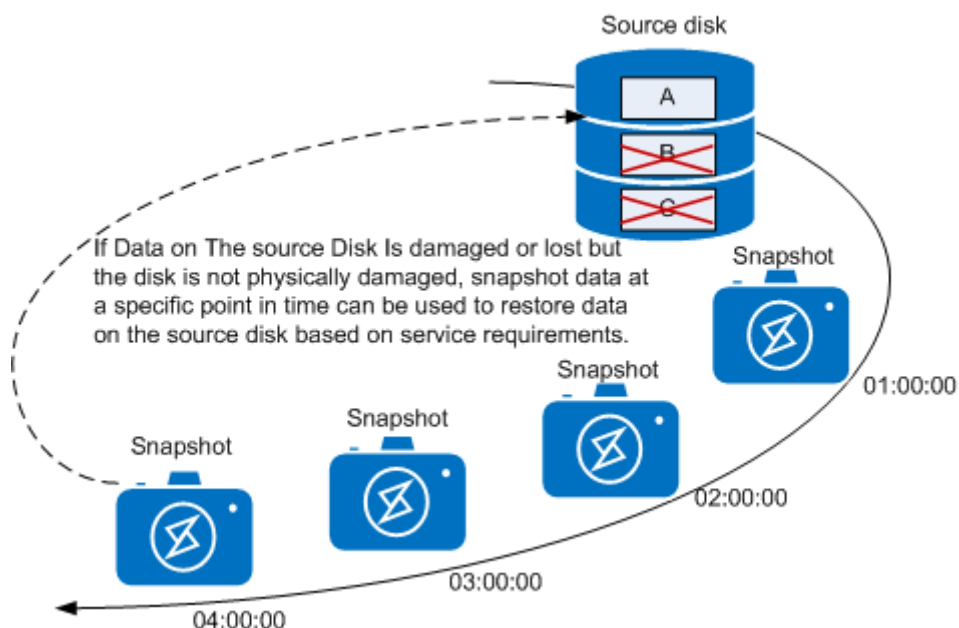
Application Scenarios

The snapshot is a convenient and efficient means of data protection, and it is recommended that you use this means of data protection in the following scenarios:

- Routine data backup and restoration

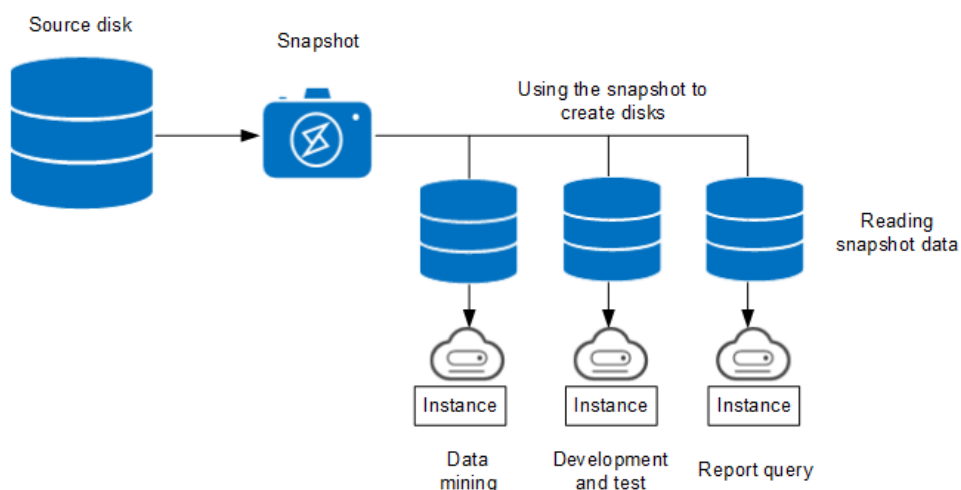
Snapshots are used to periodically back up important service data on system disks and data disks to prevent data loss caused by misoperations, attacks, or viruses.

When data loss or data inconsistency occurs on an EVS disk due to misoperations, viruses, or hacker attacks, you can use a snapshot to restore a previous normal status of the EVS disk. In addition, you are advised to create disk snapshots before a big change (such as application software upgrade and service data migration). If the operation fails, you can roll back the snapshots to restore service data, as shown in [Figure 4-14](#).

Figure 4-14 Using snapshots for routine data backup and restoration

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, such as data mining, report query, and development and test. This method protects the initial data and creates disks rapidly, meeting the diversified service data requirements. [Figure 4-15](#) shows the procedure for using a snapshot to deploy multiple services.

Figure 4-15 Using a snapshot to deploy multiple services

Recommendation Policies

You can choose an appropriate snapshot policy and retention policy based on your service type. Recommended policies are as follows:

- **Core services:** For core services that require very high Recovery Point Objective (RPO), it is recommended that data be backed up every several hours and snapshots be retained for one day.
- **Production services:** For production services, it is recommended that data be backed up every week and snapshots be retained for one month.
- **Archiving services:** For archiving services, it is recommended that data be backed up every month and snapshots be retained for one year.

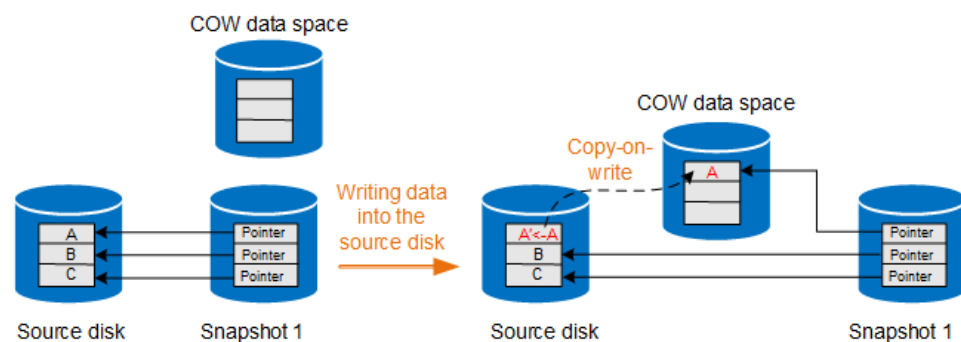
Implementation Principles

The snapshot implementation principle varies with the type of backend storage where the disk resides. Snapshot implementation principles for different backend storage types are described as follows:

- **OceanStor V3 or OceanStor V5 series as backend storage**

A snapshot is a copy of source disk data, which is generated at a specific time. A snapshot consists of a source disk, Copy-on-Write (COW) data space, and snapshot data. Snapshots are implemented using the mapping table and COW technology. [Figure 4-16](#) shows the snapshot implementation principle.

Figure 4-16 Snapshot implementation principle (OceanStor V3 or OceanStor V5 series as backend storage)



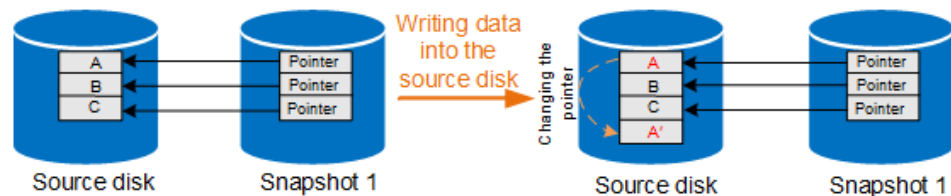
- Before creating a snapshot: When no snapshot is created for a disk, the procedure for writing data into the disk is the same as the procedure for writing data into other disks. Data changes will be directly written into disk data blocks, overwriting the original data, and the original data will not be retained.
- After creating a snapshot: After a snapshot is created, a data copy that is identical to the source disk is generated. In this step, the backend storage system dynamically allocates COW data space in the storage pool where the source disk resides, and automatically generates a snapshot. The pointer of the snapshot points to the storage location of source disk data.
- Writing data into the source disk: When an instance sends a request to write data into the source disk, the backend storage system will not write the new data immediately. Instead, the backend storage system employs the COW mechanism to copy the original data from the source disk to the COW data space, modifies the mapping in the mapping table, and writes the new data to the source disk. As shown in [Figure 4-16](#), when data A of the source disk needs to be changed, data A will be copied to the COW data space, and then the snapshot pointer will be changed to

point to the storage location of data A in the COW data space. Finally, data A' will be written into the source disk.

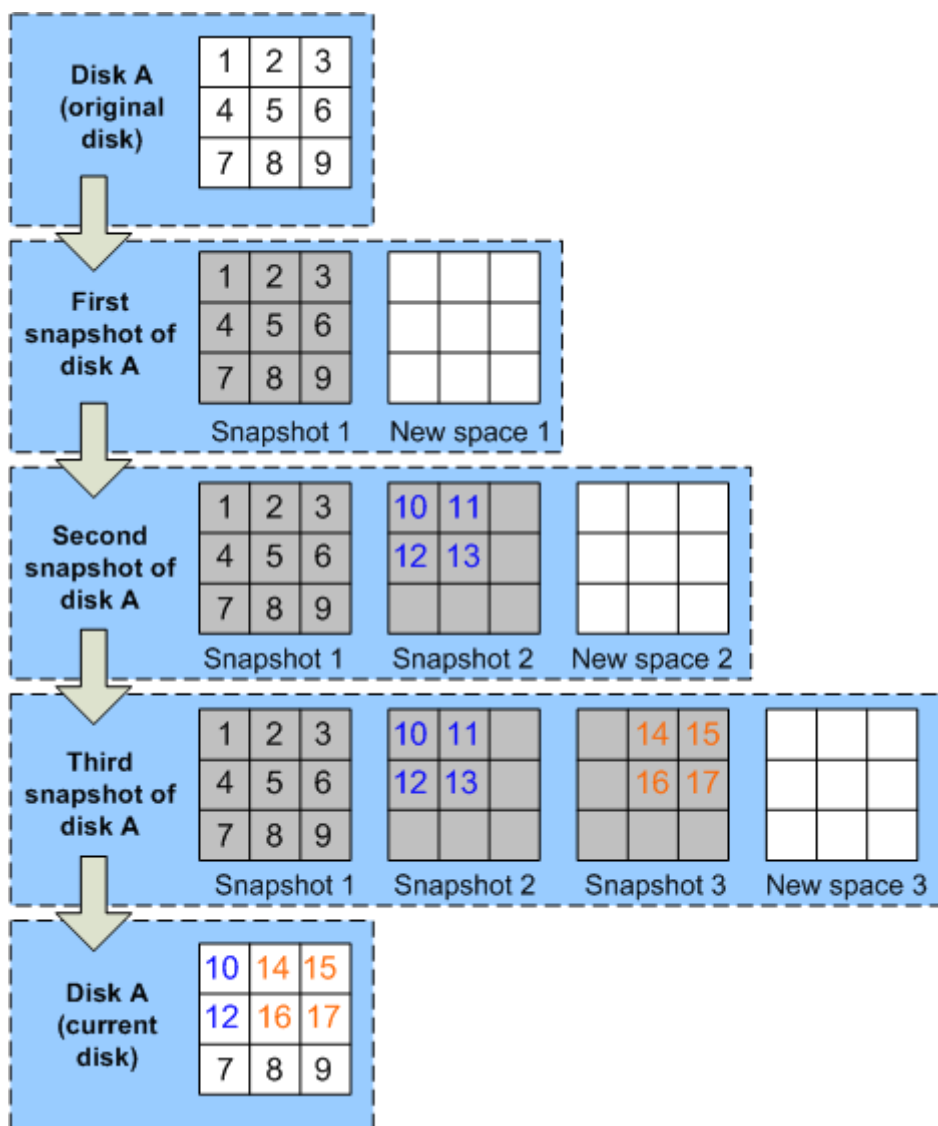
- **Dorado V3 series as backend storage**

The core technology in snapshot implementation is Redirect-on-Write (ROW). **Figure 4-17** shows the snapshot implementation principle.

Figure 4-17 Snapshot implementation principle (Dorado V3 series as backend storage)

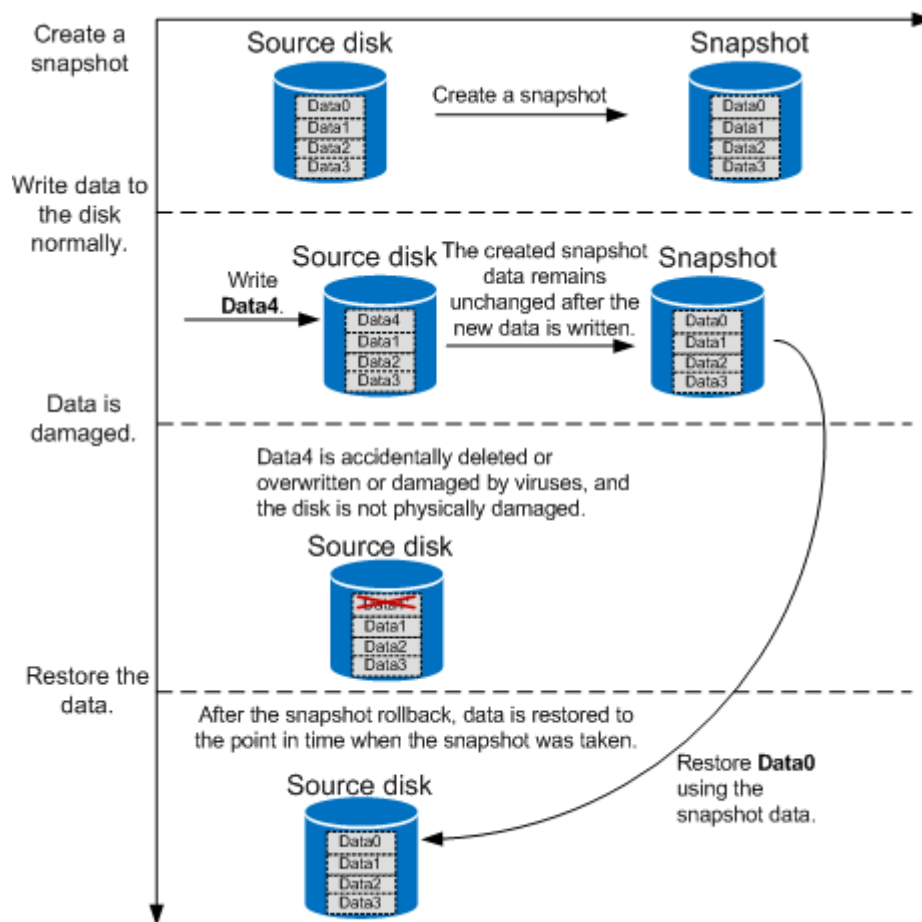


- Before creating a snapshot: When no snapshot is created for a disk, the procedure for writing data into the disk is the same as the procedure for writing data into other disks. Data changes will be directly written into disk data blocks, overwriting the original data, and the original data will not be retained.
 - After creating a snapshot: After a snapshot is created, a data copy that is identical to the source disk is generated. In this step, the backend storage system copies the pointer of the source disk to the snapshot, and the pointer of the snapshot points to the storage location of source disk data.
 - Writing data into the source disk: When an instance sends a request to write data into the source disk after a snapshot is created, the storage system uses the ROW technology to save the new data to a new location and changes the pointer of the source disk to point to the storage location of the new data. The pointer of the snapshot still points to the storage location of the original data. The source disk data at the time when the snapshot was created is saved. As shown in **Figure 4-17**, when data A of the source disk needs to be changed, data A' (new data) will be written into a new location, and the pointer of the source disk will be changed to point to the storage location of data A'. The pointer of the snapshot still points to the storage location of data A (original data).
- **Huawei Distributed Block Storage as backend storage**
Snapshot data is based on the Distributed Hash Table (DHT) mechanism. **Figure 4-18** shows the snapshot implementation principle.

Figure 4-18 Snapshot implementation principle (Huawei Distributed Block Storage as backend storage)

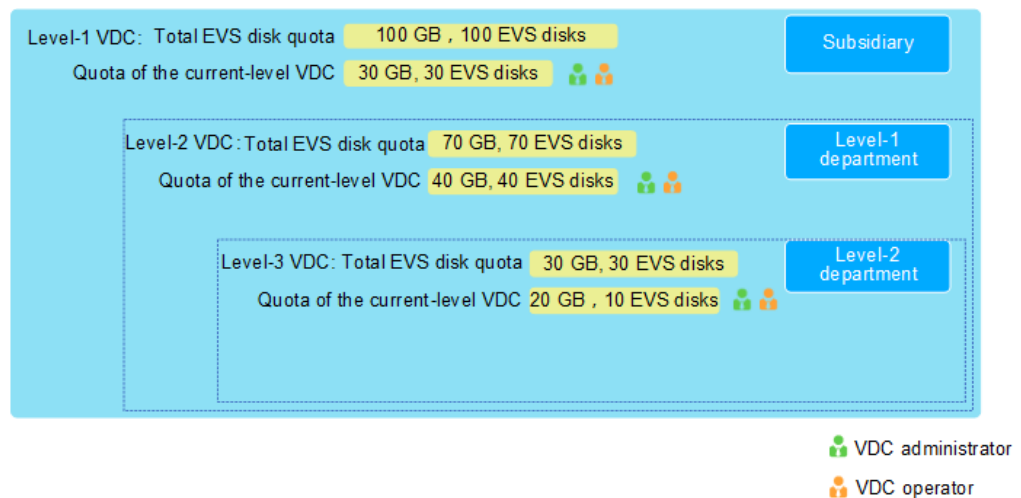
Rolling Back a Disk from a Snapshot

Snapshot rollback is a mechanism for quickly restoring data on the source disk by using the snapshot of the source disk at a certain point in time. If the data on the source disk is accidentally deleted, damaged, or infected by viruses and the source disk is not physically damaged, you can use the snapshot rollback function to quickly restore data on the source disk at the point in time when the snapshot was taken, reducing the amount of data lost. [Figure 4-19](#) shows snapshot rollback process.

Figure 4-19 Snapshot rollback

4.2.6 EVS Disk Quota

A Quota is a resource management and control technology that limits the maximum number of resources (including the resource capacity and number of resources) that can be used by a single VDC, preventing resources from being overused by users in some VDCs and affecting other VDCs. When creating a level-1 VDC, the operation administrator can set the total quota (capacity and quantity) of EVS disks in the VDC and the EVS disk quota of the current-level VDC. When creating a lower-level VDC, the VDC administrator can set the total quota of EVS disks in the lower-level VDC and the EVS disk quota of the current-level VDC. [Figure 4-20](#) shows the quota of EVS disks in VDCs of different levels.

Figure 4-20 EVS disk quota

There are three levels of VDCs in the figure.

- Users in the VDC of each level can use EVS disk resources in the quota of the current-level VDC.
- The maximum total quota of the level-2 VDC is the total quota of the level-1 VDC minus the quota of the current-level VDC corresponding to the level-1 VDC.
- The maximum total quota of the level-3 VDC is the total quota of the level-2 VDC minus the quota of the current-level VDC corresponding to the level-2 VDC.

NOTE

You can create multiple VDCs at the same level and allocate quota to each VDC. For example, when multiple second-level VDCs are created, the maximum quota of all the second-level VDCs is the maximum quota of first-level VDCs minus the quota allocated to the first-level VDCs.

4.2.7 Relationship Between the Disk Type, Backend Storage, and Storage Array

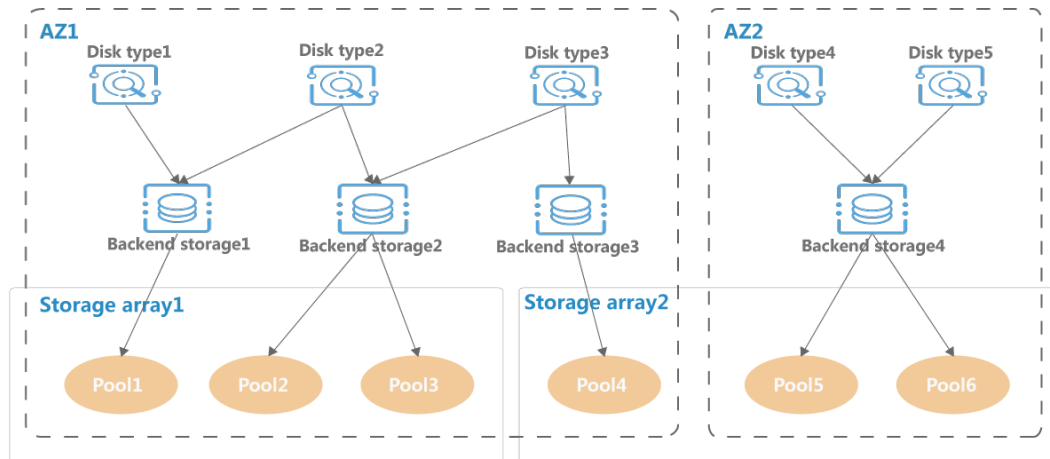
Definition

- **Disk type**
A disk type can be selected during disk creation. A disk type represents backend storage devices used by a group of disks.
For more information about disk types, see [Disk Type](#).
- **Backend storage**
Backend storage is a logical storage device that stores EVS disk resources. A backend storage device contains one or more storage pools on a storage array (Huawei Distributed Block Storage, Huawei SAN storage, or heterogeneous storage).

Relationship

Figure 4-21 shows the relationship between the disk type, backend storage, and storage array

Figure 4-21 Relationship between the disk type, backend storage, and storage array



- A storage array may belong to different availability zones. For example, storage pool 4 on storage array 2 belongs to AZ 1 while storage pools 5 and 6 belong to AZ 2.
- A backend storage device belongs to only one AZ.
- A storage pool belongs to only one backend storage device.
- A disk type belongs to only one AZ.
- A backend storage device can contain one or more storage pools, but the storage pools must be on the same storage array, for example, backend storage 2.
- A disk type can contain multiple backend storage devices from the same AZ. The storage devices can be on the same or different storage arrays but the arrays must be in the same AZ.
- Multiple disk types can be created for a backend storage device, for example, backend storage 1. Different disk types can be configured with different value-added features, such as SmartThin, SmartTier, and SmartDedupe.

4.2.8 Mapping Between Mount Points and Device Names

A block storage device is a storage device that moves data in sequences by bytes or bits (blocks). These devices support random access and wide use of cache I/O, including disks, CD-ROM, and flash drives. A block storage device can be attached to a computer or remotely accessed as it is attached to a computer. The instance supports the following block storage devices:

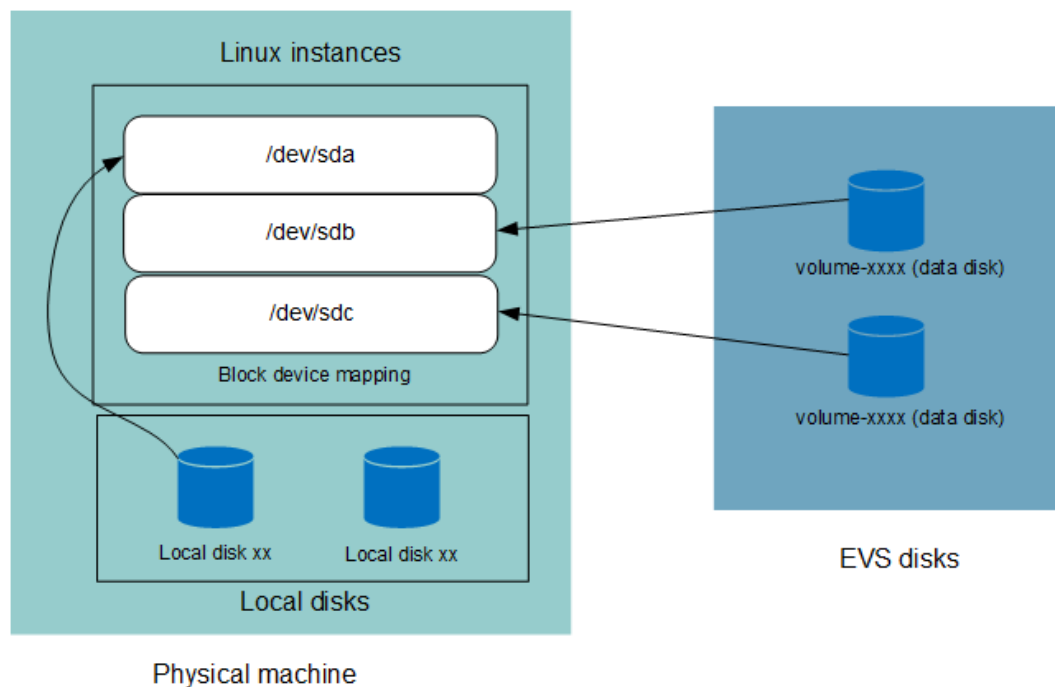
- Local disk: is the disk that is attached to the physical machine (host machine) where the instance is located and is a temporary block storage device.
- EVS disk: is a cloud disk that is attached to an instance and is a persistent block storage device.

The attachment point is the entry directory of the disk file system in Linux. It is similar to the drive letters, such as **C:**, **D:**, and **E:**, which are used to access different partitions in Windows. Each attachment point corresponds to a device name. You can attach the corresponding disk to an instance by specifying the device name of the attachment point.

Block Storage Device Mapping

The instance uses the device name (for example, **/dev/sdb**) to describe the block storage device and uses the block storage device mapping to specify the block storage device to be attached to the instance. **Figure 4-22** shows an example of mapping between EVS disks as well as local disks and instances. In the preceding figure, one local disk is attached to the Linux instance, the local disk is mapped to **/dev/sda** as the system disk, and two EVS disks are mapped to **/dev/sdb** and **/dev/sdc**, respectively, as data disks.

Figure 4-22 Example of mapping between EVS disks as well as local disks and instances



4.2.9 EVS Disk Performance

When creating a disk type, you need to select backend storage for the disk type. The I/O performance of the disk type depends on the backend storage. When applying for an EVS disk, you can select a disk type based on application performance requirements.

EVS Disk Performance Indicators

EVS disk performance indicators include Input/Output Per Second (IOPS), throughput, and latency.

- **IOPS** indicates I/O requests the system can process per unit of time, usually, per second. I/O requests typically mean data read or write operation requests. The most commonly used IOPS indicators are sequential I/O and random I/O:
 - Random I/O: Access addresses are randomly distributed in the addressing space of disks. Services that generate random I/O include OLTP, SQL, and instant messaging services.
 - Sequential I/O: Read/write operations access data continuously from adjacent addresses and are performed one by one based on logical blocks. During sequential I/O access, the disk seek time is greatly shortened because read and write heads can access the next block without moving. For example, most services such as data backup and log recording generate sequential I/O.
- **Throughput**: indicates the 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. The unit is MB/s.
- **Latency**: indicates the minimum interval between two consecutive read/write operations of an EVS disk. The unit is second.

Performance Test Methods

You can use different tools to test the EVS disk performance of instances using different OSs:

- Linux instance: You can use tools such as DD, FIO, or Sysbench to test the disk performance. For details about the test method, see [How Can I Test the Performance of a Linux EVS Disk?](#).
- Windows instance: You can use tools such as FIO and IOmeter to test the disk performance. For details about the test method, see [How Can I Test the Performance of a Windows EVS Disk?](#).

4.3 Operation Process

[Figure 4-23](#) shows the process of applying for and using an EVS disk.

Figure 4-23 Process for applying for and using an EVS disk

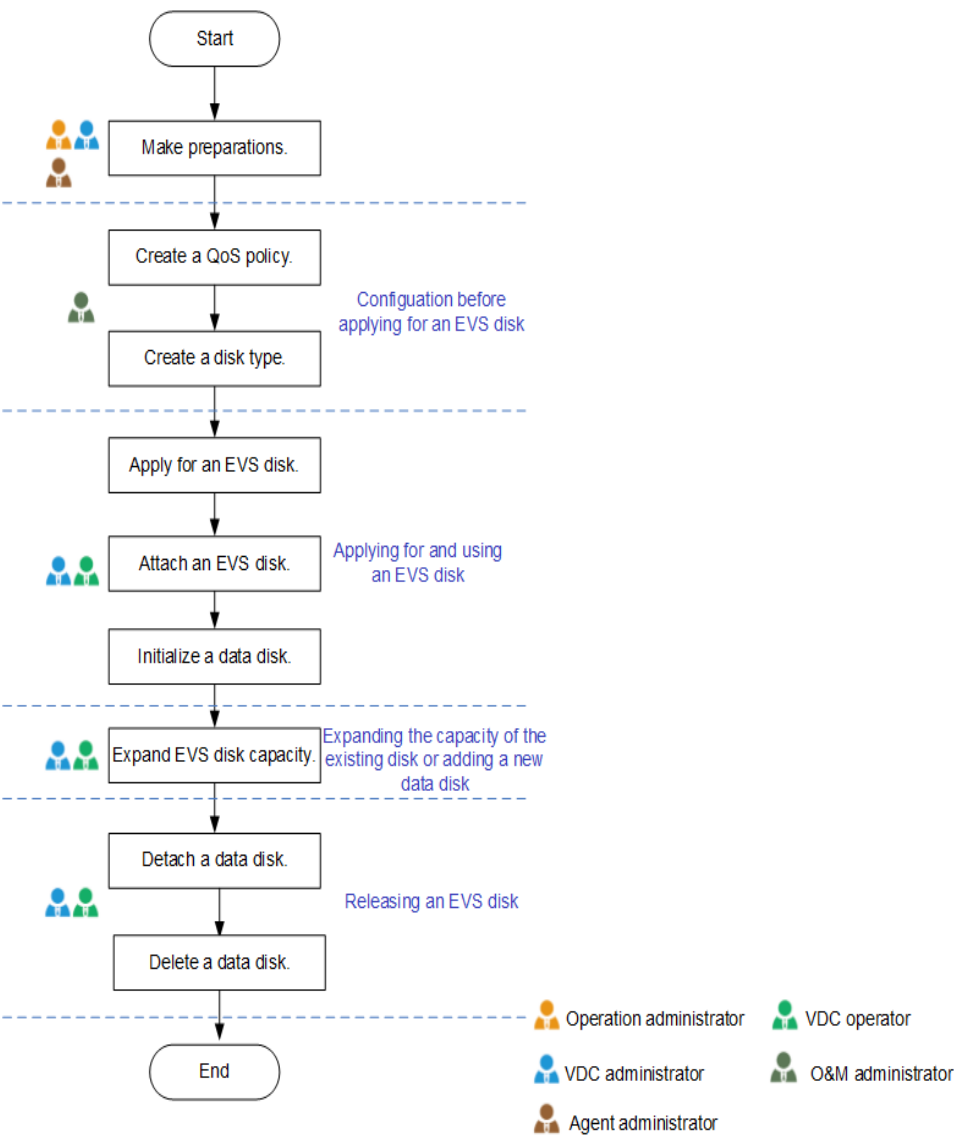


Table 4-7 lists the operators and user interfaces (UIs) required for each step.

Table 4-7 Description of the EVS disk operation process

Operation	Description and Reference	Operator and UI Portal
Preparations	<p>Before using EVS, make the following preparations:</p> <ul style="list-style-type: none">• Obtain a VDC administrator account or a VDC operator account before creating EVS disks. If no such account is available, contact an operation administrator to create a VDC and a VDC administrator account, and then use the VDC administrator account to create a VDC operator account. For details, see "VDC Tenant Modeling" in <i>HUAWEI CLOUD Stack 8.1.0 Resource Provisioning Guide</i>.• If you want to use a non-default service to apply for EVS disks, you need to apply for the EVS service on ManageOne Operation Portal (ManageOne Operation Management Portal in B2B scenarios) and release it. For details, see "Managing Services" in <i>HUAWEI CLOUD Stack 8.1.0 Resource Provisioning Guide</i>.	<ul style="list-style-type: none">• ManageOne Operation Portal (ManageOne Operation Management Portal in B2B scenarios): operation administrator• ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios): VDC administrator and agent administrator
Configuration before applying for an EVS disk	<ul style="list-style-type: none">• If you want to limit the maximum performance (IOPS or bandwidth) of burst services to prevent burst traffic from impacting other service traffic in the system, you can create a QoS policy and associate the QoS policy with a disk type. Then, the QoS policy can be applied to the disks provisioned using the disk type.• A disk type is created during installation and deployment using HUAWEI CLOUD Stack Deploy. If the disk type does not meet your requirements, contact the administrator to create a new one.	ManageOne Maintenance Portal: O&M administrator

Operation	Description and Reference	Operator and UI Portal
Applying for and using an EVS disk	You can apply for a blank EVS disk that does not contain data or apply for an EVS disk from a disk, snapshot, image, or backup and attach the EVS disk to an instance. Applying for a Data Disk Applying for an EVS Disk Using an Existing Disk Applying for an EVS Disk Using a Snapshot Applying for an EVS Disk Using a Backup	ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios): VDC operator and VDC administrator
	EVS disks are attached to instances as data disks of the instances. Attaching an EVS Disk	ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios): VDC operator and VDC administrator
	After an EVS disk is attached to a BMS as a data disk of the BMS, you need to initialize the data disk. Initializing a Data Disk	ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios): VDC operator and VDC administrator
Expanding the capacity of the existing disk or adding a new data disk	Expand the capacity of the EVS disk whose available space is insufficient. Expanding EVS Disk Capacity	ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios): VDC operator and VDC administrator
Releasing an EVS disk	If the instance to which an EVS disk is attached does not need the EVS disk anymore, you can detach the EVS disk. Detaching a Data Disk	ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios): VDC operator and VDC administrator
	After an EVS disk is detached, delete the EVS disk to release storage space. Deleting Data Disks	ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios): VDC operator and VDC administrator

4.4 Applying for an EVS Disk

4.4.1 Applying for a Data Disk

A data disk can be created together with an instance. During the instance creation, you can add data disks and set the disk type, disk capacity, and disk sharing attribute. After the instance is created, the disk will be automatically attached to the instance. This section describes how to create data disks and attach them to instances after the instance creation.

Context

ECSs do not support the merging of EVS disks. Each EVS disk is independent, and the spaces of multiple EVS disks cannot be merged through formatting. You are advised to plan the number and capacity of EVS disks before disk creation. It is not recommended that logical volumes managed by the LVM be created on the disks, because snapshots are created for independent EVS disks and creating such logical volumes will generate differential data after a snapshot rollback.

Procedure

Step 1 Log in to ManageOne as a VDC administrator or VDC operator using a browser.

URL in non-B2B scenarios: <https://Domain name of ManageOne Operation Portal>, for example, <https://console.demo.com>.

URL in B2B scenarios: <https://Domain name of ManageOne Tenant Portal>, for example, <https://tenant.demo.com>.

You can log in using a password or a USB key.

- Login using a password: Enter the username and password.
The password is that of the VDC administrator or VDC operator.
- Login using a USB key: Insert a USB key with preset user certificates, select the required device and certificate, and enter a PIN.

Step 2 Select a region and resource set, and choose **Storage > Elastic Volume Service**. The EVS console is displayed.

Step 3 Click **Apply for EVS Disk**.

Step 4 In the **Select Service** dialog box, select the target service and then click **Apply Now**.

NOTE

Services are created by administrators based on operation requirements. When creating a service, the administrator can lock service parameters (for example, apply for a fixed disk type or apply for resources in a specified AZ), specify the service publishing scope (for example, visible to only specific VDCs), and set the product approval process. For details about how to create more services, see "Creating a Common Service" in *HUAWEI CLOUD Stack 8.1.0 Resource Provisioning Guide*.

Step 5 In the **Apply for EVS Disk** dialog box, set the parameters as prompted by referring to [Table 4-8](#).

 **NOTE**

If the service you selected has parameters **AZ**, **Capacity (GB)**, **Disk Sharing**, or **Disk Type** configured, the parameter values configured for the product will be displayed for the EVS disk you apply for.

Table 4-8 Parameters required when applying for an EVS disk

Parameter	Description	Example
AZ	Specifies the availability zone (AZ) where an EVS disk is to be created. NOTE <ul style="list-style-type: none">EVS disks can be attached to instances only in the same AZ.If an AZ is bound to a tenant upon the tenant creation, the resource pools in this AZ are used as dedicated storage pools for this tenant and invisible to other tenants not bound to the AZ. Services of this tenant can run on dedicated physical devices without interference from other tenants.	az1.dc1
Data Source	Select Do not specify to apply for a blank EVS disk that does not contain data.	Do not specify
Disk	Only data disks can be applied for and attached to BMS instances.	Data disk
Disk Type	Select the disk type. You can select the disk type created in "Storage Services" > "Elastic Volume Service (EVS for BMS)" > "Configuration Before Applying for an EVS Disk" > "(Optional) Creating a Disk Type" in <i>HUAWEI CLOUD Stack 8.1.0 Resource Provisioning Guide</i> . You can also select the disk type created during automatic installation and deployment using HUAWEI CLOUD Stack Deploy. The name of the automatically created disk type is the value of business_volume_type in deployment parameter summary file <i>xxx_export_all_v2_EN.xlsx</i> .	business
Capacity (GB)	EVS disk capacity. The EVS disk capacity can neither exceed the total capacity quota of EVS disks nor the capacity quota of the current disk type.	10 GB
Device Type	Select SCSI . Only EVS disks of the SCSI type can be attached to the BMS.	SCSI

Parameter	Description	Example
Disk Sharing	<ul style="list-style-type: none">● Disable indicates the new disk will be a non-shared EVS disk.● Enable indicates the new disk will be a shared EVS disk. Such a disk can be attached to multiple ECSs. NOTE If the backend storage device is a heterogeneous storage device, no shared EVS disk can be created.	Enable
Disk Name	<p>The disk name can contain only letters, digits, underscores (_), and hyphens (-). When applying for a single EVS disk, ensure that the disk name contains less than or equal to 63 characters. When applying for EVS disks in batches, ensure that the disk name contains less than or equal to 58 characters.</p> <ul style="list-style-type: none">● If you apply for a single EVS disk, the value of this parameter is used as the name of the EVS disk.● If you apply for multiple EVS disks in batches, the value of this parameter is used as the prefix of the names of the EVS disks. The name of each EVS disk resembles <i>Disk Name-four-digit number</i>. NOTE For example, if you apply for two EVS disks and set <i>Disk Name</i> to volume , the names of the two EVS disks will be volume-0001 and volume-0002 .	volume-0001
Quantity	<p>Specifies the number of EVS disks that you apply for. The default value is 1, which means that you apply for one EVS disk.</p> <p>By default, a maximum of 100 EVS disks can be created at a time. The number of EVS disks that can be batch created varies with the current EVS disk quantity quota.</p>	1
Required Duration	<p>Specifies the validity period of the EVS disk that you apply for.</p> <ul style="list-style-type: none">● If you select Unlimited, the new EVS disk has no expiration date.● If you select 1 year, the validity period of the new EVS disk will be one year, which is subject to the expiration date displayed on the page.● If you select Custom, specify an expiration date for the EVS disk that you apply for.	Unlimited

Parameter	Description	Example
Description	Describe the EVS disk that you apply for. The length cannot exceed 63 characters.	-

Step 6 Click **Next**.

Step 7 Confirm that the information is correct and click **Add to Cart** or **Apply Now**.

If the configuration is incorrect, click **Back**.

- If you click **Add to Cart**, go to [Step 8](#).
- If you click **Apply Now**, go to [Step 9](#).

Step 8 Submit an application for the service in the shopping cart.

1. Click the shopping cart in the upper right corner of the page.
2. Select the service and click **Apply Now**.
3. Enter the order information and click **OK** to submit the application.

Step 9 If the application for an EVS disk requires approval, contact the administrator for approval. Otherwise, skip this step.

 **NOTE**

The application for an EVS disk can be rejected by the administrator. If you have entered incorrect configuration when applying for an EVS disk, you can contact the administrator to reject the application, correct the configuration, and submit the application again.

Step 10 On the **Elastic Volume Service** page, view the status of the EVS disk. After the status of the EVS disk changes to **Available**, the EVS disk is successfully created.

----End

Follow-up Procedure

If you want to use the new EVS disk, attach it to a BMS. For details, see [Attaching an EVS Disk](#).

4.4.2 Applying for an EVS Disk Using an Existing Disk

This section describes how to quickly apply for a data disk that contains the same data as an existing data disk. The new data disk can be used by a BMS instance after being attached to the instance.

Restrictions

- When you apply for an EVS disk using an existing disk, the status of the source disk must be **Available** or **In-use**.
- When OceanStor V3/V5 is used, an EVS disk can be applied for from an existing EVS disk only after the administrator imports the HyperCopy license onto the storage device.

Procedure

Step 1 [Log in to EVS Console as a VDC administrator or VDC operator.](#)

Step 2 In the EVS disk list, locate the disk used for applying for a new EVS disk.

 **NOTE**

You can also click **Apply for EVS Disk** and choose **Create from disk** to apply for an EVS disk. The restrictions on the source disk are as follows:

- Only EVS disks in the current AZ can be selected as source disks.
- If the disk capacity and type have been specified in the selected service, the capacity of the source disk must be less than or equal to that of the specified one, and the disk type must be the same as the specified one. Disks not meeting this requirement cannot be used as source disks.

Step 3 In the row where the source disk resides, choose **More > Apply for Disk**.

Step 4 In the **Select Service** dialog box, select the target service and then click **Apply Now**.

 **NOTE**

- If the disk capacity has been specified in the service, you can only select the service whose disk capacity is greater than or equal to the source disk capacity.
- If the disk type has been specified in the service, you can only select the service whose disk type is the same as the source disk type.
- If the disk capacity and disk type have been specified in the service, you can only select the service whose disk capacity is greater than or equal to the source disk capacity and whose disk type is the same as the source disk type.

Step 5 In the **Apply for EVS Disk** dialog box, configure EVS disk information as prompted.

Some parameters must be consistent with those of the source disk and cannot be modified on the GUI. Configure other parameters by referring to [Table 4-9](#). Parameters not listed in [Table 4-9](#) cannot be modified.

Table 4-9 Parameters for applying for an EVS disk using a disk

Parameter	Description	Example Value
Capacity (GB)	<p>EVS disk capacity.</p> <p>If you do not specify the capacity when applying for an EVS disk using a disk, the capacity of the new disk is the same as that of the source disk.</p> <p>The EVS disk capacity can neither exceed the total capacity quota of EVS disks nor the capacity quota of the current disk type.</p>	10 GB

Parameter	Description	Example Value
Disk Sharing	<ul style="list-style-type: none">● Disable indicates that the new disk will be a non-shared EVS disk.● Enable indicates that the new disk will be a shared EVS disk. Such a disk can be attached to multiple ECSs. NOTE <ul style="list-style-type: none">– If the backend storage device is a heterogeneous storage device, no shared EVS disk can be created.– If the source disk is a system disk, a shared disk cannot be created.	Enable
Disk Name	The disk name can contain only letters, digits, underscores (_), and hyphens (-). When applying for a single disk, ensure that the disk name contains not more than 63 characters.	volume-0001
Quantity	Number of applied EVS disks. You cannot apply for EVS disks in batches using a disk. You can apply for only one EVS disk at a time.	1
Required Duration	Specifies the validity period of the EVS disk that you apply for. <ul style="list-style-type: none">● If you select Unlimited, the new EVS disk has no expiration date.● If you select 1 year, the validity period of the new EVS disk will be one year, which is subject to the expiration date displayed on the console.● If you select Custom, specify an expiration date for the EVS disk that you apply for.	Unlimited
Description	Describes the EVS disk that you apply for. The length cannot exceed 63 characters.	-

Step 6 Click **Next**.

Step 7 Confirm that the information is correct and click **Add to Cart** or **Apply Now**.

If the configuration is incorrect, click **Back**.

- If you click **Add to Cart**, go to [Step 8](#).
- If you click **Apply Now**, go to [Step 9](#).

Step 8 Submit an application for the service in the shopping cart.

1. Click the shopping cart in the upper right corner of the page.

2. Select the service and click **Apply Now**.
3. Enter the order information and click **OK** to submit the application.

Step 9 If the application for an EVS disk requires approval, contact the administrator for approval. Otherwise, skip this step.

 **NOTE**

The application for an EVS disk can be rejected by the administrator. If you have entered incorrect configuration when applying for an EVS disk, you can contact the administrator to reject the application, correct the configuration, and submit the application again.

Step 10 On the **Elastic Volume Service** page, view the status of the EVS disk. After the EVS disk is created and its status changes to **Available**, the EVS disk is successfully created.

 **NOTE**

The time required for applying for an EVS disk using an existing disk depends on the capacity of the existing disk. A larger capacity indicates a longer time. The default speed is 10 MB/s to 20 MB/s.

----End

Follow-up Procedure

If you want to use the new EVS disk, attach it to a BMS. For details, see [Attaching an EVS Disk](#).

4.4.3 Applying for an EVS Disk Using a Snapshot

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, such as data mining, report query, and development and test. This method protects the initial data and creates disks rapidly, meeting the diversified service data requirements.

Context

You can create an EVS disk from a snapshot. The created EVS disk contains the data of the snapshot, and is a precise copy of the source EVS disk. An EVS disk created from a snapshot does not need to be partitioned or formatted, and no file system needs to be created. When the EVS disk is attached to an instance, the EVS disk can read and write data. You can but are advised not to use the snapshot of a small-capacity data disk (< 2 TB) to create large-capacity data disks (> 2 TB).

Instead, you can create large-capacity blank data disks or use the snapshot of a large-capacity data disk to create other large-capacity data disks. The reasons are as follows:

- When a large-capacity data disk is created using a snapshot of a small-capacity data disk, the disk capacity is only expanded at the block device level, but the partition format and file system are not automatically converted.
- If the snapshot of a small-capacity data disk uses the MBR partition format, the partitioning tool (**parted** on the Linux OS or the disk management module on the Windows OS) cannot convert the partition format from MBR

to GPT while retaining data. Even if you use a snapshot of a small-capacity data disk to create a large-capacity data disk, you need to delete the original data and then partition the disk in the GPT format during partition formatting.

- If a large-capacity data disk has been created using the snapshot of a small-capacity data disk, the partition format of the created disk is MBR. Before formatting, you need to delete the original disk, convert it into a GPT partition, and then partition and format the data again. That is, the data in the original snapshot will be deleted.

Procedure

Step 1 Log in to the EVS Console as a VDC administrator or VDC operator. For details, see [Logging In to the EVS Console as a VDC Administrator or VDC Operator](#).

Step 2 In the navigation pane on the left, choose **Elastic Volume Service > Snapshot**.

NOTE

You can also click **Apply for EVS Disk** and choose **Create from snapshot** to apply for an EVS disk. In this case, only snapshots in the current AZ can be used to apply for EVS disks.

Step 3 In the snapshot list, locate the row that contains the snapshot, and click **Apply for EVS Disk** in the **Operation** column.

Step 4 In the **Select Service** dialog box, select the target service and then click **Apply Now**.

NOTE

- If the disk capacity has been specified in the service, you can only select the service whose disk capacity is greater than or equal to the snapshot capacity.
- If the disk type has been specified in the service, you can only select the service whose disk type is the same as the source disk type of the snapshot.
- If the disk capacity and disk type have been specified in the service, you can only select the service whose disk capacity is greater than or equal to the snapshot capacity and whose disk type is the same as the source disk type of the snapshot.

Step 5 Click **Next**.

Step 6 Confirm that the information is correct and click **Add to Cart** or **Apply Now**.

If the configuration is incorrect, click **Back**.

- If you click **Add to Cart**, go to [Step 7](#).
- If you click **Apply Now**, go to [Step 8](#).

Step 7 Submit an application for the service in the shopping cart.

1. Click the shopping cart in the upper right corner of the page.
2. Select the service and click **Apply Now**.
3. Enter the order information and click **OK** to submit the application.

Step 8 On the **Elastic Volume Service** page, view the status of the EVS disk. After the EVS disk is created and its status changes to **Available**, the EVS disk is successfully created.

 NOTE

The time required for applying for an EVS disk using an existing disk depends on the capacity of the existing disk. A larger capacity indicates a longer time. The default speed is 10 MB/s to 20 MB/s.

----End

Follow-up Procedure

After the EVS disk is created, you can [attach it to a BMS](#).

4.4.4 Applying for an EVS Disk Using a Backup

You can use a VBS backup to create an EVS disk. After the EVS disk is created, the data on the new disk is the same as that in the VBS backup.

You can use a backup to create an EVS disk in either of the following ways:

- EVS console: To use the specified backup to create an EVS disk, perform the related operations on the EVS console.
- VBS console: To view the detailed information about a backup and then use it to create an EVS disk, perform the related operations on the VBS console.

EVS Console


Step 1 Log in to ManageOne as a VDC administrator or VDC operator using a browser.

URL in non-B2B scenarios: <https://Domain name of ManageOne Operation Portal>, for example, <https://console.demo.com>.

URL in B2B scenarios: <https://Domain name of ManageOne Tenant Portal>, for example, <https://tenant.demo.com>.

You can log in using a password or a USB key.

- Login using a password: Enter the username and password.
The password is that of the VDC administrator or VDC operator.
- Login using a USB key: Insert a USB key with preset user certificates, select the required device and certificate, and enter a PIN.

Step 2 Click  in the upper left corner, select a region and resource set, and choose **Storage > Elastic Volume Service**. The EVS console is displayed.

Step 3 Click **Apply for EVS Disk**.

Step 4 In the **Select Service** dialog box, select the target service and then click **Apply Now**.

 NOTE

Services are created by administrators based on operation requirements. When creating a service, the administrator can lock service parameters (for example, apply for a fixed disk type or apply for resources in a specified AZ), specify the service publishing scope (for example, visible to only specific VDCs), and set the product approval process. For details about how to create more services, see "Creating a Common Service" in *HUAWEI CLOUD Stack 8.1.0 Resource Provisioning Guide*.

Step 5 In the **Apply for EVS Disk** dialog box, set the parameters as prompted by referring to [Table 4-10](#).

Table 4-10 Parameters for applying for an EVS disk using a backup

Parameter	Description	Example
AZ	<p>Specifies the availability zone (AZ) where an EVS disk is to be created.</p> <p>NOTE</p> <ul style="list-style-type: none">EVS disks can be attached to instances only in the same AZ.If an AZ is bound to a tenant upon the tenant creation, the resource pools in this AZ are used as dedicated storage pools for this tenant and invisible to other tenants not bound to the AZ. Services of this tenant can run on dedicated physical devices without interference from other tenants.	az1.dc1
Data Source	<p>Select Create from backup, click Select, and select a backup as the data source of the new EVS disk.</p> <p>NOTE</p> <ul style="list-style-type: none">If the VBS service is not configured or the VBS service is configured but disabled, this parameter is not displayed. To use the VBS, configure and enable it by referring to "Storage Services" > "Elastic Volume Service (EVS for BMS)" > "FAQs" > "How Do I Enable the VBS?" in <i>HUAWEI CLOUD Stack 8.1.0 Resource Provisioning Guide</i>.If the backend storage is heterogeneous storage, you cannot select Create from backup.	Create from backup
Disk	Only data disks can be applied for and attached to BMS instances.	Data disk

Parameter	Description	Example
Disk Type	Select the disk type. You can select the disk type created in "Storage Services" > "Elastic Volume Service (EVS for BMS)" > "Configuration Before Applying for an EVS Disk" > "(Optional) Creating a Disk Type" in <i>HUAWEI CLOUD Stack 8.1.0 Resource Provisioning Guide</i> . You can also select the disk type created during automatic installation and deployment using HUAWEI CLOUD Stack Deploy. The name of the automatically created disk type is the value of business_volume_type in deployment parameter summary file <i>xxx_export_all_v2_EN.xlsx</i> .	-
Capacity (GB)	EVS disk capacity. When you apply for an EVS disk using a backup, the capacity of the new EVS disk must be greater than or equal to that of the source disk of the backup. The EVS disk capacity can neither exceed the total capacity quota of EVS disks nor the capacity quota of the current disk type.	10 GB
Device Type	Select SCSI . Only EVS disks of the SCSI type can be attached to the BMS.	SCSI
Disk Sharing	<ul style="list-style-type: none">• Disable indicates the new disk will be a non-shared EVS disk.• Enable indicates the new disk will be a shared EVS disk. Such a disk can be attached to multiple ECSs.	Disable
Disk Name	The disk name can contain only letters, digits, underscores (_), and hyphens (-). When applying for a single disk, ensure that the disk name contains not more than 63 characters.	-
Quantity	Specifies the number of applied EVS disks. You cannot create EVS disks in batches using a backup. You can apply for only one EVS disk at a time.	1

Parameter	Description	Example
Required Duration	Specifies the validity period of the EVS disk that you apply for. <ul style="list-style-type: none">If you select Unlimited, the new EVS disk has no expiration date.If you select 1 year, the validity period of the new EVS disk will be one year, which is subject to the expiration date displayed on the page.If you select Custom, specify an expiration date for the EVS disk that you apply for.	Unlimited
Description	Describe the EVS disk that you apply for. The length cannot exceed 63 characters.	-

Step 6 Click **Next**.

Step 7 Confirm that the information is correct and click **Add to Cart** or **Apply Now**.

If the configuration is incorrect, click **Back**.

- If you click **Add to Cart**, go to [Step 8](#).
- If you click **Apply Now**, go to [Step 9](#).

Step 8 Submit an application for the service in the shopping cart.

- Click the shopping cart in the upper right corner of the page.
- Select the service and click **Apply Now**.
- Enter the order information and click **OK** to submit the application.

Step 9 If the application for an EVS disk requires approval, contact the administrator for approval. Otherwise, skip this step.

 **NOTE**

The application for an EVS disk can be rejected by the administrator. If you have entered incorrect configuration when applying for an EVS disk, you can contact the administrator to reject the application, correct the configuration, and submit the application again.

Step 10 On the **Elastic Volume Service** page, view the status of the EVS disk.

During EVS disk creation, the four intermediate states **Creating**, **Available**, **Restoring**, and **Available** will be displayed in sequence. After the state has changed from **Creating** to **Available**, the disk has been successfully created. After the state has changed from **Restoring** to **Available**, backup data has been successfully restored to the created disk.

 **NOTE**

When a new EVS disk created using a VBS backup is attached to an instance, the architecture of the instance must be the same as that of the instance to which the source disk of the backup is attached. For example, they use the ARM architecture.

----End

VBS Console

For details about how to create an EVS disk using a backup, see **Operation Help Center > DR & Backup > User Guide > Restoring a Disk Using a Backup > Creating a Disk Using Backup Data.**

4.5 Attaching an EVS Disk

A created EVS disk can be used by an instance only after being attached to the instance. The EVS disk created together with an instance is automatically attached to the instance.

Restrictions

- Only SCSI EVS disks can be attached to BMSs.
- Regardless if a shared EVS disk or non-shared EVS disk is attached to an instance, the EVS disk and the instance must be in the same AZ.
- An EVS disk cannot be attached to an instance that has expired.
- An EVS disk cannot be attached to an instance that has been soft deleted.
- An EVS disk cannot be attached to an instance that has been stopped.
- If the ECS to which an EVS disk belongs has not been created, the EVS disk cannot be attached to another ECS.

Context

- You can attach the disks in the following modules:
 - EVS console: To attach multiple EVS disks to different BMSs, perform related operations on the EVS console.
 - BMS console: To attach multiple disks to one BMS, perform related operations on the BMS console.

Operations on the BMS Console

Step 1 Log in to ManageOne as a VDC administrator or VDC operator using a browser.


URL in non-B2B scenarios: **<https://Domain name of ManageOne Operation Portal>**, for example, **<https://console.demo.com>**.

URL in B2B scenarios: **<https://Domain name of ManageOne Tenant Portal>**, for example, **<https://tenant.demo.com>**.

You can log in using a password or a USB key.

- Login using a password: Enter the username and password.
The password is that of the VDC administrator or VDC operator.

- Login using a USB key: Insert a USB key with preset user certificates, select the required device and certificate, and enter a PIN.

Step 2 Click  in the upper left corner and select a region and resource set. Choose **Compute > Bare Metal Server**.

Step 3 In the BMS list, locate the BMS to which the disk will attach.

Step 4 Click the name of the queried BMS.

The page showing details about the BMS is displayed.

Step 5 Click the **EVS** tab. Then, click **Attach Disk**.

Step 6 The **Attach Disk** page is displayed. Select the EVS disk that you want to attach.

If you want to apply for a new EVS disk, click **Apply for an EVS Disk**. For details, see section [Applying for an EVS Disk](#).

Click **OK**.

NOTICE

For details about how to attach an EVS disk, see **Operation Help Center > Compute > Bare Metal Server > User Guide > Data Disks > Attaching or Detaching EVS Disks**.

----End

Operations on the EVS Console

Step 1 Select a region and resource set, and choose **Storage > Elastic Volume Service**. The EVS console is displayed.

Step 2 On the **Elastic Volume Service** page, locate the row that contains the target EVS disk, and click **Attach** in the **Operation** column.

The disk can be either a shared disk or a non-shared disk.

Step 3 In the **Attach Disk** dialog box, select the BMS to which the EVS disk is to be attached.

When the EVS disk is attached to a BMS, you cannot select a mount point.

Step 4 Click **OK**.

On the EVS disk list page, if the disk status is **In-use**, the EVS disk has been successfully attached to an instance.

NOTICE

For details about how to attach an EVS disk, see **Operation Help Center > Compute > Bare Metal Server > User Guide > Data Disks > Attaching or Detaching EVS Disks**.

----End

Follow-up Procedure

After the EVS disk is attached to the instance, perform subsequent operations by referring to [Table 4-11](#).

Table 4-11 Operations after the EVS disk is attached

EVS Disk	Instance OS	Follow-up Procedure
New blank data disk	Windows	To initialize the data disk of a Windows ECS, see Initializing a Windows Data Disk .
	Linux	<ul style="list-style-type: none">• If the disk capacity is less than 2 TB, see Initializing a Linux Data Disk (fdisk).• If the disk capacity is greater than or equal to 2 TB, see Initializing a Linux Data Disk (parted).
Data disk created from a snapshot, an existing disk, or a backup	Windows	<p>Check whether the source disk corresponding to snapshots, existing disks, and backups has been initialized.</p> <ul style="list-style-type: none">• If the conditions are met, you can directly use the EVS disk without performing any other operation.• If the conditions are not met, initialize the disk before using it. The operations are the same as those for a new blank data disk.
	Linux	<p>Check whether the source disk corresponding to snapshots, existing disks, and backups has been initialized.</p> <ul style="list-style-type: none">• If yes, log in to the instance and run the mount command (mount <i>partition mount path</i>) to mount the partition. Set automatic mounting upon system startup. Then, the EVS disk can be used properly. For details about how to set automatic mounting upon system startup, see Setting Automatic Disk Attachment Upon Instance Start.• If the conditions are not met, initialize the disk before using it. The operations are the same as those for a new blank data disk.

4.6 Initializing a Data Disk

4.6.1 Initialization Overview

After attaching an EVS disk to an instance, you need to log in to the instance to partition and initialize the disk so that the data disk can be used for the instance. [Table 4-12](#) provides common partition styles.

Table 4-12 Disk partition style

Disk Partition Style	Maximum Disk Capacity Supported	Maximum Number of Partitions Supported	Partitioning Tool Supported
Main Boot Record (MBR)	2 TB	<ul style="list-style-type: none">Four primary partitionsThree primary partitions and one extended partition	<ul style="list-style-type: none">• For Linux OS fdisk or parted• For Windows OS Disk management
GUID Partition Table (GPT)	18 EB NOTE 1 EB = 1,048,576 TB	Unlimited	<ul style="list-style-type: none">• For Linux OS parted• For Windows OS Disk management

4.6.2 Initializing a Windows Data Disk

A data disk attached to an instance or created together with an instance can be used by the instance only after being initialized. This section uses the Windows Server 2008 R2 Enterprise OS as an example. Specific initialization operations vary with OSs.

Prerequisites

- You have logged in to an ECS. For details about how to log in to the bare metal server, see **Operation Help Center > Compute > Bare Metal Server > User Guide > Logging In to a BMS**.
- A data disk has been attached to a BMS and has not been initialized.

Context

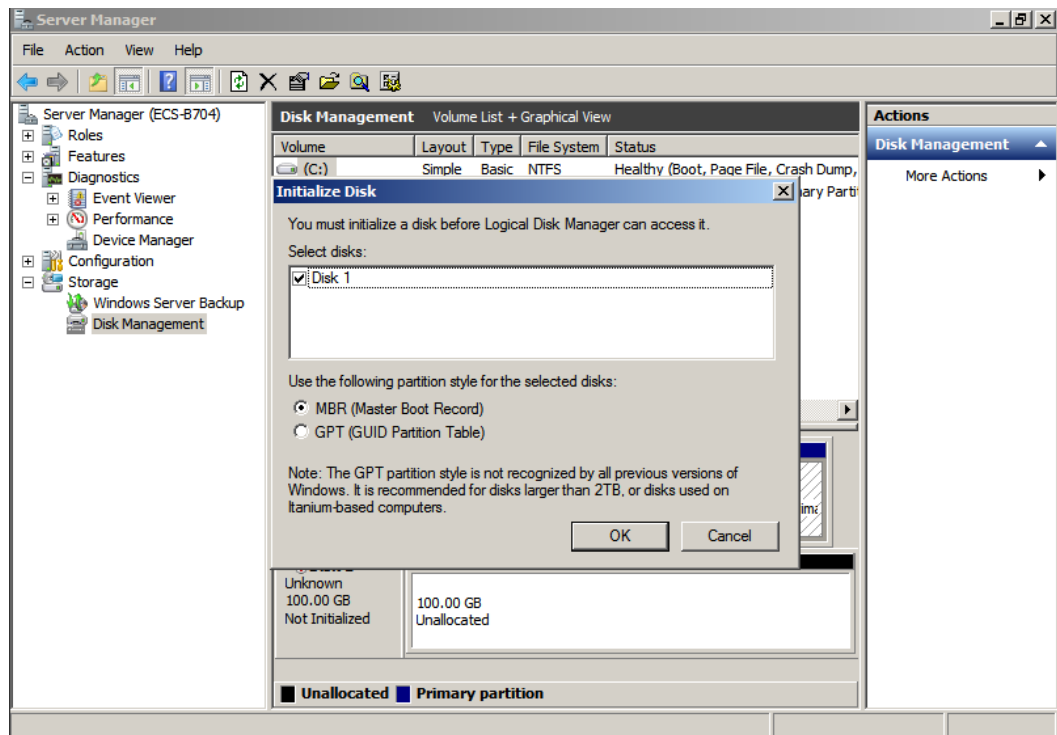
Initializing a data disk is highly risky. If there is available data on the data disk, perform the operations described in [Applying for a Snapshot](#) or [Creating a Backup](#) for the data disk first.

Procedure

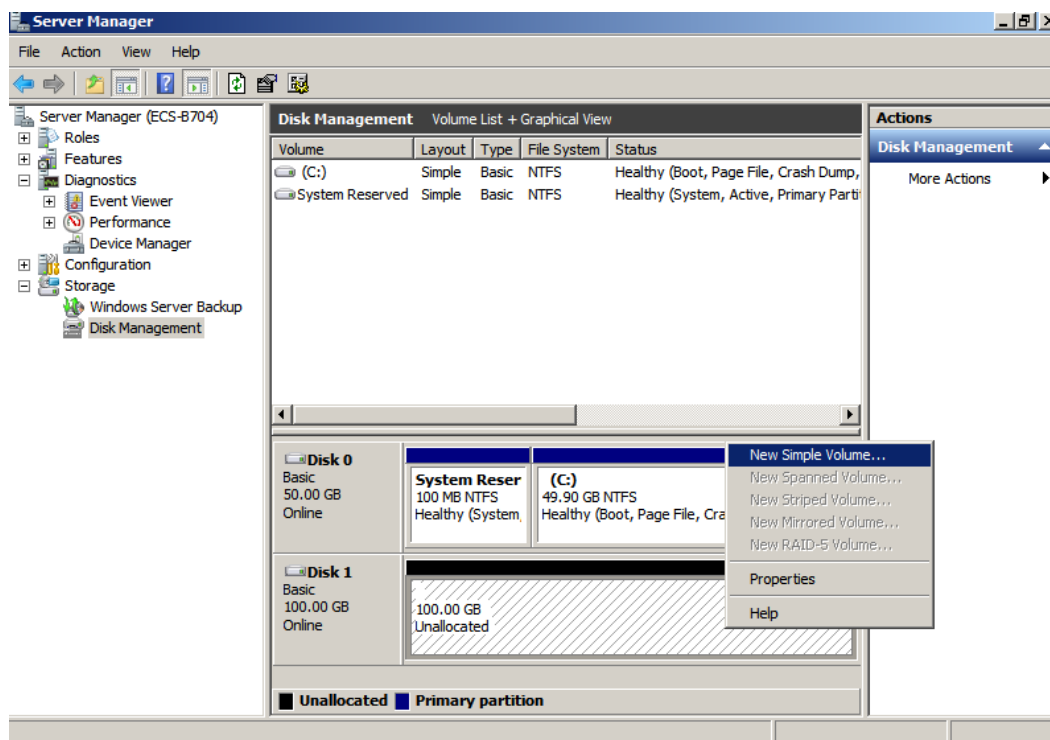
- Step 1** In desktop, right-click **Computer** and choose **Manage** from the shortcut menu.
The **Server Manager** page is displayed.
- Step 2** In the navigation pane, choose **Storage > Disk Management**.
- Step 3** If the disk to be initialized in the disk list is in **Offline** state, right-click in the disk area and choose **Online** from the shortcut menu.
Then, the disk status changes from **Offline** to **Uninitialized**.
- Step 4** Right-click in the disk area and choose **Initialize Disk** from the shortcut menu. In the displayed **Initialize Disk** dialog box, select **MBR (Master Boot Record)** and click **OK**.

NOTE

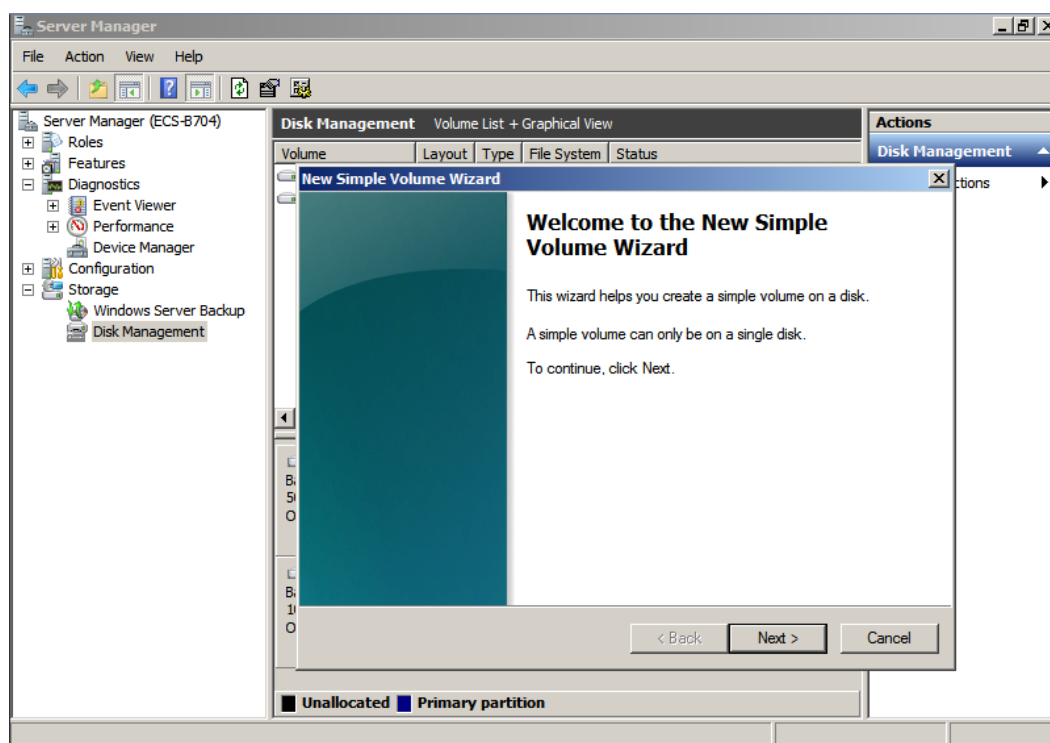
If the data disk to be initialized is larger than 2 TB, select **GPT (GUID Partition Table)** in the dialog box.



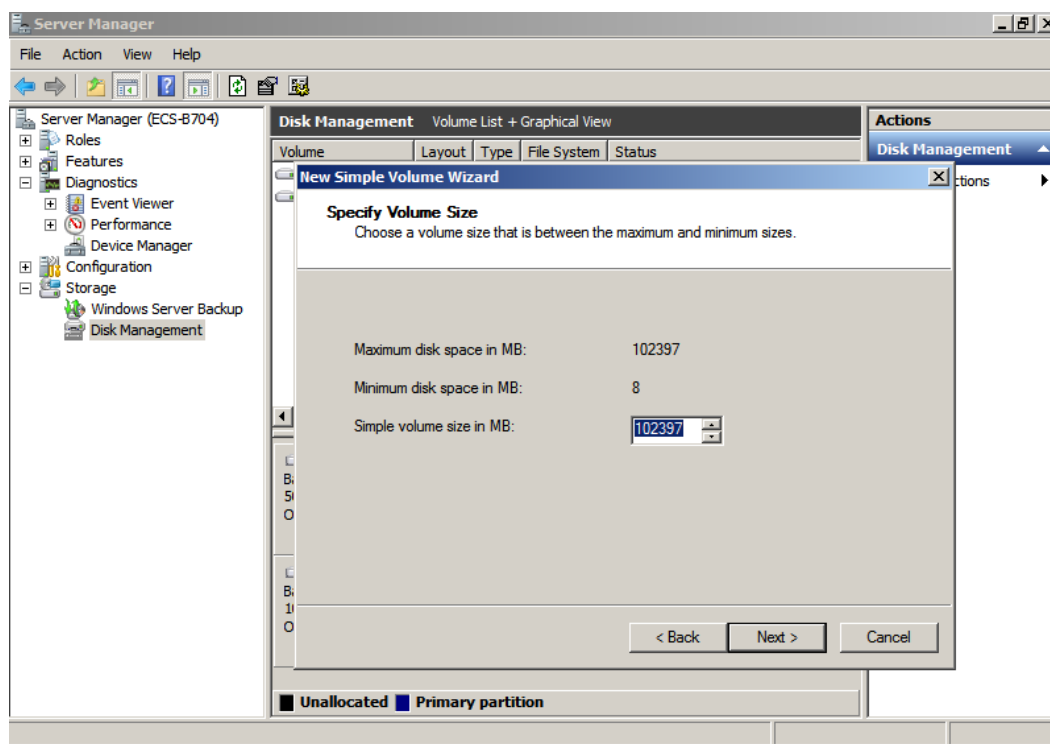
- Step 5** Right-click at the unallocated disk space and choose **New Simple Volume** from the shortcut menu.



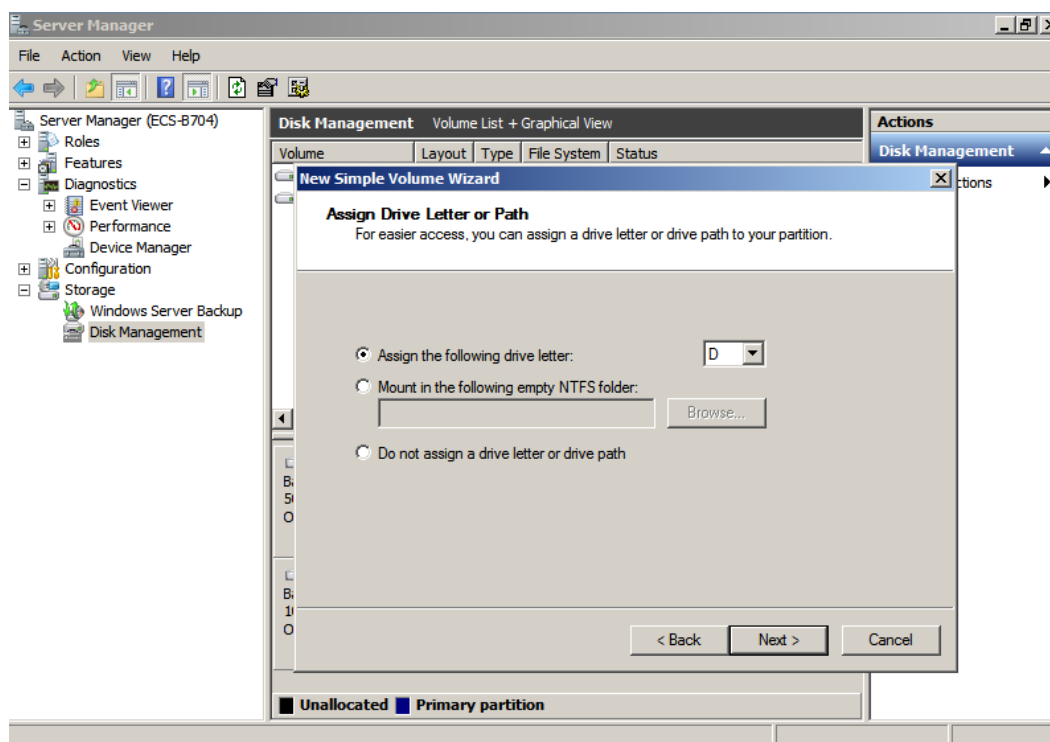
Step 6 On the displayed **New Simple Volume Wizard** page, click **Next**.



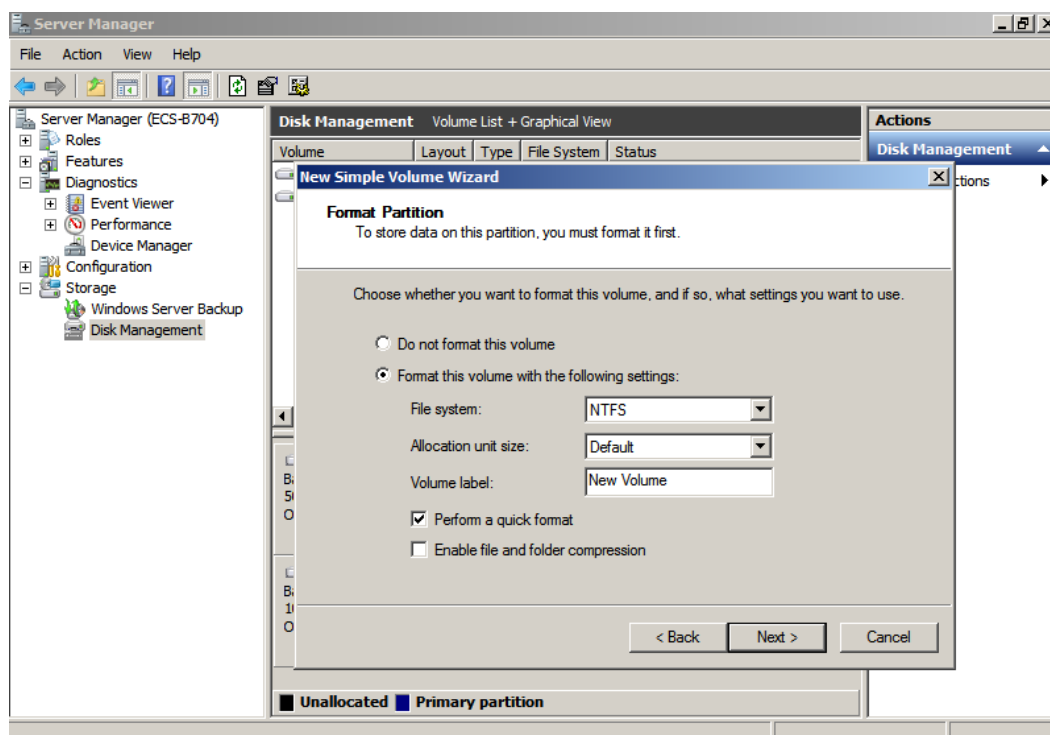
Step 7 Specify the simple volume size as required (the default value is the maximum) and click **Next**.



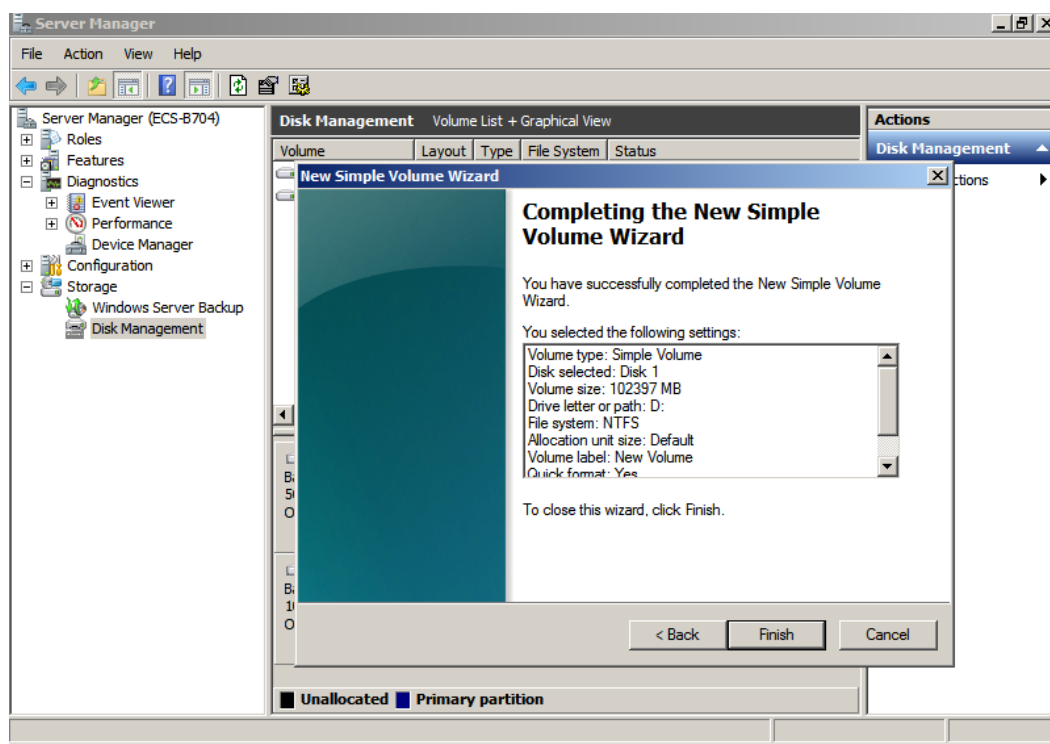
Step 8 Assign the driver letter and click **Next**.



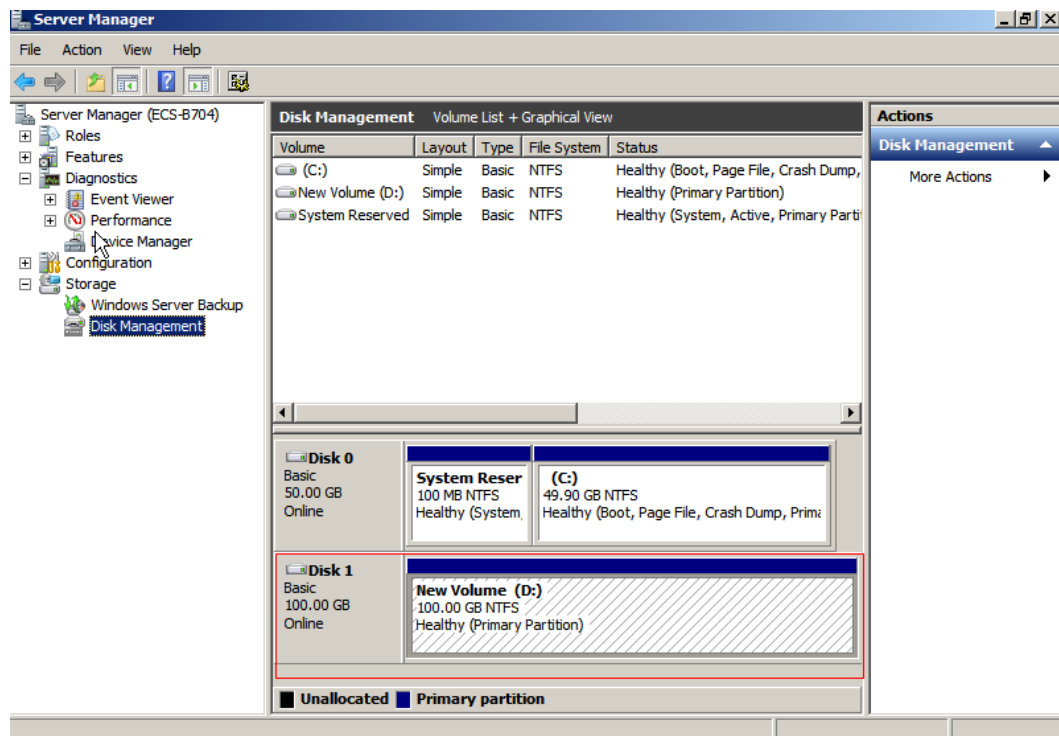
Step 9 Select **Format this volume with the following settings**, set parameters based on the actual requirements, and select **Perform a quick format**. Then click **Next**.



Step 10 Click **Finish**.



Wait for the initialization to complete. When the volume status changes to **Healthy**, the initialization has finished.



----End

4.6.3 Initializing a Linux Data Disk (fdisk)

A data disk attached to an instance or created together with an instance can be used by the instance only after being initialized. This section uses an instance running CentOS 7.0 64 bit as an example, and uses the fdisk partition tool to set up partitions for the data disk. Initialization operations vary with OSs.

Prerequisites

- You have logged in to an ECS. For details about how to log in to the bare metal server, see **Operation Help Center > Compute > Bare Metal Server > User Guide > Logging In to a BMS**.
- A data disk has been attached to a BMS and has not been initialized.

Context

Both the fdisk and parted can be used to partition a Linux data disk. For a disk larger than 2 TB, only parted can be used because fdisk cannot partition such a large disk. For details, see [Initializing a Linux Data Disk \(parted\)](#).

Creating Partitions and Mounting a Disk

The following example shows how to create a new primary partition on a new data disk that has been attached to an instance. The primary partition will be created using fdisk, and MBR is the default partition style. Furthermore, the partition will be formatted using the ext4 file system, mounted on the `/mnt/sdc` directory, and set to be automatically mounted upon a system start.

Step 1 Run the following command to view information about the added data disk:

fdisk -l

Information similar to the following is displayed: (In the command output, the server contains two disks. **/dev/xvda** is the system disk, and **/dev/xvdb** is the added data disk.)

 **NOTE**

If you do not log in to the ECS and run the **umount** command but directly detach the **/dev/xvdb** or **/dev/vdb** EVS disk on the management console, the disk name in the ECS may encounter a release delay. When you attach the disk to the server again, the mount point displayed on the management console may be inconsistent with that in the server. For example, device name **/dev/sdb** or **/dev/vdb** is selected for attachment, but **/dev/xvdc** or **/dev/vdc** may be displayed as the disk name in the OS. This issue does not adversely affect services.

```
[root@ecs-b656 test]# fdisk -l
```

```
Disk /dev/xvda: 42.9 GB, 42949672960 bytes, 83886080 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x000cc4ad
```

Device	Boot	Start	End	Blocks	Id	System
/dev/xvda1	*	2048	2050047	1024000	83	Linux
/dev/xvda2		2050048	22530047	10240000	83	Linux
/dev/xvda3		22530048	24578047	1024000	83	Linux
/dev/xvda4		24578048	83886079	29654016	5	Extended
/dev/xvda5		24580096	26628095	1024000	82	Linux swap / Solaris

```
Disk /dev/xvdb: 10.7 GB, 10737418240 bytes, 20971520 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
```

 **NOTE**

The capacity displayed here is inconsistent with the capacity of the EVS disk applied for on ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios). The reason is as follows: The capacity of EVS disks is calculated using the unit of GiB (Gibibyte), while the capacity unit in Linux OS is GB (Gigabyte). The GiB is calculated in binary mode, and the GB is calculated in decimal format. 1 GiB = 1,073,741,824 Bytes and 1 GB = 1,000,000,000 Bytes.

Step 2 Run the following command to allocate partitions for the added data disk using fdisk:

fdisk *Newly added data disk*

In this example, **/dev/xvdb** is the newly added data disk.

fdisk /dev/xvdb

Information similar to the following is displayed:

```
[root@ecs-b656 test]# fdisk /dev/xvdb
Welcome to fdisk (util-linux 2.23.2).
Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.
Device does not contain a recognized partition table
Building a new DOS disklabel with disk identifier 0xb00005bd.
Command (m for help):
```

Step 3 Enter **n** and press **Enter**.

Entering **n** creates a partition.

There are two types of disk partitions:

- Choosing **p** creates a primary partition.
- Choosing **e** creates an extended partition.

```
Command (m for help): n
Partition type:
  p   primary (0 primary, 0 extended, 4 free)
  e   extended
```

Step 4 Enter **p** and press **Enter**.

The following describes how to create a primary partition.

Information similar to the following is displayed: (**Partition number** indicates the serial number of the primary partition. The value can be **1** to **4**.)

```
Select (default p): p
Partition number (1-4, default 1):
```

Step 5 Enter the primary partition number **1** and press **Enter**.

For example, select **1** as the partition number.

Information similar to the following is displayed: (**First sector** indicates the first sector number. The value can be **2048** to **20971519**, and the default value is **2048**.)

```
Partition number (1-4, default 1): 1
First sector (2048-20971519, default 2048):
```

Step 6 Press **Enter**.

The default start sector number 2048 is used as an example.

Information similar to the following is displayed: (**Last sector** indicates the last sector number. The value can be from **2048** to **20971519**, and the default value is **20971519**.)

```
First sector (2048-20971519, default 2048):
Using default value 2048
Last sector, +sectors or +size{K,M,G} (2048-20971519, default 20971519):
```

Step 7 Press **Enter**.

The default last sector number 20971519 is used as an example.

Information similar to the following is displayed, indicating that a primary partition is created for a 10 GB data disk.

```
Last sector, +sectors or +size{K,M,G} (2048-20971519, default 20971519):
Using default value 20971519
Partition 1 of type Linux and of size 10 GiB is set
Command (m for help):
```

Step 8 Enter **p** and press **Enter** to view the details about the created partition.

Information similar to the following is displayed: (Details about the **/dev/xvdb1** partition are displayed.)

```
Command (m for help): p
```

```
Disk /dev/xvdb: 10.7 GB, 10737418240 bytes, 20971520 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0xb00005bd
```

Device	Boot	Start	End	Blocks	Id	System
/dev/xvdb1		2048	20971519	10484736	83	Linux

Command (m for help):

Step 9 Enter **w** and press **Enter** to write the changes into the partition table.

Information similar to the following is displayed: (The partition is successfully created.)

Command (m for help): **w**

The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.

Step 10 Run the following command to synchronize the new partition table to the data disk:

partprobe

Step 11 Run the following command to set the format for the file system of the newly created partition:

mkfs -t *File system format* /dev/xvdb1

For example, run the following command to set the **ext4** file system for the **/dev/xvdb1** partition:

mkfs -t ext4 /dev/xvdb1

Information similar to the following is displayed:

```
[root@ecs-b656 test]# mkfs -t ext4 /dev/xvdb1
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
655360 inodes, 2621184 blocks
131059 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2151677952
80 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632

Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

NOTE

The formatting takes a period of time. Observe the system running status and do not exit.

Step 12 Run the following command to create a mount directory:

mkdir *Mount directory*

/mnt/sdc is used in this example.

mkdir /mnt/sdc

Step 13 Run the following command to mount the new partition to the mount directory created in [Step 12](#):

mount /dev/xvdb1 *Mount directory*

For example, run the following command to mount the newly created partition on **/mnt/sdc**:

mount /dev/xvdb1 /mnt/sdc

Step 14 Run the following command to view the mount result:

df -TH

Information similar to the following is displayed. The newly created **/dev/xvdb1** partition has been mounted on **/mnt/sdc**.

```
[root@ecs-b656 test]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/xvda2      xfs       11G   7.4G  3.2G  71% /
devtmpfs        devtmpfs  4.1G   0  4.1G   0% /dev
tmpfs           tmpfs     4.1G   82k  4.1G   1% /dev/shm
tmpfs           tmpfs     4.1G   9.2M  4.1G   1% /run
tmpfs           tmpfs     4.1G   0  4.1G   0% /sys/fs/cgroup
/dev/xvda3      xfs       1.1G   39M   1.1G   4% /home
/dev/xvda1      xfs       1.1G  131M  915M  13% /boot
/dev/xvdb1      ext4      11G   38M   9.9G   1% /mnt/sdc
```

----End

Setting Automatic Disk Attachment Upon Instance Start

If you require a disk to be automatically attached to an instance when the instance is started, enable automatic disk attachment upon an instance start by referring to operations provided in this section. When enabling automatic disk attachment, you cannot directly specify **/dev/xvdb1** in **/etc/fstab**. This is because the sequence codes of the instance may change during an instance stop or start process. You are advised to use the universally unique identifier (UUID) in **/etc/fstab** to automatically attach the disk at a system start.

NOTE

The UUID of a disk is a character string that uniquely identifies a storage device in a Linux system.

Step 1 Run the following command to query the partition UUID:

blkid *Disk partition*

For example, run the following command to query the UUID of **/dev/xvdb1**:

blkid /dev/xvdb1

Information similar to the following is displayed: (The UUID of **/dev/xvdb1** is displayed.)

```
[root@ecs-b656 test]# blkid /dev/xvdb1  
/dev/xvdb1: UUID="1851e23f-1c57-40ab-86bb-5fc5fc606ffa" TYPE="ext4"
```

Step 2 Run the following command to open the **fstab** file using the vi editor:

vi /etc/fstab

Step 3 Press **i** to enter the editing mode.

Step 4 Move the cursor to the end of the file and press **Enter**. Then add the following information:

UUID=xxx attachment directory file system defaults 0 2

Assuming that the file system is **ext4** and the attachment directory is **/mnt/sdc**.
UUID=1851e23f-1c57-40ab-86bb-5fc5fc606ffa /mnt/sdc ext4 defaults 0 2

NOTICE

After automatic attachment upon instance start is configured, comment out or delete the line in the **fstab** file before detaching the disk. Otherwise, you may fail to access the OS after the disk is detached.

Step 5 Press **Esc**, enter **:wq**, and press **Enter**.

The system saves the configuration and exits the editor.

----End

4.6.4 Initializing a Linux Data Disk (parted)

A data disk attached to an instance or created together with an instance can be used by the instance only after being initialized. This section uses an instance running CentOS 7.0 (64-bit) as an example, and uses the parted partition tool to set up partitions for the data disk. Initialization operations vary with operating systems.

Prerequisites

- You have logged in to an ECS. For details about how to log in to the bare metal server, see **Operation Help Center > Compute > Bare Metal Server > User Guide > Logging In to a BMS**.
- A data disk has been attached to a BMS and has not been initialized.

Mounting Partitions to a Disk

The following example shows how to create a new primary partition on a new data disk that has been attached to an instance. The primary partition will be created using parted and GPT is the default partition style. Furthermore, the partition will be formatted using the ext4 file system, mounted on the **/mnt/sdc** directory, and set to be automatically mounted upon a system start.

Step 1 Run the following command to view information about the added data disk:

lsblk

Information similar to the following is displayed:

 **NOTE**

If you do not log in to the ECS and run the **umount** command but directly detach the **/dev/xvdb** or **/dev/vdb** EVS disk on the management console, the disk name in the ECS may encounter a release delay. When you attach the disk to the server again, the mount point displayed on the management console may be inconsistent with that in the server. For example, device name **/dev/sdb** or **/dev/vdb** is selected for attachment, but **/dev/xvdc** or **/dev/vdc** may be displayed as the disk name in the OS. This issue does not adversely affect services.

```
[root@ecs-centos-70 linux]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
xvda  202:0   0  40G  0 disk
├─xvda1 202:1   0   4G  0 part [SWAP]
└─xvda2 202:2   0  36G  0 part /
xvdb  202:16  0  10G  0 disk
```

The command output indicates that the server contains two disks. **/dev/xvda** is the system disk and **/dev/xvdb** is the new data disk.

Step 2 Run the following command to enter parted to partition the added data disk:

parted *Added data disk*

In this example, **/dev/xvdb** is the newly added data disk.

parted /dev/xvdb

Information similar to the following is displayed:

```
[root@ecs-centos-70 linux]# parted /dev/xvdb
GNU Parted 3.1
Using /dev/xvdb
Welcome to GNU Parted! Type 'help' to view a list of commands.
```

Step 3 Enter **p** and press **Enter** to view the current disk partition style.

Information similar to the following is displayed:

```
(parted) p
Error: /dev/xvdb: unrecognised disk label
Model: Xen Virtual Block Device (xvd)
Disk /dev/xvdb: 10.7GB
Sector size (logical/physical): 512B/512B
Partition Table: unknown
Disk Flags:
```

In the command output, the **Partition Table** value is **unknown**, indicating that the disk partition style is unknown.

 **NOTE**

The capacity displayed here is inconsistent with the capacity of the EVS disk applied for on ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios). The reason is as follows: The capacity of EVS disks is calculated using the unit of GiB (Gibibyte), while the capacity unit in Linux OS is GB (Gigabyte). The GiB is calculated in binary mode, and the GB is calculated in decimal format. 1 GiB = 1,073,741,824 bytes and 1 GB = 1,000,000,000 bytes.

Step 4 Run the following command to set the disk partition style:

mklabel *Disk partition style*

The disk partition styles include MBR and GPT. For example, run the following command to set the partition style to GPT:

mklabel gpt

NOTICE

If you change the disk partition style after the disk has been used, the original data on the disk will be cleared. Therefore, select a proper disk partition style when initializing the disk.

Step 5 Enter **p** and press **Enter** to view the disk partition style.

Information similar to the following is displayed:

```
(parted) mklabel gpt
(parted) p
Model: Xen Virtual Block Device (xvd)
Disk /dev/xvdb: 20971520s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number Start End Size File system Name Flags
```

Step 6 Enter **unit s** and press **Enter** to set the measurement unit of the disk to sector numbers.

Step 7 Enter **mkpart opt 2048s 100%** and press **Enter**.

In the command, **opt** is the name of the new partition, **2048s** indicates the start of the partition, and **100%** indicates the end of the partition. You can plan the number and capacity of disk partitions based on service requirements.

Information similar to the following is displayed:

```
(parted) mkpart opt 2048s 100%
Warning: The resulting partition is not properly aligned for best performance.
Ignore/Cancel? Cancel
```

If the preceding warning message is displayed, enter **Cancel** to stop the partitioning. Then, find the first sector with the best disk performance and use that value to partition the disk.

Step 8 Enter **p** and press **Enter** to view the details about the created partition.

Information similar to the following is displayed:

```
(parted) p
Model: Xen Virtual Block Device (xvd)
Disk /dev/xvdb: 20971520s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number Start End Size File system Name Flags
1 2048s 20969471s 20967424s opt
```

Details about the **/dev/xvdb1** partition are displayed.

Step 9 Enter **q** and press **Enter** to exit parted.

Step 10 Run the following command to view the disk partition information:

lsblk

Information similar to the following is displayed:

```
[root@ecs-centos-70 linux]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
xvda 202:0 0 40G 0 disk
├─xvda1 202:1 0 4G 0 part [SWAP]
└─xvda2 202:2 0 36G 0 part /
xvdb 202:16 0 100G 0 disk
└─xvdb1 202:17 0 100G 0 part
```

In the command output, **/dev/xvdb1** is the partition you created.

Step 11 Run the following command to set the format for the file system of the newly created partition:

NOTICE

The partition sizes supported by file systems vary. Therefore, you are advised to choose an appropriate file system based on your service requirements.

mkfs -t *File system format* **/dev/xvdb1**

For example, run the following command to set the **ext4** file system for the **/dev/xvdb1** partition:

mkfs -t ext4 /dev/xvdb1

Information similar to the following is displayed:

```
[root@ecs-centos-70 linux]# mkfs -t ext4 /dev/xvdb1
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
655360 inodes, 2620928 blocks
131046 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2151677925
80 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632

Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

The formatting takes a period of time. Observe the system running status, and do not exit.

Step 12 Run the following command to create a mount point:

mkdir *Mount point*

For example, run the following command to create the **/mnt/sdc** mount point:

mkdir /mnt/sdc

Step 13 Run the following command to mount the new partition to the mount point created in [Step 12](#):

mount /dev/xvdb1 *Mount point*

For example, run the following command to mount the newly created partition on **/mnt/sdc**:

```
mount /dev/xvdb1 /mnt/sdc
```

Step 14 Run the following command to view the mount result:

```
df -TH
```

Information similar to the following is displayed:

```
[root@ecs-centos-70 linux]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/xvda2      xfs       39G   4.0G   35G   11% /
devtmpfs        devtmpfs  946M    0  946M    0% /dev
tmpfs           tmpfs     954M    0  954M    0% /dev/shm
tmpfs           tmpfs     954M   9.1M   945M    1% /run
tmpfs           tmpfs     954M    0  954M    0% /sys/fs/cgroup
/dev/xvdb1      ext4      11G   38M   101G    1% /mnt/sdc
```

The newly created **/dev/xvdb1** is mounted on **/mnt/sdc**.

----End

Setting Automatic Disk Attachment Upon Instance Start

If you require a disk to be automatically attached to an instance when the instance is started, enable automatic disk attachment upon an instance start by referring to operations provided in this section. When enabling automatic disk attachment, you cannot directly specify **/dev/xvdb1** in **/etc/fstab**. This is because the sequence codes of the instance may change during an instance stop or start process. You are advised to use the universally unique identifier (UUID) in **/etc/fstab** to automatically attach the disk at a system start.

NOTE

The UUID of a disk is a character string that uniquely identifies a storage device in a Linux system.

Step 1 Run the following command to query the partition UUID:

```
blkid Disk partition
```

For example, run the following command to query the UUID of **/dev/xvdb1**:

```
blkid /dev/xvdb1
```

Information similar to the following is displayed: (The UUID of **/dev/xvdb1** is displayed.)

```
[root@ecs-b656 test]# blkid /dev/xvdb1
/dev/xvdb1: UUID="1851e23f-1c57-40ab-86bb-5fc5fc606ffa" TYPE="ext4"
```

Step 2 Run the following command to open the **fstab** file using the vi editor:

```
vi /etc/fstab
```

Step 3 Press **i** to enter the editing mode.

Step 4 Move the cursor to the end of the file and press **Enter**. Then add the following information:

```
UUID=xxx attachment directory file system defaults 0 2
```

Assuming that the file system is **ext4** and the attachment directory is **/mnt/sdc**.
UUID=1851e23f-1c57-40ab-86bb-5fc5fc606ffa /mnt/sdc ext4 defaults 0 2

NOTICE

After automatic attachment upon instance start is configured, comment out or delete the line in the **fstab** file before detaching the disk. Otherwise, you may fail to access the OS after the disk is detached.

Step 5 Press **Esc**, enter **:wq**, and press **Enter**.

The system saves the configuration and exits the vi editor.

----End

4.7 Expanding EVS Disk Capacity

4.7.1 Overview

You can expand EVS disk capacity if the disk space becomes insufficient. The capacity of a data disk can be expanded based on the original disk. You can also refer to [Adding a Data Disk](#) to add a data disk and attach it to an instance.

After expanding the capacity of a disk, you need to create partitions for the new capacity or create partitions to replace the original disk partitions.

- Creating partitions for the additional space
When using this method, you need to create partitions, without unmounting any existing partitions. This operation does not interrupt ongoing services and has minor impacts on services. This method is recommended for system disks or disks carrying services that cannot be interrupted. If the MBR partition style is used, the disk capacity must be less than 2 TB and the number of partitions does not exceed the upper limit after the expansion.
- Creating partitions to replace existing ones
If the MBR style is used and the number of disk partitions has reached the upper limit, new partitions cannot be created. In this situation, you need to unmount existing partitions and create new ones to replace them. This operation does not delete data in existing partitions, but services must be suspended during the operation, affecting ongoing services.

[Table 4-13](#) describes the post-expansion operations.

Table 4-13 Post-expansion operations for an EVS disk

Instance OS	Disk Capacity After Expansion	Current Disk Partition Style	Post-Expansion Operation
Linux	≤ 2 TB	MBR	Use the fdisk or parted tool to add partitions for the additional disk space. Operations After Expanding the Capacity of a Disk in Linux (Adding Partitions Using fdisk) Operations After Expanding the Capacity of a Disk in Linux (Adding Partitions Using parted)
			Use the fdisk or parted tool to create new partitions to replace existing ones. This operation interrupts ongoing user services and is not recommended after system disk capacity expansion. Operations After Expanding the Capacity of a Disk in Linux (Replacing the Original Partitions Using fdisk) Operations After Expanding the Capacity of a Disk in Linux (Replacing the Original Partitions Using parted)
		GPT	Use the parted tool to divide and initialize the additional disk space.
	> 2 TB	MBR	Use the parted tool to change the partition style from MBR to GPT. However, this operation will clear disk data.
		GPT	Use the parted tool to divide and initialize the additional disk space.
	-	MBR	Allocate the additional disk space to existing partitions.
Windows		GPT	Operations After Expanding the Capacity of a Disk in Windows

4.7.2 Online Disk Capacity Expansion

This section describes how to expand the capacity of an EVS disk attached to an instance.

Restrictions

- When you expand the capacity of a disk online, the instance to which the disk is attached must be in the **Running** or **Stopped** state.
- Shared EVS disks do not support online capacity expansion, that is, the capacity of a shared EVS disk can be expanded only when the disk is in the **Available** state.
- The capacity of a disk configured with the disaster recovery service (CSHA, CSDR, or VHA) cannot be expanded.
- When the backend storage is Huawei SAN storage (OceanStor V3/V5 series and OceanStor Dorado V3/6.x series) or heterogeneous storage, if the EVS disk has snapshots, capacity expansion is not supported. When the backend storage is Huawei Distributed Block Storage, capacity expansion can be performed for an EVS disk with snapshots.
- If backend storage of the disk is heterogeneous storage, online capacity expansion is not supported while offline capacity expansion is supported.
- For the OSs that support online expansion, see [Table 4-14](#).

Table 4-14 OSs

OS	Version
CentOS	7.3 64-bit
	7.2 64-bit
	6.8 64-bit
	6.7 64-bit
	6.5 64-bit
Debian	8.6.0 64-bit
	8.5.0 64-bit
Fedora	25 64-bit
	24 64-bit
SUSE	SUSE Linux Enterprise Server 12 SP2 (64-bit)
	SUSE Linux Enterprise Server 12 SP1 (64-bit)
	SUSE Linux Enterprise Server 11 SP4 (64-bit)
	SUSE Linux Enterprise Server 12 (64-bit)
OpenSUSE	42.2 64-bit
	42.1 64-bit
Oracle Linux Server release	7.3 64-bit
	7.2 64-bit
	6.8 64-bit

OS	Version
	6.7 64-bit
Ubuntu Server	16.04 64-bit
	14.04 64-bit
	14.04.4 64-bit
Windows	Windows Server 2008 R2 Enterprise 64-bit
	Windows Server 2012 R2 Standard 64-bit
	Windows Server 2016 R2 Standard 64-bit
Red Hat Enterprise Linux (RHEL)	7.3 64-bit
	6.8 64-bit

Procedure

Step 1 Log in to the EVS Console as a VDC administrator or VDC operator. For details, see [Logging In to the EVS Console as a VDC Administrator or VDC Operator](#).

Step 2 In the EVS disk list, locate the row that contains the target disk, click **More** in the **Operation** column, and choose **Expand Capacity**.

The **Expand Capacity** page is displayed.

NOTE

To view information about the instance to which the EVS disk is attached, click the instance name in the **Attaching Information** column.

Step 3 Set **Added Capacity (GB)** as prompted and click **Next**.

Step 4 On the **Confirm Specifications** page, confirm the disk specifications.

- If you do not need to modify the specifications, click **Apply Now** to start the EVS disk capacity expansion.
- If you need to modify the specifications, click **Previous** to modify parameters.

After the expansion is submitted, the disk list page is displayed.

Step 5 If the EVS disk whose capacity is to be expanded requires approval, contact the administrator for approval. Otherwise, skip this step.

On the **Elastic Volume Service** page, view the capacity of the EVS disk. If the disk capacity has increased, the expansion is successful.

----End

Follow-up Procedure

After you have expanded the capacity of an EVS disk, perform follow-up operations for the additional capacity.

 **NOTE**

If the instance to which the EVS disk is attached is stopped during capacity expansion, power on the instance before capacity expansion.

- For Windows, see [Post-Expansion Operations for a Windows EVS Disk](#).
- For Linux, refer to [Table 4-13](#) to select a proper post-expansion processing method.

4.7.3 Offline Disk Capacity Expansion

This section describes how to expand the capacity of an EVS disk not attached to any BMS.

Restrictions

- The capacity of a disk configured with the DR service (CSHA/CSDR/VHA) cannot be expanded.
- When the backend storage is Huawei SAN storage (OceanStor V3/V5 series or OceanStor Dorado V3 series/6.x) or heterogeneous storage, if the EVS disk has snapshots, capacity expansion is not supported. When the backend storage is Huawei Distributed Block Storage, capacity expansion can be performed for an EVS disk with snapshots.
- Capacity expansion cannot be performed when the disk status is **Reserved** or **Maintenance**.

Procedure

Step 1 Log in to the EVS Console as a VDC administrator or VDC operator. For details, see [Logging In to the EVS Console as a VDC Administrator or VDC Operator](#).

Step 2 If the to-be-expanded EVS disk has been attached to a BMS, detach it first. For details, see section [Detaching a Data Disk](#).

When the disk status changes to **Available**, the disk is successfully detached.

Step 3 In the EVS disk list, locate the row that contains the target disk, click **More** in the **Operation** column, and choose **Expand Capacity**.

The **Expand Capacity** page is displayed.

Step 4 Set **Added Capacity (GB)** as prompted and click **Next**.

Step 5 On the **Confirm Specifications** page, confirm the disk specifications.

- If you do not need to modify the specifications, click **Apply Now** to start the EVS expansion.
- If you need to modify the specifications, click **Previous** to modify parameters.

After the expansion is submitted, the disk list page is displayed.

Step 6 If the EVS disk whose capacity is to be expanded requires approval, contact the administrator for approval. Otherwise, skip this step.

On the **Elastic Volume Service** page, view the capacity of the EVS disk. If the disk capacity has increased, the expansion is successful.

----End

Follow-up Procedure

After you have expanded the capacity of an EVS disk, perform follow-up operations on the additional volume.

NOTE

If the BMS to which the EVS disk is attached is stopped during capacity expansion, power on the BMS before capacity expansion.

1. Attach the expanded EVS disk to a BMS. For details, see section [Attaching an EVS Disk](#).
2. Attach an EVS disk to a BMS.
 - For Windows, see [Post-Expansion Operations for a Windows EVS Disk](#).
 - For Linux, select a post-expansion processing mode by referring to [Table 4-13](#).

4.7.4 Post-Expansion Operations for a Windows EVS Disk

This section describes how to perform follow-up operations for the additional disk space after you have expanded the capacity of an EVS disk and attached the EVS disk to an instance. This section uses Windows Server 2008 R2 Enterprise as the example OS to describe how to perform post-expansion operations for an EVS disk. For other Windows OSs, see the corresponding OS product documents.

Prerequisites

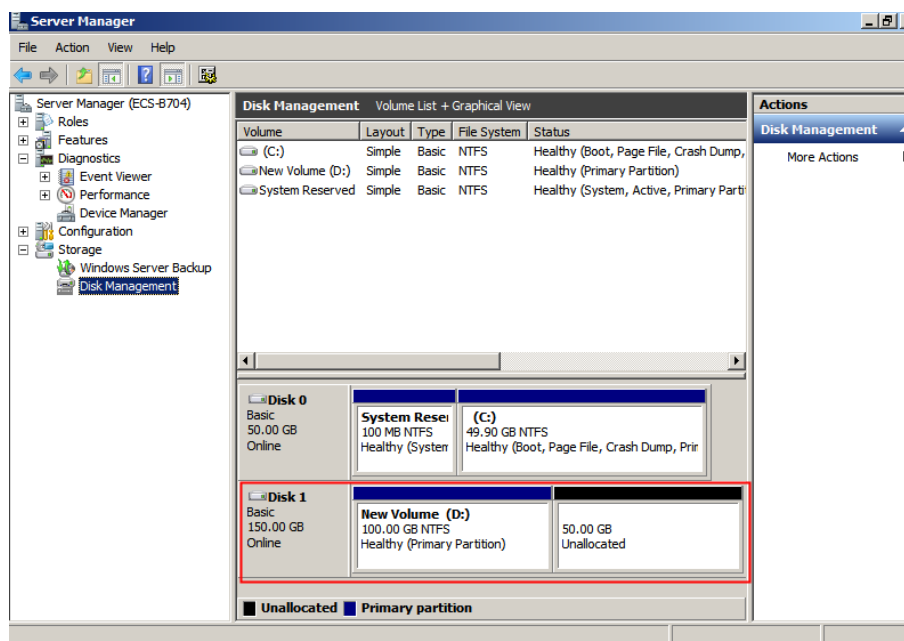
- You have logged in to an ECS. For details about how to log in to the bare metal server, see **Operation Help Center > Compute > Bare Metal Server > User Guide > Logging In to a BMS**.
- A data disk has been attached to a BMS and has not been initialized.

Context

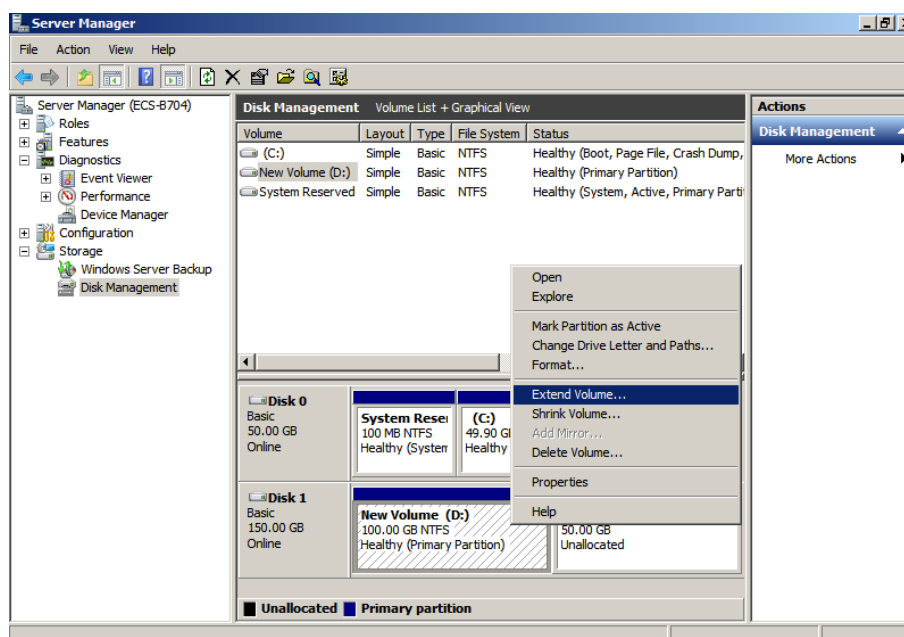
After the capacity expansion has succeeded, the new EVS disk space needs to be allocated to an existing partition or a new partition. This section describes how to allocate new EVS disk space to an existing partition.

Procedure

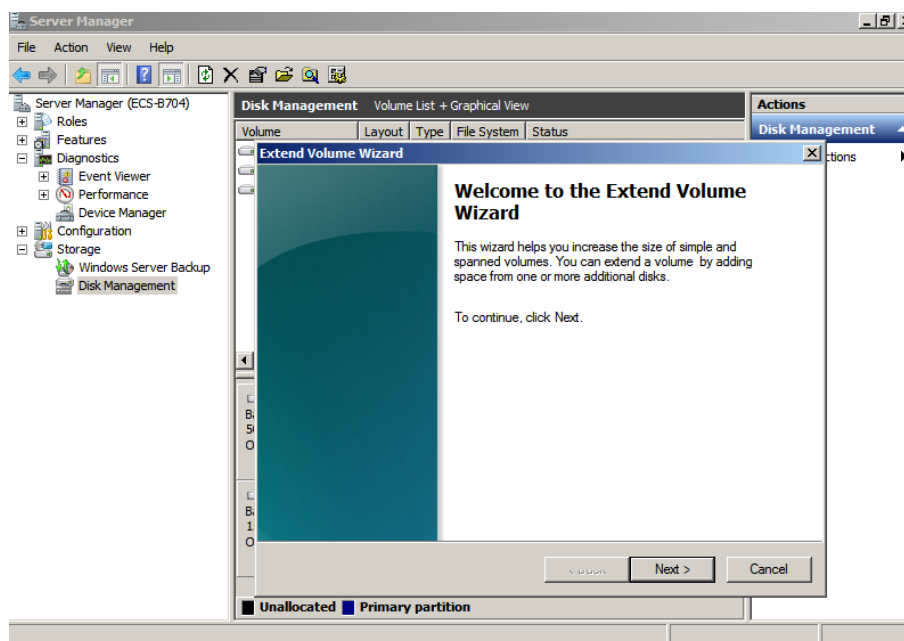
- Step 1** In desktop, right-click **Computer** and choose **Manage** from the shortcut menu.
- The **Server Manager** window is displayed.
- Step 2** In the navigation tree, choose **Storage > Disk Management**.
- The **Disk Management** page is displayed.
- Step 3** On the **Disk Management** page, select the disk and partition that needs to be extended. The current partition size and unallocated disk space are displayed.



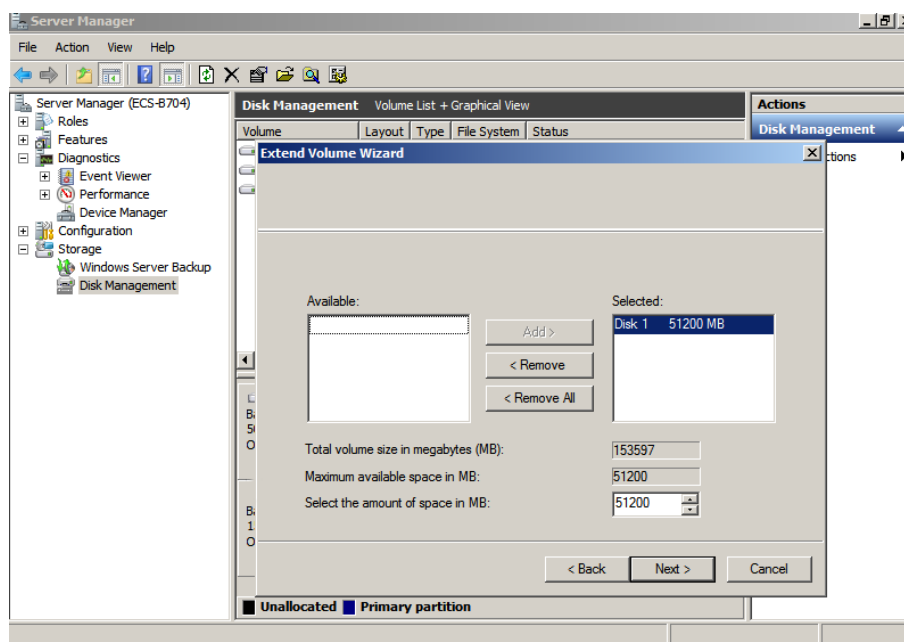
Step 4 Right-click the selected disk and choose **Extend Volume**.



Step 5 On the displayed **Extend Volume Wizard** page, click **Next**.

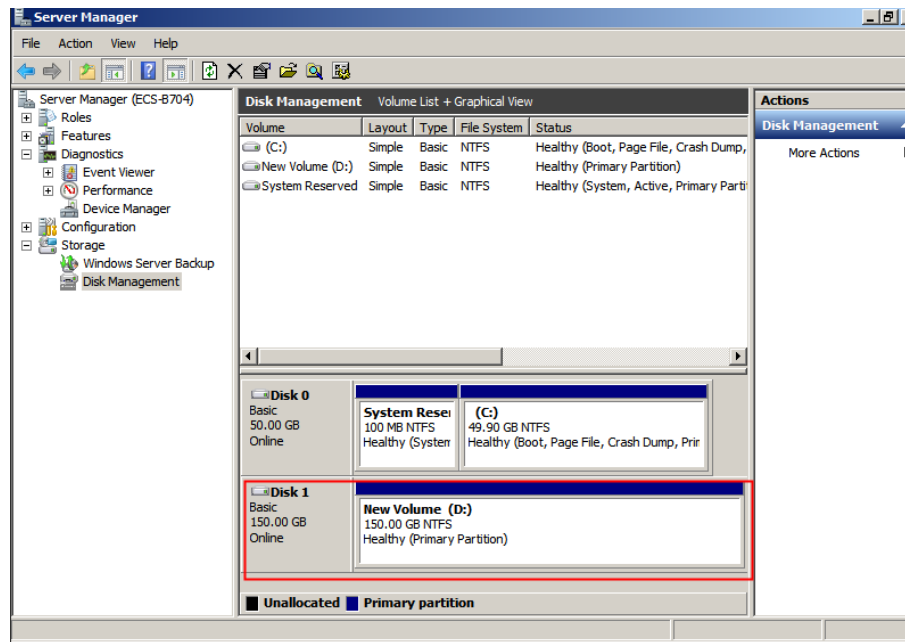


Step 6 In the text box to the right of **Select the amount of space in MB**, enter the extended capacity and click **Next**.



Step 7 Click **Finish** to complete the wizard.

After the expansion is successful, the disk capacity is greater than the original capacity.



----End

4.7.5 Operations After Expanding Disk Capacity in Linux (Adding Partitions Using fdisk)

After the capacity of an EVS disk is expanded and the disk is attached to an instance, the additional space needs to be allocated to a new partition. This section uses CentOS 7.0 (64-bit) as an example to describe how to use `fdisk` to create partitions for a disk after capacity expansion.

Prerequisites

- You have logged in to an ECS. For details about how to log in to the bare metal server, see **Operation Help Center > Compute > Bare Metal Server > User Guide > Logging In to a BMS**.
- A data disk has been attached to a BMS and has not been initialized.

Procedure

The steps below describe how to allocate the additional space to a new partition and mount the partition on `/opt`. The disk is partitioned using MBR, the disk capacity is less than 2 TB, and the number of partitions is lower than the upper limit. You can use either the `fdisk` or `parted` tool to partition the disk space. This section uses the `fdisk` tool as an example.

Step 1 Run the following command to query and view the disk information:

```
fdisk -l
```

Information similar to the following is displayed: (`/dev/xvda` is the system disk.)

```
[root@ecs-bab9 test]# fdisk -l
```

```
Disk /dev/xvda: 64.4 GB, 64424509440 bytes, 125829120 sectors  
Units = sectors of 1 * 512 = 512 bytes
```

Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x000cc4ad

Device	Boot	Start	End	Blocks	Id	System
/dev/xvda1	*	2048	2050047	1024000	83	Linux
/dev/xvda2		2050048	22530047	10240000	83	Linux
/dev/xvda3		22530048	24578047	1024000	83	Linux
/dev/xvda4		24578048	83886079	29654016	5	Extended
/dev/xvda5		24580096	26628095	1024000	82	Linux swap / Solaris

NOTE

The capacity displayed here is inconsistent with the capacity of the EVS disk applied for on ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios). The reason is as follows: The capacity of EVS disks is calculated using the unit of GiB (Gibibyte), while the capacity unit in Linux OS is GB (Gigabyte). The GiB is calculated in binary mode, and the GB is calculated in decimal format. 1 GiB = 1,073,741,824 Bytes and 1 GB = 1,000,000,000 Bytes.

Step 2 Run the following command to enter the fdisk (/dev/xvda is used as an example):

fdisk /dev/xvda

Information similar to the following is displayed:

```
[root@ecs-bab9 test]# fdisk /dev/xvda
Welcome to fdisk (util-linux 2.23.2).
```

```
Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.
```

```
Command (m for help):
```

Step 3 Enter **n** and press **Enter** to create a partition.

Because the system disk has five existing partitions, the system automatically creates the sixth one.

Information similar to the following is displayed:

```
Command (m for help): n
All primary partitions are in use
Adding logical partition 6
First sector (26630144-83886079, default 26630144):
```

Step 4 Enter the new partition's first sector number, for example the default value, and press **Enter**.

The first sector number must be greater than the last sector numbers of existing partitions.

Information similar to the following is displayed:

```
First sector (26630144-83886079, default 26630144):
Using default value 26630144
Last sector, +sectors or +size{K,M,G} (26630144-83886079, default 83886079):
```

Step 5 Enter the new partition's last sector number and press **Enter**.

The default last sector number is used as an example.

Information similar to the following is displayed:

```
Last sector, +sectors or +size{K,M,G} (26630144-83886079, default 83886079):
Using default value 83886079
```

Partition 6 of type Linux and of size 27.3 GiB is set

Command (m for help):

Step 6 Enter **p** and press **Enter** to view the created partition.

Information similar to the following is displayed:

Disk /dev/xvda: 64.4 GB, 64424509440 bytes, 125829120 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x000cc4ad

Device	Boot	Start	End	Blocks	Id	System
/dev/xvda1	*	2048	2050047	1024000	83	Linux
/dev/xvda2		2050048	22530047	10240000	83	Linux
/dev/xvda3		22530048	24578047	1024000	83	Linux
/dev/xvda4		24578048	83886079	29654016	5	Extended
/dev/xvda5		24580096	26628095	1024000	82	Linux swap / Solaris
/dev/xvda6		26630144	83886079	28627968	83	Linux

Command (m for help):

Step 7 Enter **w** and press **Enter**.

The partition result is written into the partition table, and the partitioning is complete.

Information similar to the following is displayed:

Command (m for help): **w**

The partition table has been altered!

Calling ioctl() to re-read partition table.

WARNING: Re-reading the partition table failed with error 16: Device or resource busy.
The kernel still uses the old table. The new table will be used at
the next reboot or after you run partprobe(8) or kpartx(8)
Syncing disks.

Step 8 Run the following command to synchronize the new partition table to the OS:

partprobe

Step 9 Run the following command to set the format for the file system of the new partition:

In this example, the ext4 file system is used.

mkfs -t ext4 /dev/xvda6

Information similar to the following is displayed:

```
[root@ecs-bab9 test]# mkfs -t ext4 /dev/xvda6
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
1790544 inodes, 7156992 blocks
357849 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2155872256
219 block groups
32768 blocks per group, 32768 fragments per group
```

```
8176 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000
```

```
Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

The formatting takes a while, and you need to observe the system running status. Once **done** is displayed in the command output, the formatting is complete.

- Step 10** Run the following command to mount the new partition to the space-demanding directory, for example **/opt**:

```
mount /dev/xvda6 /opt
```

Information similar to the following is displayed:

```
[root@ecs-bab9 test]# mount /dev/xvda6 /opt
[root@ecs-bab9 test]#
```

NOTE

If the new partition is mounted to a directory that is not empty, the subdirectories and files in the directory will be hidden. In this situation, you are advised to mount the new partition to an empty directory or a new directory. If the new partition must be mounted to a directory that is not empty, move the subdirectories and files in this directory to another directory temporarily. After the partition is mounted, move the subdirectories and files back.

- Step 11** Run the following command to view the mount result:

```
df -TH
```

Information similar to the following is displayed:

```
[root@ecs-bab9 test]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/xvda2      xfs       11G   7.4G  3.2G  71% /
devtmpfs        devtmpfs  4.1G   0  4.1G   0% /dev
tmpfs           tmpfs     4.1G  82k  4.1G   1% /dev/shm
tmpfs           tmpfs     4.1G  9.2M  4.1G   1% /run
tmpfs           tmpfs     4.1G   0  4.1G   0% /sys/fs/cgroup
/dev/xvda3      xfs       1.1G   39M  1.1G   4% /home
/dev/xvda1      xfs       1.1G  131M  915M  13% /boot
/dev/xvda6      ext4      29G   47M  28G   1% /opt
```

----End

Setting Automatic Disk Attachment Upon Instance Start

If you require a disk to be automatically attached to an instance when the instance is started, enable automatic disk attachment upon an instance start by referring to operations provided in this section. When enabling automatic disk attachment, you cannot directly specify **/dev/xvdb1** in **/etc/fstab**. This is because the sequence codes of the instance may change during an instance stop or start process. You are advised to use the universally unique identifier (UUID) in **/etc/fstab** to automatically attach the disk at a system start.

NOTE

The UUID of a disk is a character string that uniquely identifies a storage device in a Linux system.

Step 1 Run the following command to query the partition UUID:

blkid *Disk partition*

For example, run the following command to query the UUID of **/dev/xvdb1**:

blkid /dev/xvdb1

Information similar to the following is displayed: (The UUID of **/dev/xvdb1** is displayed.)

```
[root@ecs-b656 test]# blkid /dev/xvdb1
/dev/xvdb1: UUID="1851e23f-1c57-40ab-86bb-5fc5fc606ffa" TYPE="ext4"
```

Step 2 Run the following command to open the **fstab** file using the vi editor:

vi /etc/fstab

Step 3 Press **i** to enter the editing mode.

Step 4 Move the cursor to the end of the file and press **Enter**. Then add the following information:

UUID=xxx attachment directory file system defaults 0 2

Assuming that the file system is **ext4** and the attachment directory is **/mnt/sdc**.

```
UUID=1851e23f-1c57-40ab-86bb-5fc5fc606ffa /mnt/sdc ext4 defaults 0 2
```

NOTICE

After automatic attachment upon instance start is configured, comment out or delete the line in the **fstab** file before detaching the disk. Otherwise, you may fail to access the OS after the disk is detached.

Step 5 Press **Esc**, enter **:wq**, and press **Enter**.

The system saves the configuration and exits the vi editor.

----End

4.7.6 Operations After Expanding Disk Capacity in Linux (Adding Partitions Using parted)

This section describes how to perform follow-up operations for the additional disk space after you have expanded the capacity of an EVS disk and attached the EVS disk to an instance. This section uses EulerOS 2.0 (64-bit) to describe how to allocate the additional disk space to a partition using parted.

Prerequisites

- You have logged in to an ECS. For details about how to log in to the bare metal server, see **Operation Help Center > Compute > Bare Metal Server > User Guide > Logging In to a BMS**.
- A data disk has been attached to a BMS and has not been initialized.

Procedure

The steps below describe how to allocate the additional space to a new partition and mount the partition on **/opt**. The disk is partitioned using MBR, the disk capacity is less than 2 TB, and the number of partitions is lower than the upper limit. You can use either the **fdisk** or **parted** tool to partition the disk space. This section uses the **parted** tool as an example.

Step 1 Run the following command to view the disk partition information:

lsblk

If the following information is displayed, the total capacity of the current system disk **dev/xvda** is 80 GB. 40 GB of the disk has been allocated to partitions, and the remaining 40 GB is additional disk space and have not been allocated to any partition.

```
[root@ecs-1120 linux]# lsblk
NAME        MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
xvda        202:0    0   80G  0 disk
├─xvda1     202:1    0   40G  0 part /
├─xvdb      202:16   0  250G  0 disk
│ └─xvdb1   202:17   0  100G  0 part
│   └─xvdb2 202:18   0   50G  0 part
├─xvdc      202:32   0   40G  0 disk
│ └─xvdc1   202:33   0    8G  0 part
│   └─xvdc2 202:34   0   32G  0 part
```

Step 2 Run the following command to create a partition for the additional system disk space:

parted *System disk*

parted **/dev/xvda**

The following information is displayed:

```
[root@ecs-1120 linux]# parted /dev/xvda
GNU Parted 3.1
Using /dev/xvda
Welcome to GNU Parted! Type 'help' to view a list of commands.
```

Step 3 Enter **unit s** and press **Enter** to set the measurement unit of the disk to sector numbers.

Step 4 Enter **p** and press **Enter** to view the current disk partition style.

Partition Table specifies the partition style for existing disks. **msdos** indicates that the disk partition style is MBR, and **gpt** indicates that the disk partition style is GPT.

The following information is displayed:

```
(parted) unit s
(parted) p
Model: Xen Virtual Block Device (xvd)
Disk /dev/xvda: 167772160s
Sector size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:

Number Start End      Size      Type     File system  Flags
 1      2048s 83886079s 83884032s primary  ext4
```

Step 5 Enter **mkpart** and press **Enter** to create a partition.

Step 6 Enter **p** and press **Enter** to create a primary partition. Creating a primary partition is used as an example.

Set the file system format to **ext4**.

Set the first sector and last sector to **83886080** and **167772159** for the **dev/xvda2** partition, respectively.

Set the parameters as required.

The following information is displayed:

```
(parted) mkpart
Partition type? primary/extended? p
File system type? [ext2]? ext4
Start? 83886080
End? 167772159
```

NOTE

The file system type may fail to set in this step. You can reconfigure the file system format according to [Step 9](#).

Step 7 Enter **p** and press **Enter** to view the created partition.

Information similar to the following is displayed:

```
(parted) p
Model: Xen Virtual Block Device (xvd)
Disk /dev/xvda: 167772160s
Sector size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:

Number Start      End          Size         Type      File system  Flags
 1      2048s    83886079s   83884032s   primary  ext4
 2     83886080s 167772159s  83886080s   primary
```

Step 8 Enter **q** and press **Enter** to exit parted.

Step 9 Run the following command to set the format for the file system of the new partition:

In this example, the ext4 file system is used.

mkfs -t ext4 /dev/xvda2

Information similar to the following is displayed:

Wait for the formatting. If **done** is displayed in the command output, the formatting has been complete.

```
[[root@ecs-1120 linux]# mkfs -t ext4 /dev/xvda2
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
2621440 inodes, 10485760 blocks
524288 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2157969408
320 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
?32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
```

```
?4096000, 7962624
```

```
Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

Step 10 Run the following command to mount the new partition to the space-demanding directory, for example **/opt**:

```
mount /dev/xvda2 /opt
```

NOTE

If the new partition is mounted to a directory that is not empty, the subdirectories and files in the directory will be hidden. In this situation, you are advised to mount the new partition to an empty directory or a new directory. If the new partition must be mounted to a directory that is not empty, move the subdirectories and files in this directory to another directory temporarily. After the partition is mounted, move the subdirectories and files back.

Step 11 Run the following command to view the mount result:

```
df -TH
```

Information similar to the following is displayed:

```
[root@ecs-1120 linux]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/xvda1      ext4      43G   8.3G   33G   21% /
devtmpfs        devtmpfs  885M    0 885M    0% /dev
tmpfs           tmpfs     894M    0 894M    0% /dev/shm
tmpfs           tmpfs     894M   18M  877M    2% /run
tmpfs           tmpfs     894M    0 894M    0% /sys/fs/cgroup
tmpfs           tmpfs     179M    0 179M    0% /run/user/2000
tmpfs           tmpfs     179M    0 179M    0% /run/user/0
tmpfs           tmpfs     179M    0 179M    0% /run/user/1001
/dev/xvda2      ext4      43G   51M   40G    1% /opt
```

----End

Setting Automatic Disk Attachment Upon Instance Start

If you require a disk to be automatically attached to an instance when the instance is started, enable automatic disk attachment upon an instance start by referring to operations provided in this section. When enabling automatic disk attachment, you cannot directly specify **/dev/xvdb1** in **/etc/fstab**. This is because the sequence codes of the instance may change during an instance stop or start process. You are advised to use the universally unique identifier (UUID) in **/etc/fstab** to automatically attach the disk at a system start.

NOTE

The UUID of a disk is a character string that uniquely identifies a storage device in a Linux system.

Step 1 Run the following command to query the partition UUID:

```
blkid Disk partition
```

For example, run the following command to query the UUID of **/dev/xvdb1**:

```
blkid /dev/xvdb1
```

Information similar to the following is displayed: (The UUID of **/dev/xvdb1** is displayed.)

```
[root@ecs-b656 test]# blkid /dev/xvdb1
/dev/xvdb1: UUID="1851e23f-1c57-40ab-86bb-5fc5fc606ffa" TYPE="ext4"
```

Step 2 Run the following command to open the **fstab** file using the vi editor:

```
vi /etc/fstab
```

Step 3 Press **i** to enter the editing mode.

Step 4 Move the cursor to the end of the file and press **Enter**. Then add the following information:

```
UUID=xxx attachment directory file system defaults 0 2
```

Assuming that the file system is **ext4** and the attachment directory is **/mnt/sdc**.

```
UUID=1851e23f-1c57-40ab-86bb-5fc5fc606ffa /mnt/sdc ext4 defaults 0 2
```

NOTICE

After automatic attachment upon instance start is configured, comment out or delete the line in the **fstab** file before detaching the disk. Otherwise, you may fail to access the OS after the disk is detached.

Step 5 Press **Esc**, enter **:wq**, and press **Enter**.

The system saves the configuration and exits the vi editor.

----End

4.7.7 Operations After Expanding Disk Capacity in Linux (Replacing Original Partitions Using fdisk)

This section describes how to perform follow-up operations for the additional disk space after you have expanded the capacity of an EVS disk and attached the EVS disk to an instance.

NOTICE

For details, contact the OS vendor of the target VM. This section uses EulerOS 2.0 (64-bit) to describe how to allocate the additional disk space to a partition using fdisk.

Prerequisites

- You have logged in to an ECS. For details about how to log in to the bare metal server, see **Operation Help Center > Compute > Bare Metal Server > User Guide > Logging In to a BMS**.
- A data disk has been attached to a BMS and has not been initialized.

Procedure

The following example shows how to create new partitions to replace existing ones on a disk that has been attached to an instance. In this example, the disk partition is **/dev/xvdb1**, the partition style is MBR, and the mount directory is **/mnt/sdc**. After capacity expansion, the disk capacity is less than 2 TB. You can use either the **fdisk** or **parted** tool to partition the disk space. This section uses the **fdisk** tool as an example.

Step 1 Run the following command to view the disk partition information:

fdisk -l

Information similar to the following is displayed: (In the command output, the server contains two disks. **/dev/xvda** is the system disk, and **/dev/xvdb** is the data disk.)

```
[root@ecs-b656 test]# fdisk -l
```

```
Disk /dev/xvda: 42.9 GB, 42949672960 bytes, 83886080 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x000cc4ad
```

Device	Boot	Start	End	Blocks	Id	System
/dev/xvda1	*	2048	2050047	1024000	83	Linux
/dev/xvda2		2050048	22530047	10240000	83	Linux
/dev/xvda3		22530048	24578047	1024000	83	Linux
/dev/xvda4		24578048	83886079	29654016	5	Extended
/dev/xvda5		24580096	26628095	1024000	82	Linux swap / Solaris

```
Disk /dev/xvdb: 24.7 GB, 24696061952 bytes, 48234496 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0xb00005bd
```

Device	Boot	Start	End	Blocks	Id	System
/dev/xvdb1		2048	20971519	10484736	83	Linux

In the command output, parameter **Disk label type** indicates the disk partition style. Value **dos** indicates the MBR partition style, and value **gpt** indicates the GPT partition style.

View the **/dev/xvdb** capacity and check whether the additional space is included.

NOTE

The capacity displayed here is inconsistent with the capacity of the EVS disk applied for on ManageOne Operation Portal (ManageOne Tenant Portal in B2B scenarios). The reason is as follows: The capacity of EVS disks is calculated using the unit of GiB (Gibibyte), while the capacity unit in Linux OS is GB (Gigabyte). The GiB is calculated in binary mode, and the GB is calculated in decimal format. 1 GiB = 1,073,741,824 bytes and 1 GB = 1,000,000,000 bytes.

- If the additional space is displayed, go to [Step 2](#).
- If the new capacity is not displayed in the command output, run the following command to update the disk capacity:

```
echo 1 > /sys/class/scsi_device/%d:%d:%d:%d/device/rescan &
```

In the command, `%d:%d:%d:%d` indicates a folder in the `/sys/class/scsi_device/` directory and can be obtained using `ll /sys/class/scsi_device/`.

Information similar to the following is displayed: `(2:0:0:0)` indicates the folder to be obtained.)

```
cs-xen-02:/sys/class/scsi_device # ll /sys/class/scsi_device/
total 0
lrwxrwxrwx 1 root root 0 Sep 26 11:37 2:0:0:0-> ../../devices/xen/vscsi-2064/host2/target2:0:0/2:0:0:0/scsi_device/2:0:0:0
```

Example command:

echo 1 > /sys/class/scsi_device/2:0:0:0/device/rescan &

After the update is complete, run the **fdisk -l** command to view the disk partition information.

Step 2 Run the following command to unmount the disk:

umount /mnt/sdc

Step 3 Run the following command and enter **d** and then **1** to delete existing partition `/dev/xvdb1`. The partition to be deleted varies between scenarios.

fdisk /dev/xvdb

The command output is as follows:

```
[root@ecs-b656 test]# fdisk /dev/xvdb
Welcome to fdisk (util-linux 2.23.2).

Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.

Command (m for help): d
Selected partition 1
Partition 1 is deleted

Command (m for help):
```

NOTE

Deleting partitions will not cause data losses on the data disk.

Step 4 Enter **n** and press **Enter** to create a partition.

Entering **n** creates a partition.

There are two types of disk partitions:

- Choosing **p** creates a primary partition.
- Choosing **e** creates an extended partition.

```
Command (m for help): n
Partition type:
 p primary (0 primary, 0 extended, 4 free)
 e extended
```

Step 5 Ensure that the entered partition type is the same as that of the partition to be replaced. In this example, creating a primary partition is used. Therefore, enter **p** and press **Enter** to create a primary partition.

Information similar to the following is displayed: (**Partition number** indicates the serial number of the primary partition. The value can be **1** to **4**.)

```
Select (default p): p
Partition number (1-4, default 1):
```

- Step 6** Ensure that entered partition number is the same as that of the partition to be replaced. In this example, the partition number **1** is used. Therefore, enter **1** and press **Enter**.

Information similar to the following is displayed: (**First sector** indicates the first sector number. The value can be **2048** to **20971519**, and the default value is **2048**.)

```
Partition number (1-4, default 1): 1
First sector (2048-20971519, default 2048):
```

- Step 7** Press **Enter**.

The default start sector number is used as an example.

Information similar to the following is displayed: (**Last sector** indicates the last sector number. The value can be **2048** to **20971519**, and the default value is **20971519**.)

```
First sector (2048-20971519, default 2048):
Using default value 2048
Last sector, +sectors or +size{K,M,G} (2048-20971519, default 20971519):
```

- Step 8** Press **Enter**.

The default last sector number is used as an example.

Information similar to the following is displayed, indicating that the partition is successfully created.

```
Last sector, +sectors or +size{K,M,G} (2048-20971519, default 20971519):
Using default value 20971519
Partition 1 of type Linux and of size 10 GiB is set
Command (m for help):
```

- Step 9** Enter **p** and press **Enter** to view the details about the created partition.

Information similar to the following is displayed: (Details about the **/dev/xvdb1** partition are displayed.)

```
Command (m for help): p

Disk /dev/xvdb: 10.7 GB, 10737418240 bytes, 20971520 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0xb00005bd
```

Device	Boot	Start	End	Blocks	Id	System
/dev/xvdb1		2048	20971519	10484736	83	Linux

```
Command (m for help):
```

- Step 10** Enter **w** and press **Enter** to write the changes into the partition table.

Information similar to the following is displayed: (The partition is successfully created.)

```
Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.
```

If the following error is displayed when you write partition results to the partition table, the new partition table will be updated upon the next OS restart.

Command (m for help): **w**

The partition table has been altered!

Calling ioctl() to re-read partition table.

WARNING: Re-reading the partition table failed with error 16: Device or resource busy.

The kernel still uses the old table, The new table will be used at the next reboot or after you run partprobe(8) or kpartx(8)
Syncing disks.

Step 11 Run the following commands to check and adjust the size of the file system on **/dev/xvdb1**:

e2fsck -f /dev/xvdb1

Information similar to the following is displayed:

```
[root@ecs-b656 test]# e2fsck -f /dev/xvdb1
e2fsck 1.42.9 (28-Dec-2013)
Pass 1: Checking inodes, blocks, and sizes
Pass 2: Checking directory structure
Pass 3: Checking directory connectivity
Pass 4: Checking reference counts
Pass 5: Checking group summary information
/dev/xvdb1: 11/655360 files (0.0% non-contiguous), 83137/2621184 blocks
```

resize2fs /dev/xvdb1

Information similar to the following is displayed:

```
[root@ecs-b656 test]# resize2fs /dev/xvdb1
resize2fs 1.42.9 (28-Dec-2013)
Resizing the filesystem on /dev/xvdb1 to 6029056 (4k) blocks.
The filesystem on /dev/xvdb1 is now 6029056 blocks long.
```

----End

4.7.8 Operations After Expanding Disk Capacity in Linux (Replacing Original Partitions Using parted)

This section describes how to perform follow-up operations for the additional disk space after you have expanded the capacity of an EVS disk and attached the EVS disk to an instance.

NOTICE

For details, contact the OS vendor of the target VM. This section uses EulerOS 2.0 (64-bit) to describe how to allocate the additional disk space to a partition using parted.

Prerequisites

- You have logged in to an ECS. For details about how to log in to the bare metal server, see **Operation Help Center > Compute > Bare Metal Server > User Guide > Logging In to a BMS**.
- A data disk has been attached to a BMS and has not been initialized.

Procedure

The following example shows how to create a partition to replace the `/dev/xvdc1` partition mounted on `/mnt/sdc`. `/dev/xvdc1` is the only partition of the `/dev/xvdc` disk attached to an instance and uses the GPT partition style. There are two disks attached to the instance. During the partition creation, services will be interrupted.

NOTICE

After the disk capacity has been expanded, the additional space is added to the end of the disk. When the disk has multiple partitions, only the partition at the end of the disk can be expanded.

Step 1 Run the following command to view the disk partition information:

lsblk

Information similar to the following is displayed:

```
[root@ecs-1120 sdc]# lsblk
NAME        MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
xvda        202:0    0   80G  0 disk
├─xvda1     202:1    0   40G  0 part /
├─xvda2     202:2    0   40G  0 part /opt
└─xvdb      202:16   0  350G  0 disk
   ├─xvdb1   202:17   0  100G  0 part
   └─xvdb2   202:18   0  200G  0 part
xvdc        202:32   0   60G  0 disk
└─xvdc1     202:33   0   10G  0 part /mnt/sdc
```

Indicates that the total capacity of the current data disk `/dev/xvdc` is 60 GB and the allocated partition capacity is 10 GB. The last partition is `/dev/xvdc1`, which is attached to the `/mnt/sdc` directory.

View the `/dev/xvdc` capacity and check whether the additional space is included.

- If the additional space is included, go to [Step 2](#).
- If the new capacity is not displayed in the command output, run the following command to update the disk capacity:

```
echo 1 > /sys/class/scsi_device/%d:%d:%d:%d/device/rescan &
```

In the command, `%d:%d:%d:%d` indicates a folder in the `/sys/class/scsi_device/` directory and can be obtained using `ll /sys/class/scsi_device/`.

Information similar to the following is displayed: (`2:0:0:0` indicates the folder to be obtained.)

```
cs-xen-02:/sys/class/scsi_device # ll /sys/class/scsi_device/
total 0
lrwxrwxrwx 1 root root 0 Sep 26 11:37 2:0:0:0-> ../../devices/xen/vscsi-2064/host2/target2:0:0/2:0:0:0/
scsi_device/2:0:0:0
```

Example command:

```
echo 1 > /sys/class/scsi_device/2:0:0:0/device/rescan &
```

After the update is complete, run the `fdisk -l` command to view the disk partition information.

Step 2 Run the following command to unmount the disk partition:

```
umount /mnt/sdc
```

Step 3 Run the following command to view the unmount result:

lsblk

Information similar to the following is displayed:

```
[root@ecs-1120 linux]# umount /mnt/sdc
[root@ecs-1120 linux]# lsblk
NAME        MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
xvda        202:0    0   80G  0 disk
├─xvda1     202:1    0   40G  0 part /
├─xvda2     202:2    0   40G  0 part /opt
xvdb        202:16   0  350G  0 disk
├─xvdb1     202:17   0  100G  0 part
├─xvdb2     202:18   0  200G  0 part
xvdc        202:32   0   60G  0 disk
└─xvdc1     202:33   0   10G  0 part
```

Step 4 Run the following command to enter parted to allocate the additional space of the data disk to a partition:

parted Data disk

In this example, **/dev/xvdc** is the data disk.

parted /dev/xvdc

Information similar to the following is displayed:

```
[root@ecs-1120 linux]# parted /dev/xvdc
GNU Parted 3.1
Using /dev/xvdc
Welcome to GNU Parted! Type 'help' to view a list of commands.
```

Step 5 Enter **unit s** and press **Enter** to set the measurement unit of the disk to sector numbers.

Step 6 Enter **p** and press **Enter** to view the current disk partition style.

Information similar to the following is displayed:

```
(parted) unit s
(parted) p
Error: The backup GPT table is not at the end of the disk, as it should be.
This might mean that another operating system believes the disk is smaller.
Fix, by moving the backup to the end (and removing the old backup)?
Fix/Ignore/Cancel? Fix
Warning: Not all of the space available to /dev/xvdb appears to be used,
you can fix the GPT to use all of the space (an extra 104857600 blocks)
or continue with the current setting?
Fix/Ignore? Fix
Model: Xen Virtual Block Device (xvd)
Disk /dev/xvdc: 125829120s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start      End          Size         File system  Name  Flags
1       2048s     20969471s   20967424s    ext4         opt
```

Partition Table specifies the partition style for existing disks. **msdos** indicates that the disk partition style is MBR, and **gpt** indicates that the disk partition style is GPT.

If the preceding information is displayed, enter **Fix** to rectify the disk exception. Then take note of the first and last sectors of the **/dev/xvdc1** partition. These values will be used during the partition recreation. In this example, the partition's first sector is **2048**, and its last sector is **20969471**.

The **/dev/xvdc1** partition number is **1**. Therefore, enter **rm 1** and press **Enter** to delete the partition.

Step 7 Enter **p** and press **Enter** to check whether the **/dev/xvdc1** partition has been deleted.

```
(parted) rm 1
(parted) p
Model: Xen Virtual Block Device (xvd)
Disk /dev/xvdc: 125829120s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number Start End      Size      File system Name  Flags
```

Step 8 Enter **mkpart opt 2048s 100%** and press **Enter** to create a partition.

opt indicates the name of the created partition, **2048s** indicates the first sector recorded in [Step 6](#), and **100%** indicates that all capacity of the disk is allocated to one partition, which must be greater than or equal to the last sector recorded in [Step 6](#).

Information similar to the following is displayed:

```
(parted) mkpart opt 2048s 100%
Warning: You requested a partition from 2048s to 125829199s (sectors 2048..125829199).
The closest location we can manage is 2048s to 125829036s (sectors 2048..125829036).
Is this still acceptable to you?
Yes/No? Yes
```

Enter **Yes** as prompted to set the last sector.

If the following warning message is displayed, enter **Cancel** to stop the partitioning. Then, find the first sector with the best disk performance and use that value to partition the disk. The warning message will not be displayed if the first sector with the best disk performance has been entered. In this example, **2048s** is one of such first sectors. Therefore, the system does not display the warning message.

```
Warning: The resulting partition is not properly aligned for best performance.
Ignore/Cancel? Cancel
```

NOTE

Data will be lost if the following operations are performed:

- Select a first sector which is inconsistent with that of the original partition.
- Select a last sector which is smaller than that of the original partition.

Step 9 Enter **p** and press **Enter** to check whether the **/dev/xvdc1** partition has been recreated.

```
(parted) p
Model: Xen Virtual Block Device (xvd)
Disk /dev/xvdb: 125829120s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number Start      End        Size       File system Name  Flags
 1      2048s      125829086s 125827039s ext4      opt
```

The **/dev/xvdc1** partition has been recreated.

Step 10 Enter **q** and press **Enter** to exit parted.

Step 11 Run the following command to view the disk partition information after the partition expansion:

lsblk

If information similar to the following is displayed, the total capacity of **/dev/xvdc** is 60 GB. The new 50 GB of the capacity is allocated to the **/dev/xvdc1** partition, and the partition is mounted on the **/mnt/sdc** directory. Skip [Step 12](#), [Step 14](#), and [Step 15](#).

```
[root@ecs-1120 sdc]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
xvda 202:0 0 80G 0 disk
├─xvda1 202:1 0 40G 0 part /
└─xvda2 202:2 0 40G 0 part /opt
xvdb 202:16 0 350G 0 disk
├─xvdb1 202:17 0 100G 0 part
└─xvdb2 202:18 0 200G 0 part
xvdc 202:32 0 60G 0 disk
└─xvdc1 202:33 0 60G 0 part /mnt/sdc
```

Step 12 Run the following command to check the correctness of the file system on **/dev/xvdc1**:

e2fsck -f /dev/xvdc1

Information similar to the following is displayed:

```
[root@ecs-1120 linux]# e2fsck -f /dev/xvdc1
e2fsck 1.42.9 (28-Dec-2013)
Pass 1: Checking inodes, blocks, and sizes
Pass 2: Checking directory structure
Pass 3: Checking directory connectivity
Pass 4: Checking reference counts
Pass 5: Checking group summary information
/dev/xvdc1: 11/655360 files (0.0% non-contiguous), 83137/2620928 blocks
```

Step 13 Run the following command to expand the size of the file system on **/dev/xvdc1**:

resize2fs /dev/xvdc1

Information similar to the following is displayed:

```
[root@ecs-1120 linux]# resize2fs /dev/xvdc1
resize2fs 1.42.9 (28-Dec-2013)
Resizing the filesystem on /dev/xvdc1 to 15728379 (4k) blocks.
The filesystem on /dev/xvdc1 is now 15728379 blocks long.
```

Step 14 Run the following command to view the disk partition information after the partition expansion:

lsblk

Information similar to the following is displayed:

```
[root@ecs-1120 linux]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
xvda 202:0 0 80G 0 disk
├─xvda1 202:1 0 40G 0 part /
└─xvda2 202:2 0 40G 0 part /opt
xvdb 202:16 0 350G 0 disk
├─xvdb1 202:17 0 100G 0 part
└─xvdb2 202:18 0 200G 0 part
xvdc 202:32 0 60G 0 disk
└─xvdc1 202:33 0 60G 0 part
```

In the command output, the total capacity of the **/dev/xvdc** disk is 60 GB, in which the additional 50 GB has been allocated to the **dev/xvdc1** partition.

Step 15 Run the following command to mount the created partition to the **/mnt/sdc** directory:

```
mount /dev/xvdc1 /mnt/sdc
```

Step 16 Run the following command to view the mount result of **/dev/xvdc1**:

```
df -TH
```

Information similar to the following is displayed:

```
[root@ecs-1120 linux]# mount /dev/xvdc1 /mnt/sdc
[root@ecs-1120 linux]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/xvda1      ext4      43G   8.3G   33G   21% /
devtmpfs        devtmpfs  885M    0   885M    0% /dev
tmpfs           tmpfs     894M    0   894M    0% /dev/shm
tmpfs           tmpfs     894M   18M   877M    2% /run
tmpfs           tmpfs     894M    0   894M    0% /sys/fs/cgroup
tmpfs           tmpfs     179M    0   179M    0% /run/user/2000
tmpfs           tmpfs     179M    0   179M    0% /run/user/0
tmpfs           tmpfs     179M    0   179M    0% /run/user/1001
/dev/xvda2      ext4      43G   51M   40G    1% /opt
/dev/xvdc1      ext4      64G   55M   60G    1% /mnt/sdc
```

The **/dev/xvdc1** partition has been mounted on the **/mnt/sdc** directory.

----End

4.7.9 Adding a Data Disk

You can expand disk capacity by adding data disks.

Procedure

- Step 1** Apply for a data disk as required by services. For details, see section [Applying for a Data Disk](#).
- Step 2** Attach the data disk to the BMS requiring capacity expansion. For details, see [Attaching an EVS Disk](#).
- Step 3** After the attachment is complete, initialize the data disk by referring to section [Initializing a Data Disk](#).

----End

4.8 Releasing an EVS Disk

4.8.1 Detaching a Data Disk

This section describes how to detach a data disk from a BMS.

You can detach a data disk in the following ways:

- BMS console: To detach multiple data disks from one BMS, perform related operations on the BMS console.
- EVS console: To detach a shared disk from multiple BMSs, perform related operations on the EVS console. In other scenarios, you can perform related operations on either console.

Prerequisites

If automatic attachment upon instance start has been configured in Linux by referring to [Setting Automatic Disk Attachment Upon Instance Start](#), comment out or delete the added content in the `/etc/fstab` file before detaching the disk. Otherwise, you may fail to access Linux after the disk is detached.

Restrictions

- Data disks can be detached online, that is, data disks can be detached from BMSs in the **running** state.
- Before detaching a disk online from an instance running Windows, log in to the instance to perform the offline operation and confirm that the disk is not being read and written. Otherwise, the disk will fail to be detached.
- Before detaching a disk online from an instance running Linux, log in to the instance, run the **umount** command to cancel the relationship between the disk and the file system, and confirm that the disk is not being read and written. Otherwise, the disk will fail to be detached.

Operations on the BMS Console


Step 1 Log in to ManageOne as a VDC administrator or VDC operator using a browser.

URL in non-B2B scenarios: <https://Domain name of ManageOne Operation Portal>, for example, <https://console.demo.com>.

URL in B2B scenarios: <https://Domain name of ManageOne Tenant Portal>, for example, <https://tenant.demo.com>.

You can log in using a password or a USB key.

- Login using a password: Enter the username and password.
The password is that of the VDC administrator or VDC operator.
- Login using a USB key: Insert a USB key with preset user certificates, select the required device and certificate, and enter a PIN.

Step 2 Click  in the upper left corner and select a region and resource set. Choose **Compute > Bare Metal Server**.

Step 3 In the BMS list, locate the BMS from which you want to detach the disk.

Step 4 Click the name of the queried BMS.

The page showing details about the BMS is displayed.

Step 5 On the **EVS** tab, select the data disk to be detached and click **Detach**.

Step 6 Click **OK**.

NOTICE

For details about how to detach an EVS disk, see **Operation Help Center > Compute > Bare Metal Server > User Guide > Data Disks > Attaching or Detaching EVS Disks**.

Step 7 After the detaching is complete, the disk will no longer be displayed on the **EVS** tab of the BMS.

----End

Operations on the EVS Console

Step 1 Select a region and resource set, and choose **Storage > Elastic Volume Service**. The EVS console is displayed.

Step 2 Click **Detach** in the row where the EVS disk to be detached is located.

If you are attempting to detach a shared disk, select the instance where the shared disk resides. You can select multiple instances.

Step 3 Click **OK**.

- If a non-shared disk is detached, **Status** of the EVS disk becomes **Available**.
- If a shared disk is detached, **Status** of the EVS disk becomes **Available** only after it is detached from all instances.

NOTICE

For details about how to detach an EVS disk, see **Operation Help Center > Compute > Bare Metal Server > User Guide > Data Disks > Attaching or Detaching EVS Disks**.

----End

Follow-up Procedure

If the data disk is no longer needed after the detachment, you can delete the data disk to release space. For details, see [Deleting Data Disks](#).

4.8.2 Deleting Data Disks

If a data disk is no longer required, you can soft delete or permanently delete it. A soft deleted data disk is stored in the recycle bin. A soft deleted data disk can be restored from the recycle bin. A permanently deleted data disk cannot be restored.

Restrictions

- A disk that has been attached to an instance cannot be deleted.
- If a disk has been configured with the disaster recovery service (CSDR, CSHA, or VHA), the disk cannot be deleted.
- If an EVS disk has a snapshot, the EVS disk can be soft deleted only when the snapshot is in the **Available** or **Error** state.
- When an EVS disk is permanently deleted, all snapshots of the EVS disk are also deleted.
- A shared disk to be deleted must have been detached from all instances.
- If the ECS to which an EVS disk belongs has not been created, the EVS disk cannot be deleted.

Procedure

Step 1 Log in to the EVS Console as a VDC administrator or VDC operator. For details, see [Logging In to the EVS Console as a VDC Administrator or VDC Operator](#).

Step 2 On the **Elastic Volume Service** page, locate the row that contains the target data disk, click **More** in the **Operation** column, and choose **Delete**.

To delete multiple data disks at a time, select ☐ in front of the data disks on the **Elastic Volume Service** page.

Step 3 Delete the data disk.

- If you want to delete an EVS disk permanently, select **By selecting this option, the disk will be permanently deleted instead of being deleted to the recycle bin.** and click **OK**.

When an EVS disk is permanently deleted, all snapshots of the EVS disk are also deleted and cannot be restored.

- If you want to delete a data disk but require that the deleted data disk can be restored later, click **OK** to soft delete the data disk.

Delete Disk



Are you sure you want to delete the following disk?

This operation will delete the disk to the recycle bin. You can choose 'Console > Recycle Bin' to restore or permanently delete the disk.

If a disk is permanently deleted, all the snapshots of the disk will be deleted. Exercise caution when you perform this operation.

Name	Capacity (GB)	Snapshots	Disk Type	Created	Status
disk-1	50	0	FC_volume_t...	Dec 25, 2019 15:29:48	Available

☐ By selecting this option, the disk will be permanently deleted instead of being deleted to the recycle bin.

OK

Cancel


Step 4 If deleting a data disk requires approval, contact the administrator for approval. Otherwise, skip this step.


- After a data disk is permanently deleted, it is no longer displayed in the EVS disk list.
- After a data disk is soft deleted, the status of the data disk changes to **Soft deleted**.

----End

Follow-up Procedure

A deleted EVS disk is stored in the recycle bin and still occupies storage space.

- To restore data on a disk that has been **Soft deleted**, click  in the upper left corner of the page and choose **Recycle Bin**. On the **Recycle Bin** page, select the target EVS disk and click **Restore**.

- To permanently delete an EVS disk that has been **Soft deleted**, click  in the upper left corner of the page and choose **Recycle Bin**. On the **Recycle Bin** page, select the target EVS disk and click **Delete Permanently**.

4.9 Managing Disk Types

4.9.1 Changing the Disk Type

When the read and write performance of the storage device where the upper-layer service resides no longer suits the service, you can change the disk type to alter the type of the storage device to change the read and write performance, meeting the requirements of varying service performance of the instance. For example: Thick LUNs have been configured. It is found later that the actually used capacity is small. Therefore, Thick LUNs need to be changed to thin LUNs. SAS disks have been purchased but later services require new high-performance SSDs. In this case, services need to be migrated from SAS disks to SSDs.

Restrictions

- Changing the disk type is supported when the backend storage is OceanStor V3/V5, OceanStor Dorado V3, or Huawei Distributed Block Storage.
- If the backend storage is OceanStor V3/V5 or OceanStor Dorado V3, the disk type can be changed between different storage pools in the same storage system. If the backend storage is Huawei Distributed Block Storage, the disk type can be changed only in the same storage pool.
- If the backend storage is Huawei Distributed Block Storage, the disk type can be changed only by modifying the QoS attribute.
- The administrator needs to import the SmartMigration license on the device in advance if the backend storage is OceanStor V3/V5 or OceanStor Dorado V3.
- When changing the disk type, if the backend storage is OceanStor Dorado 6.x, the administrator needs to check whether the SmartMigration license has been imported to the device in advance. (The basic software package of OceanStor Dorado 6.x contains the SmartMigration license.)
- You can change the type of the EVS disk only in the **Available** or **In-use** state.
- If a disk has snapshots or is configured with the backup service (VBS or CSBS) or the disaster recovery service (CSDR, CSHA, or VHA), the disk type cannot be changed.
- If backend storage of a disk is heterogeneous storage, the disk type cannot be changed.

NOTICE

If you change the disk type of an EVS disk that has been attached to an instance, the service performance of the instance will be affected.

Procedure

Step 1 Log in to EVS Console as a VDC administrator or VDC operator.

Step 2 On the **Elastic Volume Service** page, locate the row that contains the target EVS disk, click **More** in the **Operation** column, and choose **Change Disk Type**.

<input type="checkbox"/>	Name	Status	Share...	Capacity ...	Disk Type	Disk properties	Created	Expires	AZ	Attaching Inf...	Operation
▶	ecs-d6ee-volume-0000	In-use	No	1	business_type...	System disk	Dec 25, 2019 19:30:56	Never		ecs-d6ee ECS	Attach Detach More ▾
▶	ecs-f35f-volume-0001	In-use	Yes	60	business_type...	Data disk	Dec 25, 2019 19:11:06	Never		ecs-f35f ECS	Expand Capacity Apply for Disk
▶	ecs-f35f-volume-0000	In-use	No	1	business_type...	System disk	Dec 25, 2019 19:10:33	Never		ecs-f35f ECS	Apply for Snapshot Delete
▶	ecs-7f9e-volume-0001	Available	No	60	business_type...	Data disk	Dec 25, 2019 17:26:05	Never		--	Extend
▶	ecs-7f9e-volume-0000	In-use	No	1	business_type...	System disk	Dec 25, 2019 17:25:30	Never		ecs-7f9e ECS	Change Disk Type

Step 3 Select the target disk type, and click **Next**.

NOTE

When the backend storage device of the source disk type is Huawei Distributed Block Storage, at least one of the backend storage devices of the selected target disk type should be Huawei Distributed Block Storage. Otherwise, the disk type cannot be changed.

Step 4 Confirm the application information and click **Apply Now**.

Step 5 If changing the disk type of this EVS disk requires approval, contact the administrator for approval. If no, skip this step.

On the **Elastic Volume Service** page, view the disk type of the EVS disk. If the disk type of the EVS disk has changed as expected, changing the disk type is successful.

NOTE

If the EVS disk stays in **Retyping** state for a long time after the disk type is changed, contact the administrator to rectify the fault by referring to "Storage Services" > "Elastic Volume Service (EVS for BMS)" > "FAQs" > "What Can I Do If an EVS Disk Whose Disk Type Is Changed Remains in the Retyping State for a Long Time?" in *HUAWEI CLOUD Stack 8.1.0 Resource Provisioning Guide*.

----End

4.10 Managing Snapshots

4.10.1 Applying for a Snapshot

A snapshot can capture the data and status of a disk at a certain time point. If a service change or application software upgrade is required, you can create a snapshot for the disk in advance. If a fault occurs during the change or upgrade, you can use the snapshot to quickly restore disk data, ensuring service continuity and security. You can also use snapshots for routine backup of disk data.

Restrictions

- If backend storage is one of OceanStor V3/V5 or OceanStor Dorado V3 series, it is necessary for the administrator to import the license for HyperSnap in advance on the device side.

- Snapshots can be created only for disks in the **Available** or **In-use** state.
- A snapshot name cannot be the same as the prefix of the temporary snapshot created by the backup service, such as Volume Backup Service (VBS) and Cloud Server Backup Service (CSBS), or the disaster recovery service, such as Cloud Server Disaster Recovery (CSDR), Cloud Server High Availability (CSHA), and VHA.
- Snapshots created using the EVS console consume the capacity quota instead of quantity quota of EVS disks.
- Snapshots created using the EVS console and temporary snapshots created by the DR and backup service (VBS, CSBS, CSDR, CSHA, or VHA) consume backend storage capacity. If a large number of snapshots are created, contact the administrator to set the thin provisioning ratio of backend storage to a large value, preventing EVS disk provisioning failures caused by excessive snapshots.
- Temporary snapshots created by the backup service (VBS or CSBS) or the disaster recovery service (CSDR, CSHA, or VHA) do not consume EVS disk quotas.
- No snapshots can be created for disks that have expired.
- No snapshots can be created for disks that have been soft deleted.
- If a task for creating a snapshot fails, the task is automatically deleted.
- If backend storage of the disk is heterogeneous storage, snapshots can be created.

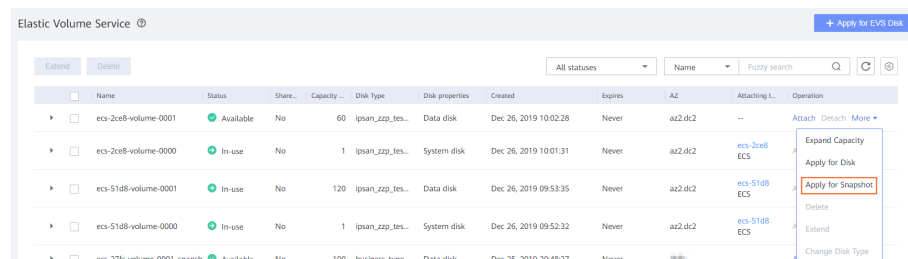
Procedure

Step 1 Log in to the EVS Console as a VDC administrator or VDC operator. For details, see [Logging In to the EVS Console as a VDC Administrator or VDC Operator](#).

Step 2 Use one of the following methods to display the **Apply for Snapshot** page.

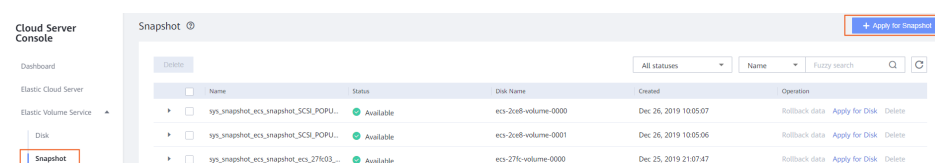
- Method 1

In the EVS disk list, locate the row that contains the target disk, click **More** in the **Operation** column, and choose **Apply for Snapshot**.



- Method 2

In the navigation pane on the left, choose **Elastic Volume Service > Snapshot**. Then, click **Apply for Snapshot**.



Step 3 Configure the basic information about the snapshot, as shown in [Table 4-15](#).

Table 4-15 Parameters for creating a snapshot


Parameter	Description	Example Value
Name	The value contains only digits, letters, underscores (_), and hyphens (-) and cannot exceed 63 characters. NOTE The snapshot name cannot be the same as the name prefix of any temporary snapshot created by the backup service (VBS or CSBS) or the DR service (CSDR, CSHA, or VHA) or the name prefix of any VM snapshot. The prefixes include autobk_snapshot , manualbk_snapshot , and sys_snapshot .	snapshot-01
EVS Disk	<ul style="list-style-type: none">When you create a snapshot for the specified disk, the disk has been determined.When you create a snapshot on the snapshot page, click Select, select the target disk, and click OK.	volume-01
Description	Description of the created snapshot. The value contains a maximum of 63 bytes.	-

Step 4 Click **Next**.

Step 5 If you do not need to modify the specifications, click **Apply Now** to start the snapshot application.

If you need to modify the specifications, click **Back** to modify parameter settings.

Step 6 If creating a snapshot requires approval, contact the administrator for approval. If no, skip this step.

Return to the **Snapshot** page and click  to view the snapshot creation information.

If the status of the snapshot is **Available**, the snapshot is successfully created.

-----End

4.10.2 Rolling Back an EVS Disk Using a Snapshot

If the data of an EVS disk is lost or damaged, or if you want to roll back the disk to a certain time point and you have created a snapshot for the disk, you can use the snapshot to roll back the disk to the time when the snapshot was created.

Restrictions

- A temporary snapshot created by the backup service (VBS or CSBS) or the disaster recovery service (CSDR, CSHA, or VHA) cannot be used to roll back the EVS disk.

- Snapshots created for disks having any DR service (CSDR/CSHA/VHA) configured cannot be rolled back.
- A snapshot can be used to roll back its source EVS disk, and cannot be used to roll back any other EVS disk.
- When the source disk of a snapshot is in the recycle bin, EVS disk rollback from the snapshot is not supported.
- If backend storage of the disk is heterogeneous storage (excluding XSKY), EVS disk rollback from a snapshot is not supported.
- The rollback operation is irreversible. After the rollback is complete, all data generated from the time when the snapshot was created to the time when the rollback operation is performed will be lost and cannot be restored. Exercise caution when performing this operation.

Procedure

- Step 1** Log in to the EVS Console as a VDC administrator or VDC operator. For details, see [Logging In to the EVS Console as a VDC Administrator or VDC Operator](#).
- Step 2** In the navigation pane on the left, choose **Elastic Volume Service > Snapshot**.
The **Snapshot** page is displayed.
- Step 3** In the snapshot list, locate the row that contains the target snapshot and click **Rollback data** in the **Operation** column.
- Step 4** After confirming that the rollback information is correct, click **OK** to perform the rollback.

Return to the **Snapshot** page. After the snapshot status changes from **Rolling back** to **Available**, the data rollback is successful.

----End

4.10.3 Deleting Snapshots

Snapshots occupy the disk capacity quota. Therefore, if some snapshots are not required, you are advised to delete them in a timely manner to release space.

Restrictions

- Users are allowed to delete a temporary snapshot created by the backup service (VBS or CSBS). After the snapshot is deleted, if users want to back up the EVS disk corresponding to the snapshot, full backup is performed for the first time.
- Temporary snapshots created by the disaster recovery service (CSDR, CSHA, or VHA) cannot be deleted.
- You can delete a snapshot only when its state is **Available**, **Deletion failed**, or **Error**.

NOTICE

The deletion operation is irreversible. Exercise caution when performing this operation.

Procedure

- Step 1** Log in to the EVS Console as a VDC administrator or VDC operator. For details, see [Logging In to the EVS Console as a VDC Administrator or VDC Operator](#).
- Step 2** In the navigation pane on the left, choose **Elastic Volume Service > Snapshot**.
- Step 3** Delete snapshots.
- Deleting one snapshot
 - a. On the **Snapshot** page, locate the row that contains the target snapshot and click **Delete** in the **Operation** column.
 - b. In the displayed dialog box, confirm the information and click **OK**.
 - Deleting multiple snapshots at a time
 - a. On the **Snapshot** page, select multiple snapshots by selecting ☐ next to them.
 - b. Click **Delete** in the upper part of the page. In the displayed dialog box, click **OK**.
- Step 4** If deleting a snapshot requires approval, contact the administrator for approval. If no, skip this step.

----End

4.11 Managing Backups

4.11.1 Creating a Backup

To prevent data loss caused by misoperations or system faults, you can create a backup to copy all or some data sets from the hard disk of the application host to the backup storage. This also ensures data correctness and security.

You can create backups for EVS disks in the following ways:

- EVS console: To create backups for a single EVS disk, perform related operations on the EVS console.
- VBS console: To create backups for multiple EVS disks in batches, perform related operations on the VBS console.

Restrictions

- Only disks in the **Available** or **In-use** state can be backed up.
- If backend storage of the disk is heterogeneous storage, backups cannot be created.

Prerequisites

The VBS service has been configured and enabled by the administrator. Otherwise, this function is unavailable. For details about how to enable the VBS, see "Storage Services" > "Elastic Volume Service (EVS for BMS)" > "FAQs" > "How Do I Enable the VBS?" in *HUAWEI CLOUD Stack 8.1.0 Resource Provisioning Guide*.

EVS Console

Step 1 [Log in to EVS Console as a VDC administrator or VDC operator.](#)

Step 2 On the **Elastic Volume Service** page, locate the row that contains the target EVS disk, click **More** in the **Operation** column, and then choose **Create Backup**.

Step 3 On the **Create VBS Backup** page, configure parameters as prompted.

For details, see **Operation Help Center > DR & Backup > User Guide > Creating a Periodic Backup Task > Creating a Backup Task**.

----End

VBS Console

For details about how to create a backup task on the VBS page, see **Operation Help Center > DR & Backup > User Guide > Creating a Periodic Backup Task > Creating a Backup Task**.

Follow-up Procedure

After an EVS disk is backed up successfully, a VBS backup of the disk is generated. You can view the details about the backup on **Operation Help Center > DR & Backup > User Guide > Managing Backups and Replicas > Viewing EVS Disk Backups and Replicas**. You can also view backup tree information on the **Elastic Volume Service** page.

4.11.2 Using a Backup to Restore Data

If the data on the current EVS disk is lost or damaged and you have created a backup for the disk, you can use the backup to restore the data on the disk to the time when the backup was created.

- You can use a backup to restore data to the original EVS disk. For details, see **Operation Help Center > DR & Backup > User Guide > Restoring a Disk Using a Backup > Restoring Backup Data to the Source EVS Disk**.
- If the current EVS disk is damaged, you can use the backup to restore the data to another EVS disk. For details, see **Operation Help Center > DR & Backup > User Guide > Restoring a Disk Using Backup > Restoring Backup Data to Another EVS Disk**.

4.12 Managing EVS Disks

4.12.1 Extending the EVS Disk Validity Period


This section describes how to extend the validity period of an EVS disk. You can extend the validity period of multiple EVS disks in batches. When an EVS disk expires, you can only delete the disk or extending the disk validity period. If your EVS disk is about to expire, you are advised to extend the validity period as soon as possible.

Restrictions

- If an EVS disk is created with an instance, the validity period of the EVS disk is unlimited.
- If the validity period of an EVS disk is unlimited, the validity period cannot be extended.
- You can extend the EVS disk validity period only when the disk is in the **available** or **In-use** status.
- If an EVS disk has expired, its snapshot cannot be used to roll back the EVS disk or create an EVS disk. To continue using this EVS disk, extend its validity period.
- When an EVS disk expires, its data will not be deleted. You can continue using this EVS disk after extending its validity period.

Procedure

Step 1 Log in to the EVS Console as a VDC administrator or VDC operator. For details, see [Logging In to the EVS Console as a VDC Administrator or VDC Operator](#).

Step 2 On the **Elastic Volume Service** page, click  to select the target EVS disk.

You can also choose **More > Extend** to extend the validity period of a single EVS disk.

Step 3 Click **Extend**.

Step 4 In the **Configuration** area, select **Extension Period** and enter the description of the EVS disk. Click **Next**.

The extension period can be set to **Unlimited**, **1 year**, or **Custom**.

Step 5 On the **Resource Details** page, confirm the extension information.

- After confirming that the information is correct, click **Apply Now** to start extending the EVS disk validity period.
- If you need to modify the specifications, click **Back** to modify parameter settings.

Step 6 If this application requires approval, contact the administrator for approval. If no, skip this step.

On the **Elastic Volume Service** page, view the expiration time of the EVS disk. If the expiration time of the EVS disk changes as expected, validity period extending has succeeded.

----End

4.12.2 Viewing Monitoring Data

If an EVS disk has been attached to an instance which is not of the VMware virtualization type and in use, you can refer to this section to view the monitoring data of the EVS disk, such as the read and write rates.

Procedure

- Step 1** Log in to the EVS Console as a VDC administrator or VDC operator. For details, see [Logging In to the EVS Console as a VDC Administrator or VDC Operator](#).
 - Step 2** Find the EVS disk whose monitoring data you want to view and click ► next to the disk. The details about the disk will be displayed.
 - Step 3** Click **View Monitoring Data** to view the read and write rates of the EVS disk.
- End

4.12.3 Changing the VDC Quota

When a ManageOne operation administrator creates a tenant and configures the VDC service, the operation administrator needs to set not only the resource pool for each region but also the resource quota for the tenant and the resource quota for the VDC of this level. The EVS disk quota consists of the total capacity and total quantity for a tenant and the total capacity and total quantity for the VDC of this level, which you can configure. When the EVS disk quota in the VDC is insufficient, contact the administrator to change the quota. For details, see .

4.13 FAQs

4.13.1 Can I Attach an EVS Disk to Multiple Instances?

A non-shared EVS disk can be attached to only one instance.

A shared EVS disk can be attached to a maximum of 16 instances by default.

4.13.2 How Can I Attach an EVS Disk to a BMS?

For details, see section [Attaching an EVS Disk](#).

4.13.3 How Many States Does an EVS Disk Have?

Disks are managed using the EVS during the entire process from disk creation to release, ensuring optimal user experience of applications or sites hosted on them. [Figure 4-24](#) shows the switching between different states of an EVS disk, and [Table 4-16](#) describes the meaning and supported operations of each state.

Figure 4-24 EVS disk status

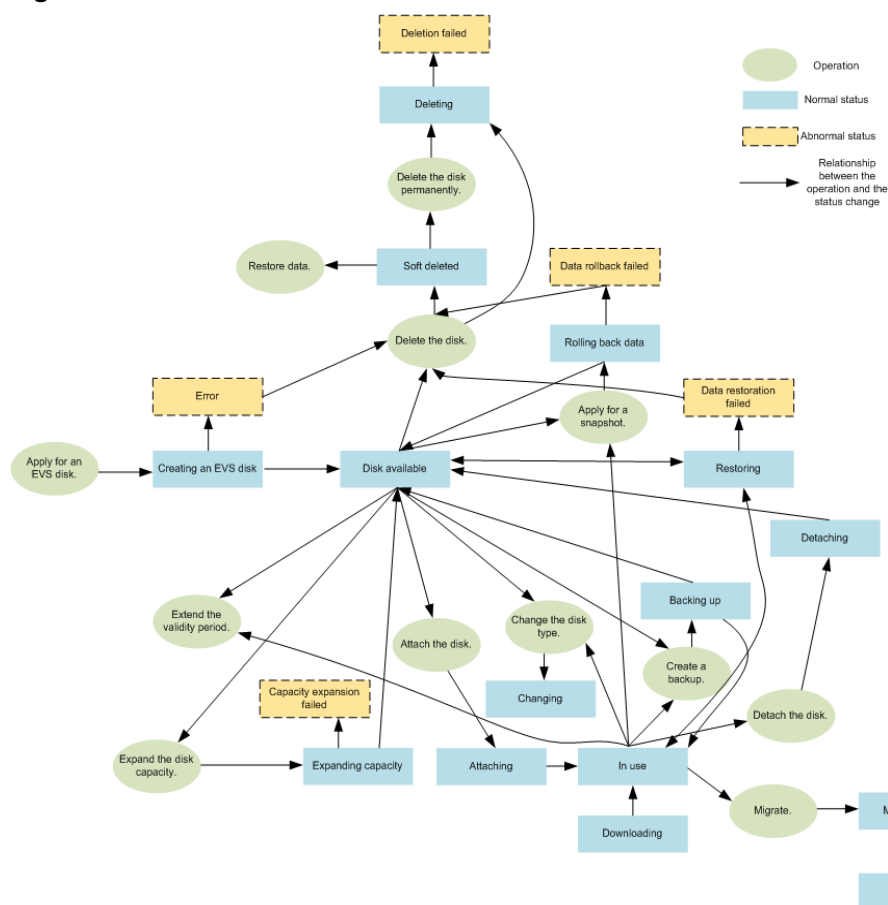


Table 4-16 EVS disk status description

EVS Disk Status	Status Description	Supported Operation
In-use	The EVS disk is attached to an instance and is in use.	<ul style="list-style-type: none"> • Detach • Expand capacity • Create snapshot • Create backup • Extend validity period of EVS disks (except EVS disks whose expiration time is Unlimited) • Change disk type
Downloading	Data is being downloaded from an image to the EVS disk. An EVS disk is in this status when you are creating an ECS.	-

EVS Disk Status	Status Description	Supported Operation
Available	The EVS disk is created and has not been attached to any server.	<ul style="list-style-type: none">• Attach• Expand capacity• Create snapshot• Create backup• Delete• Extend validity period of EVS disks (except EVS disks whose expiration time is Unlimited)• Change disk type
Creating	The EVS disk is being created.	-
Attaching	The EVS disk is being attached to an instance.	-
Detaching	The EVS disk is being detached from an instance.	-
Deleting	The EVS disk is being deleted.	-
Extending	The capacity of the EVS disk is being expanded.	-
Rolling back	The EVS disk data is being rolled back using the snapshot.	-
Backing up	The EVS disk is being backed up.	-
Restoring	The EVS disk data is being restored using the backup.	-
Retyping	The EVS disk type is being changed.	-
Soft deleted	The EVS disk has been soft deleted and moved to the recycle bin.	<ul style="list-style-type: none">• Restore• Permanently delete
Error	An error occurs when you are creating an EVS disk. You can delete the EVS disk and create it again.	Delete
Deletion failed	An error occurs when you are deleting an EVS disk. Contact the administrator.	None

EVS Disk Status	Status Description	Supported Operation
Expansion failed	An error occurs when you are expanding the capacity of an EVS disk. The administrator will contact you and help you handle this error. Do not perform any operations on the disk before the administrator contact you. If you require that the error be handled as soon as possible, contact the administrator.	-
Rollback failed	An error occurs during an EVS disk rollback.	Delete
Restoration failed	An error occurs during EVS disk restoration. The administrator will contact you and help you handle this error. Do not perform any operations on the disk before the administrator contact you. If you require that the error be handled as soon as possible, contact the administrator.	Delete
Maintenance	The EVS disk is being migrated.	-
Reserved	Status of the source disk after the migration task is complete	-

4.13.4 Does an EVS Disk or Snapshot Generate Metering Information in All States?

EVS disks and snapshots do not generate metering information in the following states:

- EVS disk state: Creating, Error, Restoration failed, Downloading, Expansion failed, Deleting, Deletion failed, Rollbacking, Rollback failed, or Retyping
- Snapshot state: Creating, Error, Deleting, or Deletion failed

4.13.5 Can I Change the EVS Disk Capacity?

EVS disk capacity can be expanded but cannot not be reduced at present.

For details about how to expand the capacity of an EVS disk, see [Expanding EVS Disk Capacity](#).

4.13.6 Will Data in an EVS Disk Be Lost When the EVS Disk Is Detached?

The data will not be lost.

To prevent data loss, you are advised to perform the following operations:

- Windows operating system: Before the detachment, log in to the instance and perform the offline operation.
- Linux operating system: Before the detachment, log in to the instance and run the **umount** command to unmount the disk partition.

4.13.7 How Can I Test the Performance of a Windows EVS Disk?

This section describes how to test the performance of an EVS disk on an instance that runs on the Windows OS.

Prerequisites

Before testing the performance, log in to the Iometer official website to download and install the Iometer performance test tool.

Precautions

- In the EVS disk performance test, if the start cylinder number is not 4-KB aligned, the EVS disk performance will be greatly affected. Ensure that the start cylinder number is 4-KB aligned before you start the test.
- Do not perform the fio test on system disks to avoid damaging important system files.
- You are advised to perform the fio test on idle disks that do not store important data and to create a file system after the test is complete.
- Do not perform the test on data disks. Otherwise, disk data may be damaged if the metadata of the underlying file system is damaged.

Procedure

The method for testing an EVS disk varies depending on the in-use server OS. This section uses Windows 7 Professional 64-bit as the example OS to describe how to test the EVS disk performance.

Step 1 Log in to an instance.

Step 2 Press **win+R** to open the **Run** window. Enter **msinfo32** and click **OK**.

The system information window is displayed.

Step 3 Choose **Component > Storage > Disk**. View the start offset of the partition in the right pane, and check whether the start cylinder number of the disk partition is 4-KB aligned.

- If 4096 can be divided by the start cylinder number, the partition is 4-KB aligned. Go to [Step 4](#) to start the test.

- If 4096 cannot be divided by the start cylinder number, the partition is not 4-KB aligned. If you need to continue to perform the test, ensure the 4-KB alignment for the partition.

NOTICE

Data losses will occur when you delete the partition and select another start cylinder number for the 4-KB alignment. Exercise caution when performing such an operation.

Step 4 Use Iometer to test the EVS disk performance. For details, see the Iometer product document.

When the disk IOPS and throughput are tested, the parameter configurations for Iometer and fio are the same. For details, see [Table 4-17](#).

----End

4.13.8 How Can I Test the Performance of a Linux EVS Disk?

This section describes how to test the performance of an EVS disk on an instance that runs on the Linux OS.

Prerequisites

fio is a tool for testing disk performance. It is used to perform pressure tests on hardware. You are advised to use the I/O engine of Libaio to perform the test. You can log in to the official website to download and install fio and Libaio.

Precautions

- When you use an old version Linux OS, for example CentOS 6.5, and run **fdisk** to create partitions, the default start cylinder number is not 4-KB aligned, which will greatly affect the test performance. For that reason, if such an OS is used, you are advised to select a new start cylinder number that is 4-KB aligned when creating partitions.
- Do not perform the fio test on system disks to avoid damaging important system files.
- You are advised to perform the fio test on idle disks that do not store important data and to create a file system after the test is complete.
- Do not perform the test on data disks. Otherwise, disk data may be damaged if the metadata of the underlying file system is damaged.
- When testing disk performance, you are advised to directly test raw disks. When testing the performance of the file system, you are advised to specify a specific file for the test.

Procedure

The method for testing an EVS disk varies depending on the in-use server OS. This section uses CentOS 7.2 64-bit as the example OS to describe how to test the EVS disk performance.

Step 1 Log in to an instance and switch to the **root** user.

Step 2 Before you start the test, run the following command to check whether the start cylinder number is 4-KB aligned:

fdisk -lu

Information similar to the following is displayed:

```
[root@ecs-centos sdc]# fdisk -lu
```

```
Disk /dev/xvda: 10.7 GB, 10737418240 bytes, 20971520 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x7db77aa5
```

Device	Boot	Start	End	Blocks	Id	System
/dev/xvda1	*	2048	20968919	10483436	83	Linux

```
Disk /dev/xvdb: 10.7 GB, 10737418240 bytes, 20971520 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
```

```
Disk /dev/xvdc: 53.7 GB, 53687091200 bytes, 104857600 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x3cf3265c
```

Device	Boot	Start	End	Blocks	Id	System
/dev/xvdc1		2048	41943039	20970496	83	Linux

- If 8 can be divided by the start cylinder number, the partition is 4-KB aligned. Go to [Step 3](#) to start the test.
- If 8 cannot be divided by the start cylinder number, the partition is not 4-KB aligned. If you need to continue to perform the test, select another start cylinder number to ensure the 4-KB alignment for the partition.

NOTICE

Data losses will occur when you delete the partition and select another start cylinder number for the 4-KB alignment. Exercise caution when performing such an operation.

Step 3 Run commands and use fio to test the EVS disk performance:

- To test random write IOPS, run **fio -direct=1 -iodepth=128 -rw=randwrite -ioengine=libaio -bs=4k -size=10G -numjobs=1 -runtime=600 -group_reporting -filename=/dev/[device] -name=Rand_Write_IOPS_Test**.
- To test random read IOPS, run **fio -direct=1 -iodepth=128 -rw=randread -ioengine=libaio -bs=4k -size=10G -numjobs=1 -runtime=600 -group_reporting -filename=/dev/[device] -name=Rand_Read_IOPS_Test**.
- To test write throughput, run **fio -direct=1 -iodepth=32 -rw=write -ioengine=libaio -bs=1024k -size=10G -numjobs=1 -runtime=600 -group_reporting -filename=/dev/[device] -name=Write_BandWidth_Test**.

- To test read throughput, run **fio -direct=1 -iodepth=32 -rw=read -ioengine=libaio -bs=1024k -size=10G -numjobs=1 -runtime=600 -group_reporting -filename=/dev/[device] -name=Read_BandWidth_Test**.

[Table 4-17](#) describes the parameters.

Table 4-17 Parameters for testing the disk performance

Parameter	Description
direct	Indicates whether to use the direct I/O. The options are as follows: <ul style="list-style-type: none">If the value is 0, it indicates that the buffered I/O is used.If the value is 1, it indicates that the direct I/O is used.
iodepth	Defines the depth of the I/O queue during the test. The default value is 1. This queue depth refers to the queue depth of each thread. That said, when multiple threads are tested, the parameter defines the queue depth of each thread. Total number of concurrent I/Os of fio = iodepth x numjobs
rw	Defines the test read/write policy. The values are as follows: <ul style="list-style-type: none">Random read: randreadRandom write: randwriteSequential read: readSequential write: writeMixed random read/write: randrw
ioengine	Defines how the fio delivers I/O requests. There are synchronous I/O requests and asynchronous I/O requests. <ul style="list-style-type: none">Only one synchronous I/O request can be delivered at a time. The request is returned only after the kernel processing completes. In this way, the I/O queue depth for a single thread is always less than 1. However, multi-thread concurrent processing can be used. Generally, 16 to 32 threads are used to work at the same time to stuff the I/O queue depth.Usually, a batch of asynchronous I/O requests is submitted at a time by using the libaio method. After a batch of requests is processed, the number of interaction times is reduced, which is more efficient.
bs	Defines the I/O block size. The unit is k, K, m, or M. The default size is 4 KB.

Parameter	Description
size	Defines the amount of data processed by the test I/Os. If parameters, such as runtime , are not specified, the test ends until fio has processed all the specified data amount. The value can be a number with a unit or percentage. A number with a unit indicates the read/write data amount, for example size=10G , indicating a 10-GB read/write data amount. A percentage indicates the ratio of read/write data amount to the capacity of total files, for example size=20% , indicating the read/write data amount takes 20% of the total file space.
numjobs	Defines the number of concurrent threads.
runtime	Defines the test time. If this parameter is not specified, the test ends until the specified data amount is processed by the block size defined using parameter size .
group_reporting	Defines the test result display mode. The parameter value displays the statistics on a single thread instead of that on all jobs.
filename	Defines the name of the test file or device.
name	Defines the test task name.

----End

4.13.9 Logging In to the EVS Console as a VDC Administrator or VDC Operator

Procedure

Step 1 Log in to ManageOne as a VDC administrator or VDC operator using a browser.

URL in non-B2B scenarios: **https://Domain name of ManageOne Operation Portal**, for example, **https://console.demo.com**.

URL in B2B scenarios: **https://Domain name of ManageOne Tenant Portal**, for example, **https://tenant.demo.com**.

You can log in using a password or a USB key.

- Login using a password: Enter the username and password.
The password is that of the VDC administrator or VDC operator.
- Login using a USB key: Insert a USB key with preset user certificates, select the required device and certificate, and enter a PIN.

Step 2 Select a region and resource set, and choose **Storage > Elastic Volume Service**.
The EVS console is displayed.

----End