

Elastic Volume Service

User Guide

Date 2023-06-12

Contents

1 Overview	1
1.1 What Is EVS?	1
1.2 Disk Types and Performance	4
1.3 Device Types and Usage Instructions	8
1.4 Shared EVS Disks and Usage Instructions	9
1.5 EVS Encryption	13
1.6 EVS Backup	18
1.7 EVS Snapshot	19
1.8 Differences Between EVS Backups and EVS Snapshots	21
1.9 EVS Three-Copy Redundancy	22
1.10 Permissions	25
1.11 Constraints	27
1.12 EVS and Other Services	30
1.13 Basic Concepts	32
1.13.1 EVS Concepts	32
1.13.2 Region and AZ	32
2 Getting Started	34
2.1 Process Overview	34
2.2 Create an EVS Disk	36
2.3 Attach an EVS Disk	42
2.3.1 Attaching a Non-Shared Disk	42
2.3.2 Attaching a Shared Disk	45
2.4 Initialize an EVS Data Disk	47
2.4.1 Introduction to Data Disk Initialization Scenarios and Partition Styles	47
2.4.2 Initializing a Windows Data Disk (Windows Server 2008)	49
2.4.3 Initializing a Windows Data Disk (Windows Server 2019)	56
2.4.4 Initializing a Linux Data Disk (fdisk)	64
2.4.5 Initializing a Linux Data Disk (parted)	70
2.4.6 Initializing a Windows Data Disk Larger Than 2 TiB (Windows Server 2008)	76
2.4.7 Initializing a Windows Data Disk Larger Than 2 TiB (Windows Server 2012)	84
2.4.8 Initializing a Linux Data Disk Larger Than 2 TiB (parted)	92
3 Permissions Management	99

3.1 Creating a User and Granting EVS Permissions	
3.2 EVS Custom Policies	
4 Creating and Using an EVS Disk	102
5 Disk Capacity Expansion	105
5.1 Expansion Overview	105
5.2 Expanding Capacity for an In-use EVS Disk	108
5.3 Expanding Capacity for an Available EVS Disk	
5.4 Extending Disk Partitions and File Systems (Windows Server 2008)	
5.5 Extending Disk Partitions and File Systems (Windows Server 2016)	
5.6 Extending Disk Partitions and File Systems (Linux)	
5.6.1 Partition and File System Extension Preparations (Linux)	
5.6.3 Extending Partitions and File Systems for Data Disks (Linux)	
5.6.4 Extending Partitions and File Systems for SCSI Disks (Linux)	
6 Detaching an EVS Disk	
6.1 Detaching a System Disk	
6.2 Detaching a Data Disk	
7 Attaching an Existing Disk	
7.1 Attaching an Existing System Disk	
7.2 Attaching an Existing Non-Shared Disk	
7.3 Attaching an Existing Shared Disk	
8 Deleting EVS Disks	
9 Viewing EVS Disk Details	222
10 Managing Encrypted EVS Disks	225
11 Managing Shared EVS Disks	228
12 Managing EVS Backups	230
13 Managing EVS Snapshots	232
13.1 Snapshot Overview	232
13.2 Creating a Snapshot	233
13.3 Deleting a Snapshot	236
13.4 Viewing Details of a Snapshot	
13.5 Rolling Back Data from a Snapshot	
13.6 Creating an EVS Disk from a Snapshot	238
14 Managing EVS Transfers	240
15 Managing a Tag	
15.1 Tag Overview	
15.2 Adding a Tag	
15.3 Modifying a Tag	244

15.4 Deleting a Tag	
15.5 Searching Disks by Tags	
16 Changing EVS Disk Name	247
17 Viewing EVS Monitoring Data	249
18 Viewing EVS Monitoring Data (Agent Installed and Simplified Monitoring Metrics Used)	251
19 Auditing	255
20 Managing Quotas	256
20.1 Querying EVS Resource Quotas	
20.2 Increasing EVS Resource Quotas	
21 FAQ	
21.1 Summary	
21.2 General	
21.2.1 How Do I Start Using a Newly Created Disk?	
21.2.2 Can EVS Disks Be Used Directly for Storage?	
21.2.3 Can EVS Disks Be Used Alone?	
21.2.4 How Can I View My Disk Details?	260
21.2.5 Can I Change the AZ of My Disk?	
21.2.6 Can I Change the Disk Type, Device Type, or Sharing Attribute of My Disk?	
21.2.7 What Should I Do If an Error Occurs on My EVS Disk?	262
21.2.8 How Can I Obtain ECS NIC Information?	263
21.2.9 Why Do Some of My EVS Disks Not Have WWN Information?	263
21.2.10 How Can I Migrate Data from an EVS Disk?	264
21.2.11 What Are the Differences Between System Disks and Data Disks?	264
21.2.12 How Can I Download My EVS Disk Data to a Local PC?	265
21.2.13 How Can I Export the Original Data After I Changed My Server OS from Windows to CentC	S?.265
21.2.14 What Are the Differences Between MBR and GPT Partition Styles?	265
21.2.15 What Does the "reserveVolume" Trace Mean in CTS?	266
21.2.16 How Can I Upload Files to My EVS Disk?	
21.3 Capacity Expansion	
21.3.1 Can I Reduce or Temporarily Expand the Disk Capacity?	
21.3.2 What Are the Differences Between Expanding Capacity by Expanding an EVS Disk and Creati New EVS Disk?	
21.3.3 Will My Disk Data Be Lost After I Expand the Disk Capacity?	267
21.3.4 Can I Use Backups or Snapshots Created Before Capacity Expansion to Restore Data on Expa Disks?	
21.3.5 Do I Need to Restart the Server After Expanding the Disk Capacity?	268
21.3.6 Do I Need to Detach an EVS Disk Before Expanding Its Capacity?	
21.3.7 What Should I Do If My Disk Capacity Exceeds 2 TiB After Expansion?	
21.3.8 How Can I Allocate Newly Added Space to a New Partition?	
21.3.9 How Can I Allocate Newly Added Space to an Existing Partition?	270

21.3.10 Why Did My Disk Capacity Remain Unchanged on the Server After Capacity Expansion?	270
21.3.11 What Can't I Expand Capacity for My Disk?	270
21.3.12 How Do I Extend the File System of an Unpartitioned Data Disk in Linux?	271
21.3.13 How Do I Extend the Root Partition of a Quickly Provisioned BMS?	273
21.4 Attachment	276
21.4.1 Why Can't I View the Attached Data Disk on the Server?	276
21.4.2 Why Can't I Attach My Disk to a Server?	279
21.4.3 Can I Attach a Disk to Multiple Servers?	281
21.4.4 Can I Attach a Disk to a Server in Another AZ?	281
21.4.5 How Can I Add a Data Disk to an Existing Server?	282
21.4.6 Can I Attach Different Types of Disks to the Same Server?	282
21.4.7 What Should I Do If a Linux EVS Disk Is Attached to a Windows Server?	282
21.4.8 Can I Change the Function of a System Disk or Data Disk Created Along with a Server?	282
21.4.9 How Do I Obtain My Disk Name in the ECS OS Using the Device Identifier Provided on the Console?	283
21.5 Detachment	
21.5.1 If I Detach a Disk, Will I Lose the Data on My Disk?	
21.5.2 Why Can't I Detach My Disk?	
21.6 Deletion	
21.6.1 How Can I Recover Data from a Disk That Was Accidentally Deleted?	
21.7 Capacity	
21.7.1 What Is the Maximum Capacity Supported for the System and Data Disks?	
21.7.2 What Should I Do If My Disk Starts to Run Out of Space?	
21.7.3 What Can I Do If the Capacity of My Disk Reaches the Maximum But I Still Need More Spa	
21.7.4 What Should I Do If I Use fdisk to Initialize a Disk Larger Than 2 TiB and Then the Space in	
of 2 TiB Cannot Be Displayed?	
21.7.5 How Can I View My Disk Usage?	289
21.7.6 How Can I Monitor My Disk Usages?	296
21.7.7 Can I Transfer the Data Disk Capacity to a System Disk?	300
21.8 Performance	301
21.8.1 How Can I Test My Disk Performance?	301
21.8.2 Why Does My Disk Performance Test Using Fio Have Incorrect Results?	308
21.8.3 How Can I Handle a Slowdown in Disk Read/Write Speed and Increased I/O?	310
21.8.4 Why My Disk's Read IOPS Can't Reach the Theoretical Maximum IOPS When the Disk I/O L Almost 100%?	
21.9 Sharing	312
21.9.1 Do I Have to Deploy a Cluster to Use Shared Disks?	
21.9.2 How Many Servers Can I Attach a Shared Disk to?	
21.9.3 How Can I Attach a Shared Disk to Multiple Servers?	
21.9.4 Can a Shared Disk Be Attached to Servers That Belong to Different Accounts?	
21.9.5 Can I Attach a Shared Disk to Servers Running Different OSs?	312
21.10 Snapshot	
21.10.1 What Are the Typical Causes of a Snapshot Creation Failure?	

2023-06-12 v

21.10.2 Does EVS Support Automatic Snapshot Creation?		
21.10.4 How Is a Snapshot Created for My Disk? 21.10.5 Why Can't I Roll Back My Disk Data from a Snapshot? 21.10.6 Can I Roll Back the Data from a Snapshot After Reinstalling the OS or Formatting the Disk? 21.10.7 How Is the Snapshot Size Calculated? 21.10.8 Do Snapshots Take Space on the Disk? 21.10.9 Can I Perform Multiple Rollback Operations for a Snapshot? 21.10.10 Can I Replicate Snapshots to Other Regions or Accounts? 314 21.10.11 Why Can't I Find My Snapshot? 314 21.10.12 Can I Use a Snapshot to Create a Disk and Roll Back the Snapshot at Almost the Same Time? 314 21.11 Backup. 315 21.11.1 Do I Need to Stop the Server Before Performing a Backup? 315 21.11.2 Can I Back Up and Restore My EVS Disk to a Different Region? 316 317 A Appendix. 317 A.1 EVS Disk Status. 318	21.10.2 Does EVS Support Automatic Snapshot Creation?	313
21.10.5 Why Can't Roll Back My Disk Data from a Snapshot?	21.10.3 Can I Create Snapshots for Multiple Disks at a Time?	313
21.10.6 Can I Roll Back the Data from a Snapshot After Reinstalling the OS or Formatting the Disk?314 21.10.7 How Is the Snapshot Size Calculated?	21.10.4 How Is a Snapshot Created for My Disk?	313
21.10.7 How Is the Snapshot Size Calculated?	21.10.5 Why Can't I Roll Back My Disk Data from a Snapshot?	313
21.10.8 Do Snapshots Take Space on the Disk?	21.10.6 Can I Roll Back the Data from a Snapshot After Reinstalling the OS or Formatting the Disk?	314
21.10.9 Can I Perform Multiple Rollback Operations for a Snapshot?	21.10.7 How Is the Snapshot Size Calculated?	314
21.10.10 Can I Replicate Snapshots to Other Regions or Accounts?	21.10.8 Do Snapshots Take Space on the Disk?	314
21.10.11 Why Can't I Find My Snapshot?	21.10.9 Can I Perform Multiple Rollback Operations for a Snapshot?	314
21.10.12 Can I Use a Snapshot to Create a Disk and Roll Back the Snapshot at Almost the Same Time? 314 21.11 Backup	21.10.10 Can I Replicate Snapshots to Other Regions or Accounts?	314
314 21.11 Backup	21.10.11 Why Can't I Find My Snapshot?	314
21.11 Backup		
21.11.1 Do I Need to Stop the Server Before Performing a Backup?		
21.11.2 Can I Back Up and Restore My EVS Disk to a Different Region?	21.11 Backup	315
21.11.3 How Do I View My Backup Data? 315 A Appendix 317 A.1 EVS Disk Status 317 A.2 EVS Snapshot Status 319	21.11.1 Do I Need to Stop the Server Before Performing a Backup?	315
A Appendix	21.11.2 Can I Back Up and Restore My EVS Disk to a Different Region?	315
A.1 EVS Disk Status	21.11.3 How Do I View My Backup Data?	315
A.2 EVS Snapshot Status	A Appendix	317
A.2 EVS Snapshot Status		
B Change History321		
	B Change History	321

2023-06-12 vi

1 Overview

1.1 What Is EVS?

Overview

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

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

EVS disks are sometimes just referred to as disks in this document.

Create backups. EVS disk **ECSs** backups EVS disk Attach Create disks to new servers. disks. Restore EVS disk data. EVS disk Create BMSs EVS snapshots snapshots.

Figure 1-1 EVS architecture

EVS Advantages

EVS has the following advantages:

Table 1-1 EVS advantages

Advantage	Description	Related Knowledge
Various disk types	EVS provides a variety of disk types for you to choose from, and EVS disks can be used as data disks and system disks for servers. You can select an appropriate disk type that best suits your budget and service requirements.	Disk Types and Performance
Elastic scalability	The EVS disk capacity ranges from 10 GiB to 32 TiB. When it no longer meets your needs, you can expand the disk capacity up to 32 TiB in increments of 1 GiB, without interrupting your applications.	Expansion Overview
	Besides the disk capacity limit, the additional space you can add cannot exceed the remaining quota. You can increase the quota if the remaining quota is insufficient.	Querying EVS Resource Quotas
High security and reliability	Data redundancy functions, such as backups and snapshots, safeguard the disk data, preventing incorrect data caused by application exceptions or attacks.	EVS Backup EVS Snapshot
Real-time monitoring	On Cloud Eye, you can monitor the disk health and operating status at any time.	Viewing EVS Monitoring Data

Differences Among EVS, SFS, and OBS

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

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

Dimensio n	SFS	OBS	EVS
Concept	SFS provides on- demand high- performance file storage, which can be shared by multiple servers. SFS is similar to a remote directory for Windows or Linux OSs.	OBS provides massive, secure, reliable, and costeffective data storage for users to store data of any type and size.	EVS provides scalable block storage that features high reliability and high performance to meet a variety of service requirements. An EVS disk is similar to a hard disk on a PC.

Dimensio n	SFS	OBS	EVS
Data storage logic	Stores files. Data is sorted and displayed in files and folders.	Stores objects. Files can be stored directly to OBS. The files automatically generate corresponding system metadata. You can also customize the metadata if needed.	Stores binary data and cannot directly store files. To store files, you need to format the file system first.
Access method	SFS file systems can be accessed only after being mounted to ECSs or BMSs through NFS or CIFS. A network address must be specified or mapped to a local directory for access.	OBS buckets can be accessed through the Internet or Direct Connect. The bucket address must be specified for access, and transmission protocols HTTP and HTTPS are used.	EVS disks can be used and accessed from applications only after being attached to ECSs or BMSs and initialized.
Applicatio n Scenarios	Gene sequencing, image rendering, media processing, file sharing, content management, and web services	Big data analytics, static website hosting, online video on demand (VoD), gene sequencing, and intelligent video surveillance	Industrial design, energy exploration, critical clustered applications, enterprise application systems, and development and testing
Capacity	Petabytes	Exabytes	Terabytes
Latency	3–10 ms	10 ms	1 ms
IOPS/TPS	10,000 for a single file system	Tens of millions	50,000 for a single disk
Bandwidt h	GiB/s	TiB/s	MiB/s
Data sharing	Supported	Supported	Supported
Remote access	Supported	Supported	Not supported
Used independe ntly	Supported	Supported	Not supported

Methods of Access

The public cloud system provides a web-based management console and HTTPS-based APIs for you to access the EVS service.

APIs

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

Management console

Use the management console if you do not need to integrate EVS with a third-party system. Log in to the management console and choose **Elastic Volume Service** from the service list.

1.2 Disk Types and Performance

EVS disks are classified based on the disk I/O performance. EVS disks differ in performance and price. Choose the disk type most appropriate for your applications.

EVS Performance

EVS performance metrics include:

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

Single-queue access latencies of different types of EVS disks are as follows:

- High I/O: 1 ms to 3 ms
- Ultra-high I/O: 1 ms

Table 1-3 EVS performance data

Parameter	High I/O	Ultra-high I/O
IOPS per GiB/EVS disk	8	50
Max. IOPS/EVS disk	5,000	50,000
Baseline IOPS/EVS disk	1,800	1,800
IOPS burst limit/EVS disk	5,000	16,000
Disk IOPS	Min. (5,000, 1,800 + 8 x Capacity)	Min. (50,000, 1,800 + 50 x Capacity)
Max. throughput	150 MiB/s	350 MiB/s
Disk throughput	Min. (150, 100 + 0.15 × Capacity) MiB/s	Min. (350, 120 + 0.5 × Capacity) MiB/s

Parameter	High I/O	Ultra-high I/O
API name NOTE This API name indicates the value of the volume_type parameter in the EVS API. It does not represent the type of the underlying hardware device.	SAS	SSD
Typical application scenarios	Mainstream applications requiring high performance and high reliability, such as largescale development and test environments, web server logs, and enterprise applications. Typical enterprise applications include SAP applications, Microsoft Exchange, and Microsoft SharePoint.	Read/write-intensive applications that require ultra-high I/O and throughput, such as distributed file systems used in HPC scenarios or NoSQL and relational databases used in I/O-intensive scenarios. Typical databases include MongoDB, Oracle, SQL Server, MySQL, and PostgreSQL databases.

EVS disk performance is closely related with the data block size. An EVS disk can achieve either the maximum IOPS or maximum throughput depending on which one is reached first.

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

The following uses an ultra-high I/O disk as an example. According to the formula, when the size of an ultra-high I/O disk is greater than or equal to 964 GiB, the disk theoretically can reach either the maximum IOPS 50,000 or the maximum throughput 350 MiB/s. However, this is not the case in practice. The maximum IOPS and maximum throughput that a disk can reach also vary with the data block size. For details, see **Table 1-4**.

Table 1-4 Maximum performance of an ultra-high I/O EVS disk

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

Data Block Size	Max. IOPS	Max. Throughput (MiB/s)
32 KiB	About 11,200	About 350

Disk IOPS Calculation Formula

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

The following example uses an ultra-high I/O EVS disk with a maximum IOPS of 50,000.

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

Disk Burst Capability and Principles

EVS disks have burst capability, which allows a small-capacity disk to surpass its maximum IOPS within a certain period of time. This IOPS applies to individual disks.

Disks with burst capability are well-suited for speeding up server startup. In most cases, system disks have small capacities. For example, the IOPS of a 50-GiB ultrahigh I/O disk without burst capability can only reach up to 4,300, calculated as follows: IOPS = Min. $(50,000, 1,800 + 50 \times Capacity)$. If the disk has burst capability, its IOPS can burst up to 16,000.

The following example uses an ultra-high I/O EVS disk with the IOPS burst limit of 16,000.

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

The following describes the burst IOPS consumption and reservation.

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

Number of initial tokens = Burst duration x IOPS burst limit

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

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

tokens/s, which is the larger value between the disk burst IOPS and maximum IOPS.

Consumption principles

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

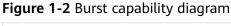
Reservation principles

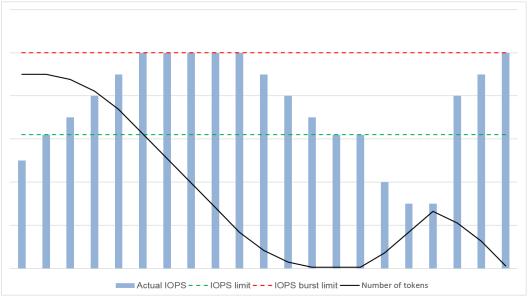
When the token consumption rate is smaller than the production rate, the number of tokens increases accordingly, enabling the disk to regain the burst capability. In this example, if the disk is suspended for approximately 4,235 seconds (28,800,000/6,800), the token bucket will be filled up with tokens.

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

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

- When the number of tokens is greater than zero, the disk IOPS can exceed 6,800 and has the capability to reach 16,000, the IOPS burst limit.
- When the number of tokens is zero, the disk does not have the burst capability, and the maximum IOPS is 6,800.
- When the disk IOPS is less than 6,800, the number of tokens starts to increase, and the disk can regain the burst capability.





Performance Test Method

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

1.3 Device Types and Usage Instructions

What Device Types Are Available?

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

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

Device type is configured during creation. It cannot be changed after the disk has been created.

Common Application Scenarios and Usage Instructions of SCSI EVS Disks

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

SCSI reservations take effect only when shared SCSI EVS disks are attached to ECSs in the same ECS group. For more information about shared EVS disks, see **Shared EVS Disks and Usage Instructions**.

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

To use SCSI EVS disks, you need to install a driver for certain server OSs.

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

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

□ NOTE

ECS virtualization types are categorized into KVM and Xen. For details, see **Product Introduction** > **ECS Types** in the *Elastic Cloud Server User Guide*.

Xen FCS

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

However, a few images support SCSI EVS disks on Xen ECSs. For the supported images, see **Table 1-5**.

□ NOTE

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

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

Table 1-5 OSs supporting SCSI EVS disks

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

1.4 Shared EVS Disks and Usage Instructions

What Are Shared EVS Disks?

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

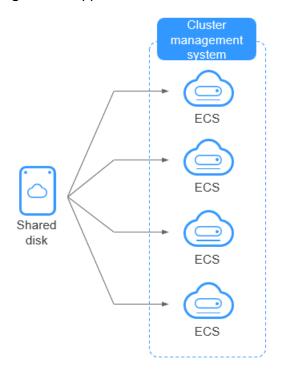
A shared EVS disk can be attached to a maximum of 16 servers. Servers that EVS supports include ECSs and BMSs. To share files, you need to deploy a shared file

system or a cluster management system, such as Windows MSCS, Veritas VCS, or CFS.

NOTICE

You must set up a shared file system or cluster management system before using shared EVS disks. If you directly attach a disk to multiple servers, the sharing function will not work and data may be overwritten.

Figure 1-3 Application scenario of shared EVS disks



Usage Precautions

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

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

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

NOTICE

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

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

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

Advantages

- Multiple attachments: A shared EVS disk can be attached to a maximum of 16 servers.
- High-performance: The random read/write IOPS of a shared ultra-high I/O disk can reach up to 160,000.
- High-reliability: Shared EVS disks support both manual and automatic backup, delivering highly reliable data storage.
- Wide range of use: Shared EVS disks can be used for Linux RHCS clusters where only VBD EVS disks are needed. They can also be used for Windows MSCS and Veritas VCS clusters that require SCSI reservations.

Specifications and Performance

Shared EVS disks have the same specifications and performance as non-shared EVS disks. For details, see **Disk Types and Performance**.

Data Sharing Principle and Common Usage Mistakes

A shared EVS disk is essentially the disk that can be attached to multiple servers for use, which is similar to a physical disk in that the disk can be attached to multiple physical servers, and each server can read data from and write data into any space on the disk. If the data read/write rules, such as the read/write

sequence and meaning, between these servers are not defined, data read/write interference between servers or other unpredictable errors may occur.

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

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

Data inconsistency caused by read/write conflicts

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

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

Data inconsistency caused by data caching

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

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

Before you attach a shared EVS disk to multiple servers, the disk device type needs to be determined. The device type can be either VBD or SCSI. Shared SCSI EVS disks support SCSI reservations. Before using SCSI reservations, you need to install a driver in the server OS and ensure that the OS image is included in the compatibility list.

NOTICE

If you simply attach a shared EVS disk to multiple servers, files cannot be shared between the servers as shared EVS disks do not have the cluster capability. Therefore, build a shared file system or deploy a cluster management system if you need to share files between servers.

1.5 EVS Encryption

What Is EVS Encryption?

In case your services require encryption for the data stored on EVS disks, EVS provides you with the encryption function. You can encrypt newly created EVS disks. Keys used by encrypted disks are provided by the Key Management Service (KMS) of Data Encryption Workshop (DEW), which is secure and convenient. So you do not need to establish and maintain the key management infrastructure.

NOTICE

The encryption attribute of a disk cannot be changed after the disk is created. For details about how to create an encrypted disk, see **Create an EVS Disk**.

Keys Used for EVS Encryption

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

- Default Master Key: A key that is automatically created by EVS through KMS and named evs/default.
 - The Default Master Key cannot be disabled and does not support scheduled deletion.
- CMKs: Keys created by users. You may use existing CMKs or create new CMKs to encrypt disks. For details, see **Key Management Service** > **Creating a CMK** in the *Data Encryption Workshop User Guide*.

When an encrypted disk is attached, EVS accesses KMS, and KMS sends the data key (DK) to the host memory for use. The disk uses the DK plaintext to encrypt and decrypt disk I/Os. The DK plaintext is only stored in the memory of the host housing the ECS and is not stored persistently on the media. If the CMK is disabled or deleted in KMS, the disk encrypted using this CMK can still use the DK plaintext stored in the host memory. If this disk is later detached, the DK plaintext will be deleted from the memory, and data cannot be read from or written to the disk. Before you re-attach this encrypted disk, ensure that the CMK is enabled.

If you use a CMK to encrypt disks and this CMK is then disabled or scheduled for deletion, data cannot be read from or written to these disks or may never be restored. See **Table 1-6** for more information.

CMK Status	Impact	How to Restore
Disabled	For an encrypted disk already attached: Reads and writes to the disk are normal unless	Enable the CMK. For details, see Managing CMKs > Enabling One or More CMKs in the Data Encryption Workshop User Guide.
Scheduled deletion	the disk is detached. Once detached, the disk cannot be attached again. For an encrypted disk not attached: The disk cannot be attached anymore.	Cancel the scheduled deletion for the CMK. For details, see Managing CMKs > Canceling the Scheduled Deletion of One or More CMKs in the Data Encryption Workshop User Guide.
Deleted		Data on the disks can never be restored.

Table 1-6 Impact of CMK unavailability

Encryption Scenarios

System disk encryption

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

Table 1-7 Encryption	relationship) between imag	es and system disks

Creating Server Using an Encrypted Image	Whether the System Disk Will Be Encrypted	Description
Yes	Yes	The encryption attribute of a systems disk is the same as that of the image. For details, see Managing Private Images > Encrypting Images in the Image Management Service User Guide.
No	No	-

Data disk encryption

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

Table 1-8 Encryption relationship between backups, snapshots, images, and data disks

Where the Disk Is Created	Method of Creation	Whether the Data Disk Will Be Encrypted	Description
On the Cloud Server Console	Created together with the server	Yes/No	When a data disk is created together with a server, you can choose to encrypt the disk or not. For details, see Getting Started > Creating an ECS > Step 1: Configure Basic Settings in the Elastic Cloud Server User Guide.
Created Separately on the EVS console	No data source selected	Yes/No	When an empty disk is created, you can choose whether to encrypt the disk or not. The encryption attribute of the disk cannot be changed after the disk has been created.
	Creating from a backup	Yes/No	When a disk is created from a backup, you can choose whether to encrypt the disk or not. The encryption attributes of the disk and backup do not need to be the same.
			When you create a backup for a system or data disk, the encryption attribute of the backup will be the same as that of the disk.
	Creating from a snapshot (The snapshot's source disk is encrypted)	Yes	A snapshot created from an encrypted disk is also encrypted.
	Creating from a snapshot (The snapshot's source disk is not encrypted)	No	A snapshot created from a non-encrypted disk is not encrypted.

Where the Disk Is Created	Method of Creation	Whether the Data Disk Will Be Encrypted	Description
	Creating from an image (The image's source disk is encrypted)	Yes	-
	Creating from an image (The image's source disk is not encrypted)	No	-

Who Can Use the Encryption Function?

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

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

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

Application Scenarios of EVS Encryption

Figure 1-4 shows the user relationships under regions and projects from the perspective of a tenant. The following example uses region B to describe the two scenarios of using the encryption function.

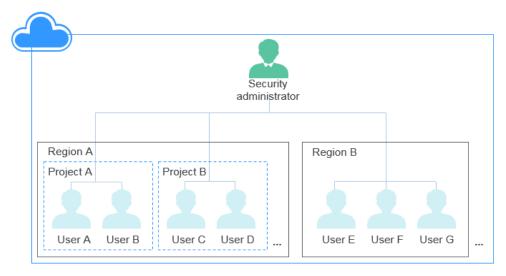


Figure 1-4 User relationships

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

After the KMS access rights have been granted, the system automatically creates a Default Master Key and names it **evs/default**. You can use the Default Master Key to encrypt EVS disks.

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

b. Select a key.

You can select one of the following keys:

- Default Master Key: evs/default
- CMKs: Existing or newly created CMKs. For details, see Key
 Management Service > Creating a CMK in the Data Encryption
 Workshop User Guide.

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

- If User E (common user) uses the encryption function for the first time ever, the operation process is as follows:
 - When user E uses encryption, and the system prompts a message indicating that the KMS access rights have not been granted to EVS.
 - b. Contact the security administrator to grant the KMS access rights to EVS.

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

1.6 EVS Backup

What Is EVS Backup?

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

For more information, see the *Cloud Backup and Recovery User Guide*.

CBR Architecture

CBR involves backups, vaults, and policies.

Backup

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

- Cloud disk backup: provides snapshot-based backups for EVS disks.
- Cloud server backup: uses the consistency snapshot technology to protect data for ECSs and BMSs. Backups of non-database servers are non-database server backups, and those of database servers are application-consistent backups.
- SFS Turbo backup: backs up data of SFS Turbo file systems.

Vault

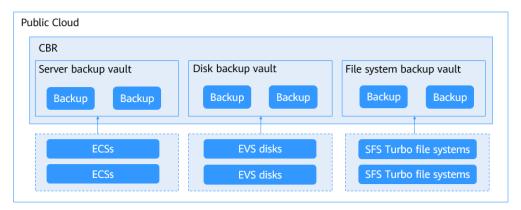
CBR stores backups in vaults. Before creating a backup, you need to create at least one vault and associate the resources you want to back up with the vaults. Then the resources can be backed up to the associated vaults.

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

Policy

 A backup policy defines when you want to take a backup and for how long you would retain each backup.

Figure 1-5 CBR architecture



Application Scenarios

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

Usage Instructions

For how to back up EVS disks, see **Managing EVS Backups** or *Cloud Backup and Recovery User Guide*.

1.7 EVS Snapshot

What Is EVS Snapshot?

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

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

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

Snapshot Principle

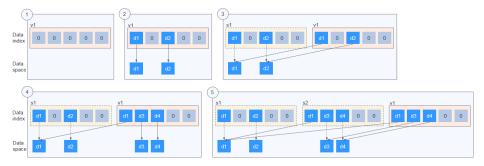
Snapshots and backups are different in that a backup saves the data as another copy in the storage system other than on the disk, whereas a snapshot establishes a relationship between the snapshot and disk data.

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

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

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





Application Scenarios

The snapshot function helps address your following needs:

Routine data backup

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

Rapid data restoration

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

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

When rolling back data from snapshots, data can only be rolled back to the original disk, and a rollback to a different disk is not possible.

Multi-service quick deployment

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

Usage Instructions

For details about the snapshot usages, see Managing EVS Snapshots.

1.8 Differences Between EVS Backups and EVS Snapshots

Both EVS backups and EVS snapshots provide redundancies for improved disk data reliability. **Table 1-9** lists the differences between them.

Table 1-9 Differences between backups and snapshots

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

Metric	Storage Solution	Data Synchronizati on	DR Range	Service Recovery
Snapshot	Snapshots are stored on the same disk as the original data. NOTE Creating a backup requires a certain amount of time because data needs to be transferred to OBS. Creating or rolling back a snapshot consumes less time than creating a backup.	A snapshot is the state of a disk at a specific point in time and is stored on the same disk. If the disk is deleted, all its snapshots will also be deleted. For example, if you reinstalled or changed the server OS, snapshots of the system disk were also automatically deleted. Snapshots of the data disks can be used as usual.	A snapshot and its source disk reside in the same AZ.	You can use a snapshot to roll back its original disk or create a disk from the snapshot.

1.9 EVS Three-Copy Redundancy

What Is the Three-Copy Redundancy?

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

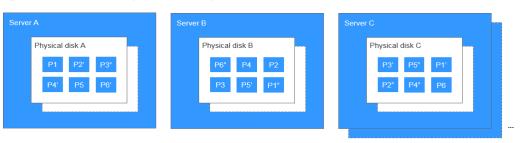
Three-copy redundancy has the following characteristics:

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

For example, for data block P1 on physical disk A of server A, the storage system backs up its data to P1" on physical disk B of server B and to P1' on physical disk C of server C. Data blocks P1, P1', and P1" are the three copies of the same data

block. If physical disk A where P1 resides is faulty, P1' and P1" can continue providing storage services, ensuring service continuity.

Figure 1-7 Three-copy redundancy



How Does the Three-Copy Redundancy Keep Data Consistency?

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

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

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

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

How Does Three-Copy Redundancy Rapidly Rebuild Data?

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

Figure 1-8 shows the data rebuild process.

Figure 1-8 Data rebuild process

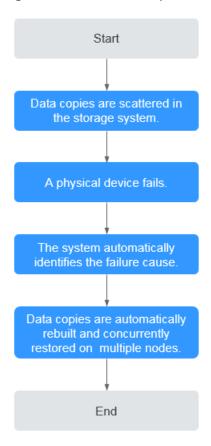


Figure 1-9 shows the data rebuild principle. For example, if physical disks on server F are faulty, the data blocks on these physical disks will be rebuilt on the physical disks of other servers.

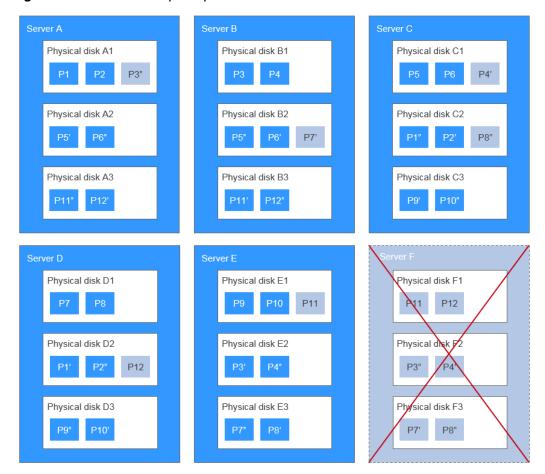


Figure 1-9 Data rebuild principle

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

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

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

1.10 Permissions

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

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

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

IAM is a free service. You only pay for the resources in your account. For more information about IAM, see section "Service Overview" in the *Identity and Access Management User Guide*.

EVS Permissions

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

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

You can grant users permissions by using roles and policies.

- Roles: A type of coarse-grained authorization mechanism that defines
 permissions related to user responsibilities. This mechanism provides only a
 limited number of service-level roles for authorization. When using roles to
 grant permissions, you need to also assign other roles on which the
 permissions depend to take effect. However, roles are not an ideal choice for
 fine-grained authorization and secure access control.
- Policies: A type of fine-grained authorization mechanism that defines
 permissions required to perform operations on specific cloud resources under
 certain conditions. This mechanism allows for more flexible policy-based
 authorization, meeting requirements for secure access control. For example,
 you can grant ECS users only the permissions for managing a certain type of
 ECSs. Most policies define permissions based on APIs. For the API actions
 supported by EVS, see section "Permissions Policies and Supported Actions" in
 the Elastic Volume Service API Reference.

Table 1-10 lists all the system-defined roles and policies supported by EVS.

Table 1-10 System-defined roles and policies supported by EVS

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

Role/Policy Name	Description	Туре	Dependen cy
Server Administrato r	Full permissions for EVS	System role	None

Table 1-11 lists the common operations supported by each system-defined policy of EVS. Select the policies as required.

Table 1-11 Common operations supported by each system-defined policy of EVS

Operation	EVS FullAccess	EVS ReadOnlyAccess
Creating disks	√	x
Viewing disk list	√	√
Viewing disk details	√	√
Attaching disks	√	x
Detaching disks	√	x
Deleting disks	√	x
Expanding disk capacities	√	x
Adding tags for disks	√	x
Modifying tags	√	х
Deleting tags	√	х
Searching for disks by tag	√	√
Changing disk names	√	х

Related Links

- Permissions management sections in the *Identify and Access Management User Guide*
- Creating a User and Granting EVS Permissions
- Section "Permissions Policies and Supported Actions" in the Elastic Volume Service API Reference

1.11 Constraints

This section describes the constraints on using EVS.

Table 1-12 Constraints on using EVS

Scenario	Item	Restrictions
Disk	High I/O	Maximum IOPS per disk: 5,000
performance		Maximum throughput per disk: 150 MiB/s
	Ultra-high I/O	Maximum IOPS per disk: 50,000
		Maximum throughput per disk: 350 MiB/s
Disk creation	Maximum number of disks that can be created at a time	100
	Disk creation from snapshot	The disk type of the new disk is the same as that of the snapshot's source disk.
		The device type of the new disk is the same as that of the snapshot's source disk.
		The encryption attribute of the new disk is the same as that of the snapshot's source disk.
		Batch creation is not supported. One can create only one disk from a snapshot at a time.
	Disk creation from backup	Batch creation is not supported. One can create only one disk from a backup at a time.
		One backup cannot be used for concurrent disk creation operations at the same time. For example, if you are creating disk A from a backup, this backup can be used to create another disk only after disk A has been created.
		 If a disk is created from a backup of a system disk, the new disk can be used as a data disk only.
	Disk creation from image	The device type of the new disk is the same as that of the image's source disk.
		The encryption attribute of the new disk is the same as that of the image's source disk.
	Device type	The device type of a disk cannot be changed after the disk has been created.

Scenario	Item	Restrictions
	Disk sharing	The sharing attribute of a disk cannot be changed after the disk has been created.
	Disk encryption	The encryption attribute of a disk cannot be changed after the disk has been created.
Disk attachment	Constraints on region and AZ	The disk and server must be in the same region and AZ.
	Maximum number of servers that a non-shared disk can be attached to	1
	Maximum number of servers that a shared disk can be attached to	16
	Maximum number of disks that can be attached to an ECS	This value varies with ECS types. For details, see section "Can I Attach Multiple Disks to an ECS?" in the <i>Elastic Cloud Server User Guide</i> .
	Maximum number of disks that can be attached to a BMS	60 (1 system disk and 59 data disks)
Disk capacity expansion	Capacity expansion	Disk capacities can only be expanded, not reduced.
	Capacity expansion of non-shared disks	Some server OSs support the capacity expansion of non-shared, In-use disks. For details, see Expanding Capacity for an In-use EVS Disk.
	Capacity expansion of shared disks	A shared disk must be detached from all its servers before expansion. That is, the shared disk status must be Available .
	Expansion increment	1 GiB
Disk detachment	System disk detachment	A system disk can only be detached offline, which means that the server must be in the Stopped state.
	Data disk detachment	A data disk can be detached online or offline, that is, its server can either be in the Running or Stopped state.

Scenario	Item	Restrictions
Disk deletion	Deletion of disks	 The disk status is Available, Error, Expansion failed, Restoration failed, or Rollback failed. The disk is not locked by any service. The shared disk has been detached from all its servers.
Snapshot	Maximum number of snapshots that can be created for a disk	7
	Maximum number of disks that can be created from a snapshot	128
	Snapshot data rollback to disk	A snapshot can be rolled back only to its source disk. Rollback to another disk is not possible.
		 A snapshot can be rolled back only when the snapshot status is Available and the source disk status is Available (not attached to any server) or Rollback failed.
		A snapshot whose name starts with autobk_snapshot_vbs_, manualbk_snapshot_vbs_, autobk_snapshot_csbs_, or manualbk_snapshot_csbs_ is automatically generated during backup. Such a snapshot can only be viewed. It cannot be used to roll back the disk data. If the first till the back automatically generated to roll back the disk data.
		 If you first roll back the snapshot to the original disk, you cannot use the snapshot to create a new disk then.
Tag	Maximum number of tags that can be added for a disk	10

1.12 EVS and Other Services

Figure 1-10 shows the relationships between EVS and other services.

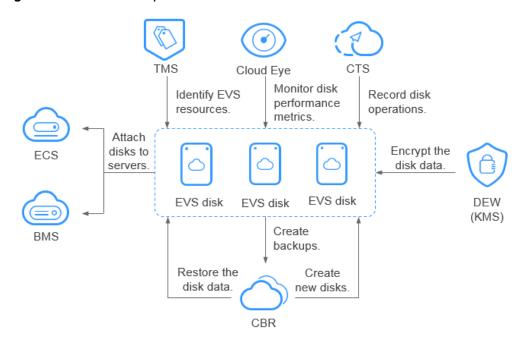


Figure 1-10 Relationships between EVS and other services

Table 1-13 EVS and other services

Interactive Function	Related Service	Reference
EVS disks can be attached to ECSs and used as scalable block storage devices.	ECS	• Attaching a Non-Shared Disk
SCSI EVS disks can be attached to BMSs and used as scalable block storage devices.	BMS	Attaching a Shared Disk
EVS disk encryption depends on the KMS service in DEW. Keys provided by KMS can be used to encrypt EVS disks (both system and data disks), thus improving EVS disk data security.	DEW	 EVS Encryption Managing Encrypted EVS Disks
After EVS is enabled, the performance metrics of monitored disks can be viewed through Cloud Eye without installing any additional plug-in. The monitored metrics include Disk Read Rate, Disk Write Rate, Disk Read Requests, and Disk Write Requests.	Cloud Eye	Viewing EVS Monitoring Data
Cloud Trace Service (CTS) records operations of EVS resources, facilitating user query, audit, and backtracking.	CTS	Auditing

Interactive Function	Related Service	Reference
Tag Management Service (TMS) tags are used to identify EVS resources for purposes of easy categorization and quick search.	TMS	Adding a Tag

1.13 Basic Concepts

1.13.1 EVS Concepts

Table 1-14 EVS concepts

Concept	Description
IOPS	Number of read/write operations performed by an EVS disk per second
Throughput	Amount of data read from and written into an EVS disk per second
Read/write I/O latency	Minimum interval between two consecutive read/write operations of an EVS disk
Burst capability	The burst capability allows the IOPS of a small-capacity disk to reach the disk IOPS burst limit, which can surpass the disk IOPS limit within a certain period of time.
VBD	A device type of EVS disks. VBD EVS disks only support basic SCSI read/write commands.
SCSI	A device type of EVS disks. SCSI EVS disks support transparent SCSI command transmission and allow the server OS to directly access the underlying storage media.

1.13.2 Region and AZ

Concept

A region and availability zone (AZ) identify the location of a data center. You can create resources in a specific region and AZ.

- A region is a physical data center, which is completely isolated to improve fault tolerance and stability. The region that is selected during resource creation cannot be changed after the resource is created.
- An AZ is a physical location where resources use independent power supplies and networks. A region contains one or more AZs that are physically isolated

2023-06-12

but interconnected through internal networks. Because AZs are isolated from each other, any fault that occurs in one AZ will not affect others.

Figure 1-11 shows the relationship between regions and AZs.

Figure 1-11 Regions and AZs



Selecting a Region

Select a region closest to your target users for lower network latency and quick access.

Selecting an AZ

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

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

Regions and Endpoints

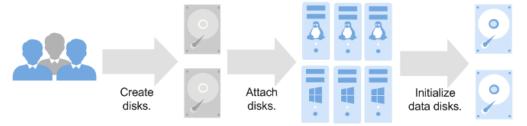
Before you use an API to call resources, specify its region and endpoint. For more details, see **Regions and Endpoints**.

2 Getting Started

2.1 Process Overview

Figure 2-1 shows the EVS process overview.

Figure 2-1 Process overview



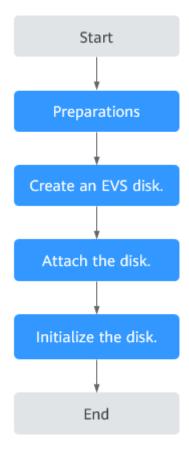
EVS disks can be attached to servers and be used as system disks or data disks. Table 2-1 lists the disk purchasing methods.

Table 2-1 Method of creation

Disk	Description	Method
System disk	System disks are created along with servers and cannot be created separately.	 See section "Creating an ECS" in the Elastic Cloud Server User Guide. See section "Creating a BMS" in the Bare Metal Server User Guide.
Data disk	Data disks can be created along with servers or separately.	 See section "Creating an ECS" in the Elastic Cloud Server User Guide. See section "Creating a BMS" in the Bare Metal Server User Guide. Create an EVS Disk

Figure 2-2 shows how to purchase a data disk separately.

Figure 2-2 Process overview



- 1. **Preparations**: Register an account on the console and obtain the permissions required to create servers and disks.
- 2. **Create a disk.** Configure the disk parameters, including the disk type, size, name, and other information. For details, see **Create an EVS Disk**.
- 3. **Attach the data disk.** Attach the separately created disk to a server. For details, see the following sections:
 - Attaching a Non-Shared Disk
 - Attaching a Shared Disk
- 4. **Initialize the data disk.** Log in to the server and initialize the data disk before using it. For details about how to initialize the disk, see the following sections:
 - Introduction to Data Disk Initialization Scenarios and Partition Styles
 - Windows
 - Initializing a Windows Data Disk (Windows Server 2008)
 - Initializing a Windows Data Disk (Windows Server 2019)
 - Initializing a Windows Data Disk Larger Than 2 TiB (Windows Server 2008)
 - Initializing a Windows Data Disk Larger Than 2 TiB (Windows Server 2012)

- Linux
 - Initializing a Linux Data Disk (fdisk)
 - Initializing a Linux Data Disk (parted)
 - Initializing a Linux Data Disk Larger Than 2 TiB (parted)

2.2 Create an EVS Disk

Scenarios

EVS disks can be used as system disks or data disks for servers.

- System disks can only be created together with servers and are automatically attached.
- A system disk can have a maximum of 1,024 GiB, and a data disk 32,768 GiB.
- Data disks created together with servers are automatically attached.
- After servers are created, any data disks added on the cloud server console are automatically attached.
- After servers are created, any data disks added on the EVS console need to be manually attached.

■ NOTE

- By default, disks created with ECSs are VBD disks, and those created with BMSs are SCSI disks.
- Disks and their attached servers must reside in the same AZ. The AZ cannot be changed after a disk is created.

This section describes how to separately create disks on the EVS console.

Procedure

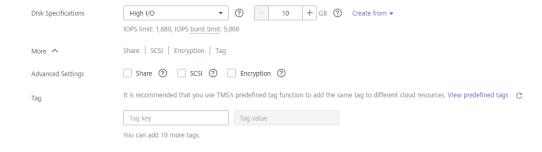
- **Step 1** Log in to the management console.
- Step 2 Under Storage, click Elastic Volume Service.

The disk list page is displayed.

Step 3 Click Create Disk.

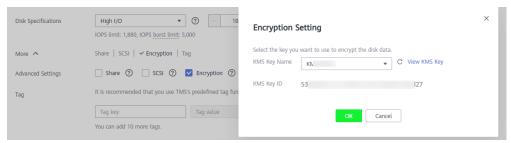
• The following figure shows the parameter setting of non-encrypted disks.

Figure 2-3 Non-encrypted disks



• The following figure shows the parameter setting of encrypted disks.

Figure 2-4 Encrypted disks



Step 4 Configure basic disk information according to **Table 2-2**.

Table 2-2 Disk parameters

Paramete r	Sub- Paramete r	Description	Example Value
Region	-	Mandatory Resources are region-specific and cannot be used across regions through internal network connections. For low network latency and quick resource access, select the nearest region.	-
AZ	-	Mandatory The availability zone (AZ) where you want to create the disk. NOTE Disks can only be attached to the servers in the same AZ. The AZ of a disk cannot be changed after the disk has been created.	-
Disk Specificati ons	Disk Type	Mandatory The available disk types are as follows: • High I/O • Ultra-high I/O NOTE When a disk is created from a snapshot, the disk type of the new disk will be consistent with that of the snapshot's source disk.	Ultra-high I/O
	Disk Size (GiB)	Mandatory The disk size. Only data disks can be created on the current page, and the disk size ranges from 10 GiB to 32,768 GiB.	100 GiB

Paramete r	Sub- Paramete r	Description	Example Value
	Create from Backu p Snapsh ot Image	 Optional If you choose Create from Backup, the backup data is used to create the disk. Click Create from and choose Backup. On the displayed page, select the target backup and click OK. NOTE One backup cannot be used for concurrent disk creation operations at the same time. For example, if you are creating disk A from a backup, this backup can be used to create another disk capturafter disk. 	• Create from Backup: autobacku p-001
		create another disk only after disk A has been created. If a disk is created from a backup of a system disk, the new disk can be used as a data disk only. If you choose Create from Snapshot, the snapshot data is used to create the disk. Click Create from and choose Snapshot. On the displayed page, select the target snapshot and click OK.	
		 NOTE The device type of the new disk is the same as that of the snapshot's source disk. For details about the disk creation from snapshots, see Creating an EVS Disk from a Snapshot. If you choose Create from Image, the image data is used to create the disk. Click Create from and choose Image. On the displayed page, select the target image and click OK. NOTE The device type of the new disk is the same as that of the image's source disk. 	

Sub- Paramete r	Description	Example Value
Advanced Settings Share SCSI Encryption	 Optional Share If you select Share, a shared disk is created. A shared disk can be attached to up to 16 servers. If you do not select Share, a non-shared disk is created, and the disk can be attached to one a server only. If you select both SCSI and Share, a shared SCSI disk is created. NOTE The sharing attribute of a disk cannot be changed after the disk has been created. For details about shared EVS disks, see Managing Shared EVS Disks. SCSI If you select SCSI, a SCSI disk is created. Such disks allow the server OS to directly access the underlying storage media and send SCSI commands to the disks. If you do not select SCSI, a VBD disk is created. That said, the disk device type is VBD, the default device type. NOTE	
	Management Service User Guide. To use the disk encryption function, select Encryption . Then configure the encryption parameters as prompted: - Create Agency An agency is a trust relationship	
	Paramete r Advanced Settings Share SCSI Encryp	Advanced Settings Share ScSI Encryption Share SCSI Share Share SCSI Share Share Share SCSI Share Share SCSI Share SCSI Share SCSI Share SCSI Share, a non-shared disk is created, and the disk can be attached to one a server only. If you select both SCSI and Share, a shared SCSI disk is created. NOTE The sharing attribute of a disk cannot be changed after the disk has been created. For details about shared EVS disks, see Managing Shared EVS Disks. SCSI If you select SCSI, a SCSI disk is created. Such disks allow the server OS to directly access the underlying storage media and send SCSI commands to the disks. If you do not select SCSI, a VBD disk is created. That said, the disk device type is VBD, the default device type. NOTE The device type of a disk cannot be changed after the disk has been created. For details about the ECS types, OSs, and ECS software supported by SCSI EVS disks, see Device Types and Usage Instructions. Encryption The encryption function on this console is used for data disk encryption relies on the image. For details, see the Image Management Service User Guide. To use the disk encryption function, select Encryption. Then configure the encryption parameters as prompted: Create Agency

Paramete r	Sub- Paramete r	Description	Example Value
		A tenant can create an agency to grant resource access rights to another tenant or service. If the KMS access rights are not granted to EVS, the Create Agency dialog box will be displayed. Otherwise, it will not be displayed.	
		Click Yes to grant the KMS access rights to EVS. After the rights have been granted, EVS can obtain KMS keys to encrypt or decrypt EVS disks.	
		After the KMS access rights have been granted, follow-up operations do not require the rights to be granted again.	
		 KMS Key Name KMS Key Name is displayed only after the KMS access rights have been granted. 	
		KMS Key Name is the identifier of the key, and you can use KMS Key Name to specify the KMS key that is to be used for encryption. One of the following keys can be used:	
		Default Master Key: After the KMS access rights have been granted to EVS, the system automatically creates a Default Master Key and names it evs/default.	
		CMKs: Existing or newly created CMKs. For details, see Key Management Service > Creating a CMK in the <i>Data Encryption Workshop User Guide</i> .	

Paramete r	Sub- Paramete r	Description	Example Value
		NOTE - Before you use the encryption function, KMS access rights need to be granted to EVS. If you have the right to grant the permission, grant the KMS access rights to EVS directly. If you do not have this permission, contact a user with the security administrator rights to grant KMS access rights to EVS, then repeat the preceding operations. - The encryption attribute of a disk cannot be changed after the disk has been created.	
	Tag	Optional During the EVS disk creation, you can tag the EVS resources. Tags identify cloud resources for purposes of easy categorization and quick search. A tag is composed of a key-value pair. Key: Mandatory if the disk is going to be tagged A tag key can contain a maximum of 36 characters, including letters,	-
		 digits, and underscores (_). Value: Optional if the disk is going to be tagged A tag value can contain a maximum of 43 characters, including letters, digits, underscores (_), periods (.), and hyphens (-). 	
		 NOTE A maximum of 10 tags can be added for an EVS disk. Tag keys of the same EVS disk must be unique. Except for tagging the disk during disk creation, you can also add, modify, or delete tags for existing disks. For details, see Managing a Tag. For details about tags, see the Tag Management Service User Guide. 	

Paramete r	Sub- Paramete r	Description	Example Value
Disk Name	-	 If you create disks individually, this parameter value is used as the actual disk name. The name can contain a maximum of 64 characters. If you create disks in a batch, this parameter value is used as the prefix of disk names, and one disk name will be composed of this parameter value and a four-digit number. The name can contain a maximum of 59 characters. 	For example, if you create two disks and set volume for Disk Name, the EVS disk names will be volume-0001 and volume-0002.
Quantity	-	Optional The number of disks to be created. The default value is set to 1, which means only one disk is created. Currently, you can create up to 100 disks at a time.	1

Step 5 Click Next.

Step 6 On the **Details** page, check the disk details.

- If you do not need to modify the specifications, click **Submit**.
- If you need to modify the specifications, click **Previous**.

Step 7 In the disk list, view the disk status.

When the disk status changes to **Available**, the disk is successfully created.

----End

2.3 Attach an EVS Disk

2.3.1 Attaching a Non-Shared Disk

Scenarios

Separately created EVS disks are data disks. In the disk list, the function of such disks is displayed as **Data disk**, and the status is displayed as **Available**. In this case, you need to attach the data disks to servers for use.

A system disk must be created together with an ECS and is automatically attached. In the disk list, the function of such disks is displayed as **System disk**, and the status is displayed as **In-use**. After a system disk is detached from an ECS, the disk function changes to **Bootable disk**, and the status changes to **Available**.

Bootable disks are the system disks detached from servers. A bootable disk can be reattached to a server and be used as a system disk or data disk depending on the disk function selected. For details, see **Attaching an Existing System Disk**.

This section describes how to attach a non-shared disk.

Constraints

- Cloud servers created from ISO images are only used for OS installation. They have limited functions and cannot have EVS disks attached.
- A non-shared disk can be attached to one server only.
- The disk and the server must be in the same region and AZ.

Attaching the Disk on the EVS Console

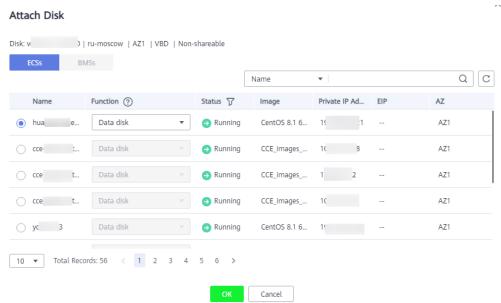
- **Step 1** Log in to the management console.
- Step 2 Under Storage, click Elastic Volume Service.

The disk list page is displayed.

- **Step 3** Locate the target disk in the list and click **Attach**.
- **Step 4** Select the server and then select the disk function from the drop-down list. Ensure that the disk and server are in the same AZ.

One device name can be used for one disk only. For how to obtain the disk name in the OS, see section "How Do I Obtain My Disk Name in the ECS OS Using the Device Identifier Provided on the Console?" in the *Elastic Cloud Server User Guide*.

Figure 2-5 Attach Disk



Step 5 Click **OK** to return to the disk list page.

The status of the disk is **Attaching**, indicating that the disk is being attached to the server. When the disk status changes to **In-use**, the disk is successfully attached.

Step 6 Initialize the disk.

After the disk has been attached to a server, the disk can be used only after you have initialized it. For details, see Introduction to Data Disk Initialization Scenarios and Partition Styles.

----End

Attaching the Disk on the ECS Console

- 1. Log in to the management console.
- 2. Click in the upper left corner and select your region and project.
- 3. Under Computing, click Elastic Cloud Server.
- 4. In the search box above the upper right corner of the ECS list, enter the ECS name, IP address, or ID for search.
- Click the name of the target ECS.
 The page providing details about the ECS is displayed.
- 6. Click the **Disks** tab. Then, click **Attach Disk**.

The **Attach Disk** dialog box is displayed.

- 7. Select the target disk and specify it as the system disk or a data disk.
 - For KVM ECSs, you can specify the disk as the system disk or a data disk but cannot specify a specific device name.
 - For Xen ECSs, you can specify a specific device name, such as /dev/vdb.

- If no disks are available, click Create Disk in the lower part of the list.
- For the restrictions on attaching disks, see FAQ "What Are the Requirements for Attaching an EVS Disk to an ECS?" in the *Elastic Cloud Server User Guide*.
- 8. Click OK.

After the disk is attached, you can view the information about it on the **Disks** tab.

Follow-Up Operations

If you are attaching a new disk, you must then log in to the server and initialize the disk before it can be used. To learn how to initialize disks, see **Introduction to Data Disk Initialization Scenarios and Partition Styles**.

Related Operations

If your disk cannot be attached to a server, see Why Can't I Attach My Disk to a Server?.

If the disk you are going to attach contains data, see Attaching an Existing Disk.

If the attached data disk is not showing up, see Why Can't I View the Attached Data Disk on the Server?.

2.3.2 Attaching a Shared Disk

Scenarios

Separately created shared EVS disks are data disks. In the disk list, the function of such a disk is displayed as **Data disk**, and the status is displayed as **Available**. In this case, you need to attach the shared data disk to servers for use.

Constraints

NOTICE

If you simply attach a shared disk to multiple servers, files cannot be shared among them. Because there are no mutually agreed data read/write rules among servers, read and write operations from them may interfere with each other, or unpredictable errors may occur. To share files between servers, set up a shared file system or a clustered management system first.

- A shared disk can be attached to a maximum of 16 servers. These servers and the shared disk must be in the same AZ within a region.
- A shared, In-use disk can be attached to other servers only when the maximum number of servers that the disk can be attached to has not been reached.
- All the servers attached with a shared disk must run either Windows or Linux.
 For example, if you attach a shared disk to multiple Windows servers and then detach it from these servers, the shared disk cannot be attached to Linux servers later. This is because Windows and Linux support different file systems and cannot identify the original file system on the disk. Improper operations may damage the original file system.
- A shared disk can only be used as a data disk. It cannot be used as a system disk.
- Cloud servers created from ISO images are only used for OS installation. They have limited functions and cannot have EVS disks attached.

Attaching the Disk on the EVS Console

- **Step 1** Log in to the management console.
- **Step 2** Under **Storage**, click **Elastic Volume Service**.

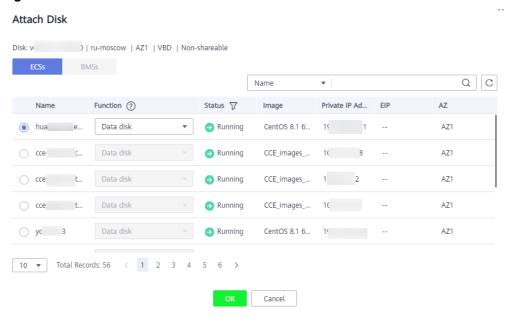
The disk list page is displayed.

Step 3 Locate the target disk in the list and click **Attach**.

Shared disks support batch attachment so that you can attach a shared disk to multiple servers. The left area in the **Attach Disk** dialog box shows the server list. After you select the target servers, the selected servers will be displayed in the right area.

Step 4 Select the target servers you want to attach the shared disk. Ensure that the disk and servers are in the same AZ. After you select the servers, **Data disk** is automatically entered as the disk function for each server.

Figure 2-6 Attach Disk



Step 5 Click **OK** to return to the disk list page.

The status of the disk is **Attaching**, indicating that the disk is being attached to the servers. When the disk status changes to **In-use**, the disk is successfully attached.

----End

Attaching the Disk on the ECS Console

- 1. Log in to the management console.
- 2. Click in the upper left corner and select the desired region and project.
- 3. Under Compute, click Elastic Cloud Server.
- 4. In the search box above the upper right corner of the ECS list, enter the ECS name, IP address, or ID for search.
- Click the name of the target ECS.
 The page providing details about the ECS is displayed.
- 6. Click the **Disks** tab. Then, click **Attach Disk**.
 - The **Attach Disk** page is displayed.
- 7. Select the target disk and specify it as the system disk or a data disk.
 - For Xen ECSs, you can specify a specific device name, such as /dev/sdb.
 - For KVM ECSs, you can specify the disk as a system disk or data disk but cannot specify a device name.

□ NOTE

- If no disks are available, click Create Disk in the lower part of the list.
- For the restrictions on attaching disks, see FAQ "What Are the Requirements for Attaching an EVS Disk to an ECS?" in the *Elastic Cloud Server User Guide*.

8. Click OK.

After the disk is attached, you can view the information about it on the **Disks** tab

Follow-Up Operations

If you are attaching a new disk, you must then log in to the server and initialize the disk before it can be used. To learn how to initialize disks, see **Introduction to Data Disk Initialization Scenarios and Partition Styles**.

Related Operations

If your disk cannot be attached to a server, see Why Can't I Attach My Disk to a Server?.

If the disk you are going to attach contains data, see Attaching an Existing Disk.

If the attached data disk is not showing up, see Why Can't I View the Attached Data Disk on the Server?.

2.4 Initialize an EVS Data Disk

2.4.1 Introduction to Data Disk Initialization Scenarios and Partition Styles

Scenarios

After a disk is attached to a server, you need to log in to the server to initialize the disk, that is, format the disk. You must initialize a disk before accessing it.

System disk

A system disk does not require manual initialization because it is automatically created and initialized upon server creation. The default partition style is master boot record (MBR).

- Data disk
 - If a data disk is created along with a server, it will be automatically attached to the server.
 - If a data disk is created separately, you need to manually attach it to a server.

In both cases, you must initialize the data disk before using it. Choose an appropriate partition style based on your service plan.

Constraints

A disk created from a data source does not need to be initialized. Such a disk contains the data of the source in the beginning. Initializing the disk may clear the initial data on it.

Disk Partition Styles

Table 2-3 lists the common disk partition styles. In Linux, different partition styles require different partitioning tools.

Table 2-3 Disk partition styles

Disk Partition Style	Maximum Disk Capacity Supported	Maximum Number of Partitions Supported	Linux Partitioning Tool
Master Boot	2 TiB	4 primary partitions	• fdisk
Record (MBR)		3 primary partitions and 1 extended partition	• parted
		With MBR, you can create several primary partitions and one extended partition. The extended partition must be divided into logical partitions before use. For example, if 6 partitions need to be created, you can create them in the following two ways:	
		3 primary partitions and 1 extended partition, with the extended partition divided into 3 logical partitions	
		1 primary partition and 1 extended partition, with the extended partition divided into 5 logical partitions	
GUID	18 EiB	Unlimited	parted
Partition Table (GPT)	1 EiB = 1048576 TiB	Disk partitions created using GPT are not categorized.	

NOTICE

The maximum disk size supported by MBR is 2 TiB, and that supported by GPT is 18 EiB. Because an EVS data disk currently supports up to 32 TiB, use GPT if your disk size is larger than 2 TiB.

If the partition style is changed after the disk has been used, all data on the disk will be lost, so take care to select an appropriate partition style when initializing the disk. If you must change the partition style to GPT after a disk has been used, it is recommended that you back up the disk data before the change.

Partitioning Operation Guide

For a disk smaller than 2 TiB, see the following sections:

- Initializing a Windows Data Disk (Windows Server 2008)
- Initializing a Windows Data Disk (Windows Server 2019)
- Initializing a Linux Data Disk (fdisk)
- Initializing a Linux Data Disk (parted)

For a disk larger than 2 TiB, see the following sections:

- Initializing a Windows Data Disk Larger Than 2 TiB (Windows Server 2008)
- Initializing a Windows Data Disk Larger Than 2 TiB (Windows Server 2012)
- Initializing a Linux Data Disk Larger Than 2 TiB (parted)

2.4.2 Initializing a Windows Data Disk (Windows Server 2008)

Scenarios

This section uses Windows Server 2008 R2 Enterprise 64bit to describe how to initialize a data disk attached to a server running Windows.

The maximum disk capacity supported by MBR is 2 TiB, and that supported by GPT is 18 EiB. Therefore, use the GPT partition style if your disk capacity is larger than 2 TiB. For details, see Initializing a Windows Data Disk Larger Than 2 TiB (Windows Server 2008). To learn more about disk partition styles, see Introduction to Data Disk Initialization Scenarios and Partition Styles.

The method for initializing a disk varies slightly depending on the OS running on the server. This document is used for reference only. For the detailed operations and differences, see the product documents of the corresponding OS.

NOTICE

When using a disk for the first time, if you have not initialized it, including creating partitions and file systems, the additional space added to this disk in an expansion later may not be normally used.

Prerequisites

- A data disk has been attached to a server and has not been initialized.
- You have logged in to the server.
 - For how to log in to an ECS, see the *Elastic Cloud Server User Guide*.
 - For how to log in to a BMS, see the Bare Metal Server User Guide.

Procedure

Step 1 On the desktop of the server, right-click **Computer** and choose **Manage** from the shortcut menu.

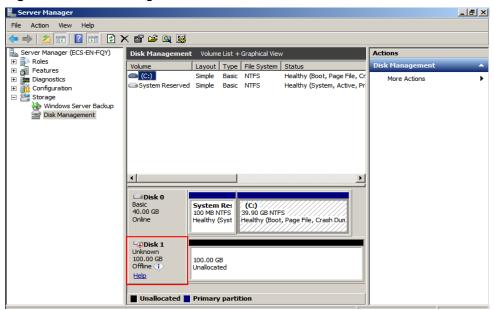
The **Server Manager** window is displayed.

Step 2 In the navigation tree, choose **Storage** > **Disk Management**.

The **Disk Management** window is displayed.

- If Figure 2-7 is displayed, the new disk is offline. Go to Step 3.
- If Figure 2-10 is displayed, the Initialize Disk window is prompted. Go to Step 5.

Figure 2-7 Disk Management



Step 3 Disks are displayed in the right pane. In the **Disk 1** area, right-click **Offline** and choose **Online** from the shortcut menu to online the disk.

_ B × Action View Help Server Manager (ECS-EN-FQY)

Roles Disk Management Actions Volume List + Graphical View Layout Type File System Status Features
Diagnostics
Configuration
Storage Volume Healthy (Boot, Page File, Cr More Actions System Reserved Simple Basic NTFS Healthy (System, Active, Pr 🐌 Windows Server Backup Disk Management □Disk 0 (C:) 39.90 GB NTFS Healthy (Boot, Page File, Crash Dun Basic 40.00 GB System Res Healthy (Syst Disk 1 Unknown 100 00 GB Offline Help

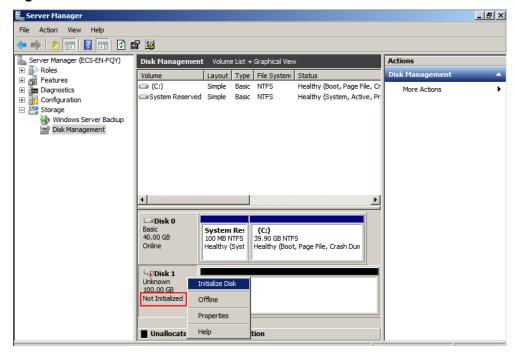
Figure 2-8 Online the disk

Ⅲ NOTE

If the disk is offline, you need to online the disk before initializing it.

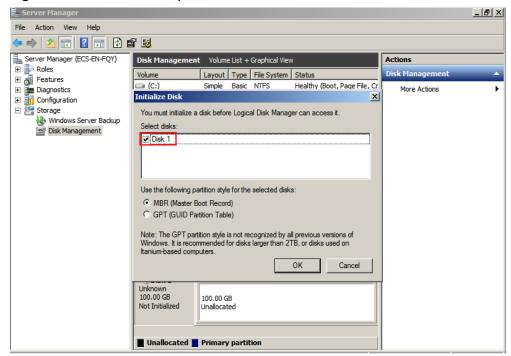
Step 4 After making the disk online, the disk status changes from **Offline** to **Not Initialized**. Right-click the disk status and choose **Initialize Disk** from the shortcut menu, as shown in **Figure 2-9**.

Figure 2-9 Initialize Disk



Step 5 In the Initialize Disk dialog box, select the target disk, click MBR (Master Boot Record) or GPT (GUID Partition Table), and click OK, as shown in Figure 2-10.

Figure 2-10 Unallocated space



NOTICE

The maximum disk size supported by MBR is 2 TiB, and that supported by GPT is 18 EiB. Because an EVS data disk currently supports up to 32 TiB, use GPT if your disk size is larger than 2 TiB.

If the partition style is changed after the disk has been used, all data on the disk will be lost, so take care to select an appropriate partition style when initializing the disk. If you must change the partition style to GPT after a disk has been used, it is recommended that you back up the disk data before the change.

Step 6 Right-click at the unallocated space and choose **New Simple Volume** from the shortcut menu, as shown in **Figure 2-11**.

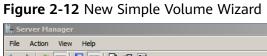
Server Manager _ B × File Action View Help (- Lack Server Manager (ECS-B704) Actions Volume List + Graphical Vie Roles
Features
Diagnostics
Expert Viewer
OPErformance Layout Type File System Status Volume Simple Basic NTFS Healthy (Boot, Page File, Crash Dump, More Actions System Reserved Simple Basic NTFS Healthy (System, Active, Primary Parti Device Manager

Configuration

Storage Windows Server Backup Disk Management Disk 0 Basic 50.00 GB Online System Reser 100 MB NTFS Healthy (Boot, Page File, Cra Healthy (System Disk 1 Properties Basic 100.00 GB 100.00 GB Help ■ Unallocated ■ Primary partition

Figure 2-11 New Simple Volume

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





Step 8 Specify the volume size and click **Next**. The default value is the maximum size.

2023-06-12 53

File Action View Help 🗦 🔷 | 🚈 📊 | 🛭 🕠 🗗 🗳 📽 Server Manager (ECS-B704)

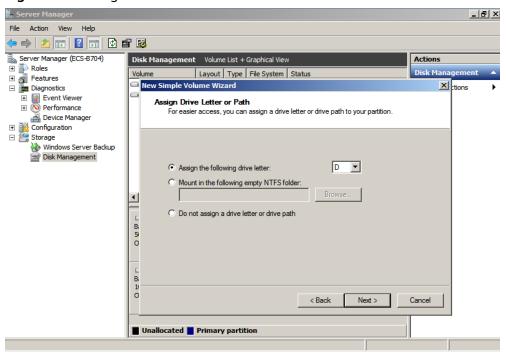
Roles ent Volume List + Graphical Viev Actions Layout Type File System Status × tions Specify Volume Size Choose a volume size that is between the maximum and minimum sizes Device Manager

Configuration Storage
Windows Server Backup Disk Management Maximum disk space in MB: Minimum disk space in MB: 8 Simple volume size in MB: 102397 Bi 5i O < Back Next > Cancel Unallocated Primary partition

Figure 2-13 Specify Volume Size

Step 9 Assign the driver letter and click Next.





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

Figure 2-15 Format Partition

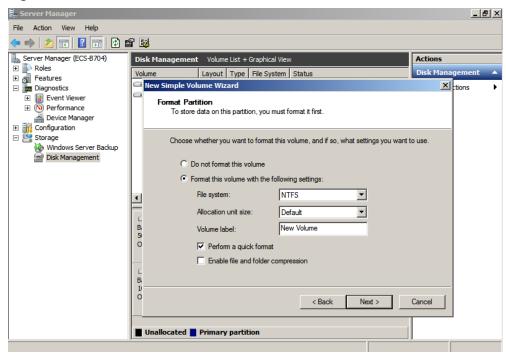
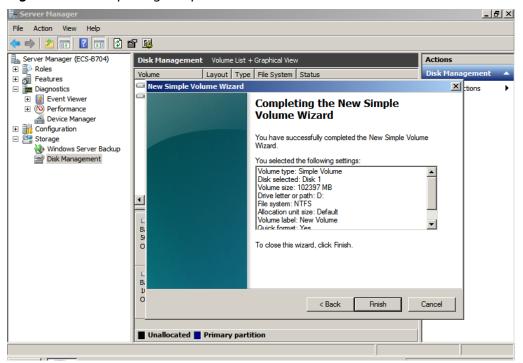


Figure 2-16 Completing the partition creation

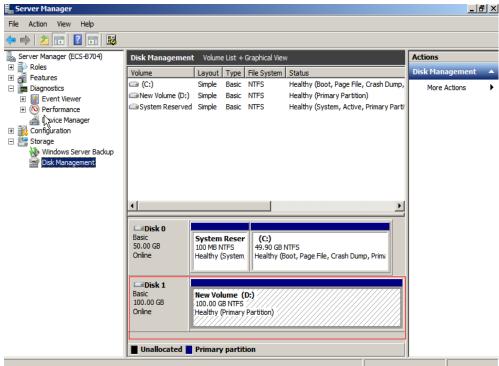


NOTICE

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

Step 11 Click Finish. Wait for the initialization to complete. When the volume status changes to **Healthy**, the initialization has finished successfully, as shown in Figure 2-17.

Figure 2-17 Disk initialization succeeded



----End

2.4.3 Initializing a Windows Data Disk (Windows Server 2019)

Scenarios

This section uses Windows Server 2019 Standard 64bit to describe how to initialize a data disk attached to a server running Windows.

The maximum disk capacity supported by MBR is 2 TiB, and that supported by GPT is 18 EiB. Therefore, use the GPT partition style if your disk capacity is larger than 2 TiB. For details, see Initializing a Windows Data Disk Larger Than 2 TiB (Windows Server 2008). To learn more about disk partition styles, see Introduction to Data Disk Initialization Scenarios and Partition Styles.

The method for initializing a disk varies slightly depending on the OS running on the server. This document is used for reference only. For the detailed operations and differences, see the product documents of the corresponding OS.

NOTICE

When using a disk for the first time, if you have not initialized it, including creating partitions and file systems, the additional space added to this disk in an expansion later may not be normally used.

Prerequisites

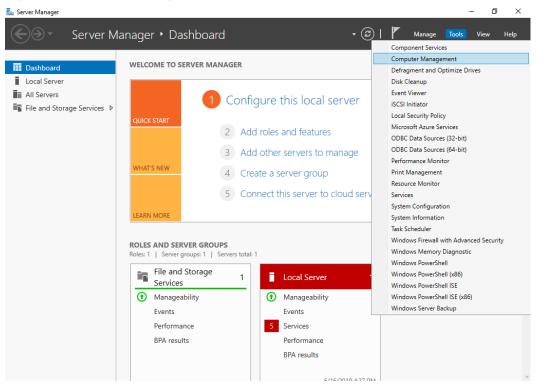
- A data disk has been attached to a server and has not been initialized.
- You have logged in to the server.
 - For how to log in to an ECS, see the Elastic Cloud Server User Guide.
 - For how to log in to a BMS, see the Bare Metal Server User Guide.

Procedure

- **Step 1** On the desktop of the server, click the start icon in the lower left corner.
 - The Windows Server window is displayed.
- Step 2 Click Server Manager.

The **Server Manager** window is displayed.

Figure 2-18 Server Manager



Step 3 In the upper right corner, choose **Tools** > **Computer Management**.

The Computer Management window is displayed.

Figure 2-19 Computer Management

Step 4 Choose **Storage** > **Disk Management**.

Disks are displayed in the right pane. If there is a disk that is not initialized, the system will prompt you with the **Initialize Disk** dialog box.

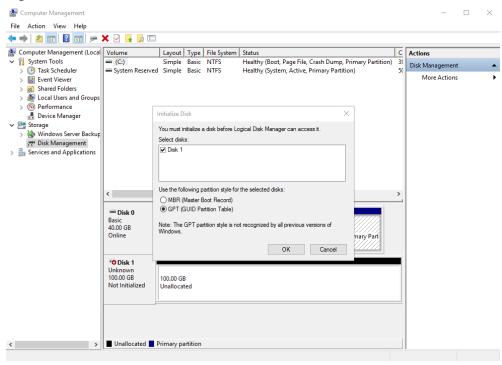
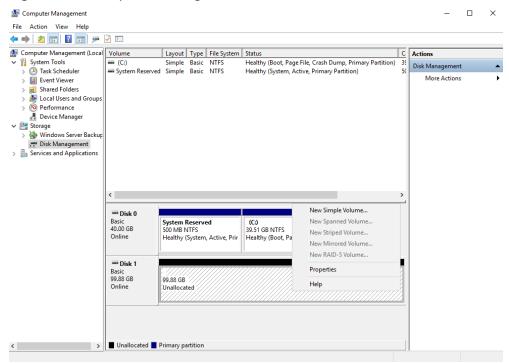


Figure 2-20 Disk list

Step 5 In the **Initialize Disk** dialog box, the to-be-initialized disk is selected. Select a disk partition style and click **OK**. In this example, **GPT (GUID Partition Table)** is selected.

The **Computer Management** window is displayed.

Figure 2-21 Computer Management



NOTICE

The maximum disk size supported by MBR is 2 TiB, and that supported by GPT is 18 EiB. Because an EVS data disk currently supports up to 32 TiB, use GPT if your disk size is larger than 2 TiB.

If the partition style is changed after the disk has been used, all data on the disk will be lost, so take care to select an appropriate partition style when initializing the disk. If you must change the partition style to GPT after a disk has been used, it is recommended that you back up the disk data before the change.

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

The **New Simple Volume Wizard** window is displayed.

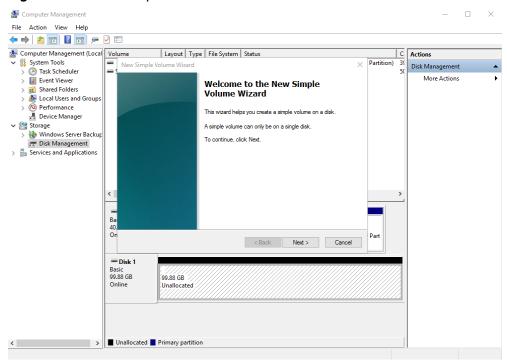
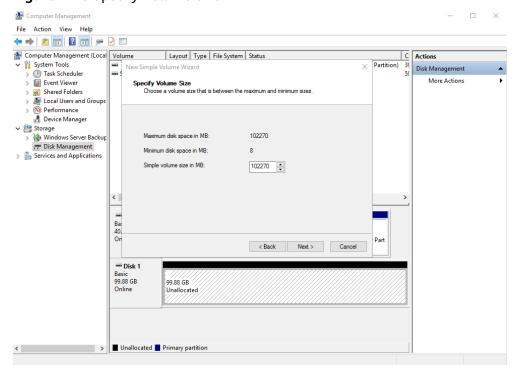


Figure 2-22 New Simple Volume Wizard

Step 7 Follow the prompts and click **Next**.

The **Specify Volume Size** page is displayed.

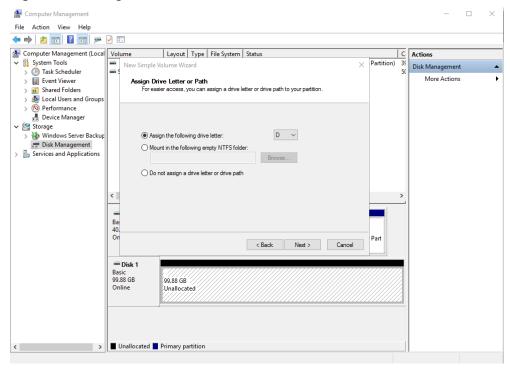
Figure 2-23 Specify Volume Size



Step 8 Specify the volume size and click **Next**. The system selects the maximum volume size by default. You can specify the volume size as required. In this example, the default setting is used.

The Assign Drive Letter or Path page is displayed.

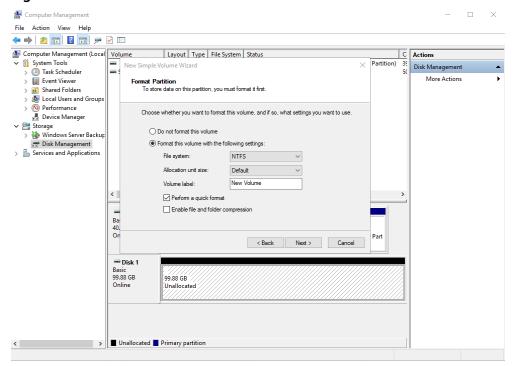
Figure 2-24 Assign Driver Letter or Path



Step 9 Assign a drive letter or path to your partition and click **Next**. The system assigns drive letter D by default. In this example, the default setting is used.

The Format Partition page is displayed.

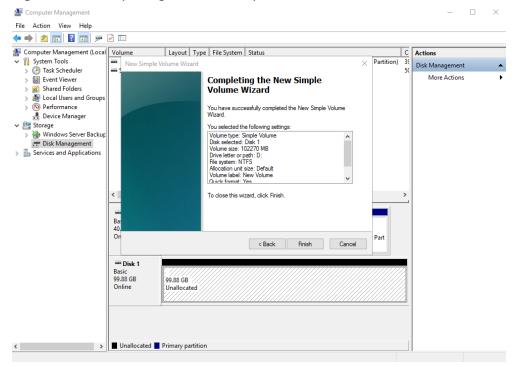
Figure 2-25 Format Partition



Step 10 Specify format settings and click **Next**. The system selects the NTFS file system by default. You can specify the file system type as required. In this example, the default setting is used.

The Completing the New Simple Volume Wizard page is displayed.

Figure 2-26 Completing the New Simple Volume Wizard



NOTICE

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

Step 11 Click Finish.

Wait for the initialization to complete. When the volume status changes to **Healthy**, the initialization has finished successfully, as shown in **Figure 2-27**.

2023-06-12

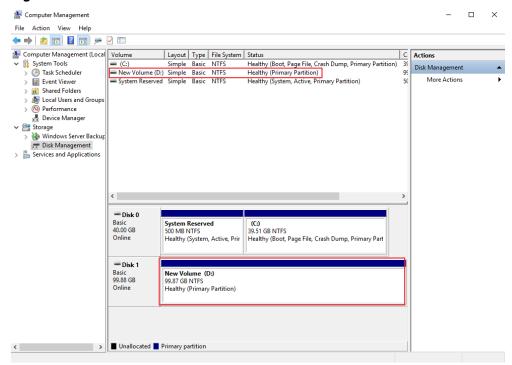
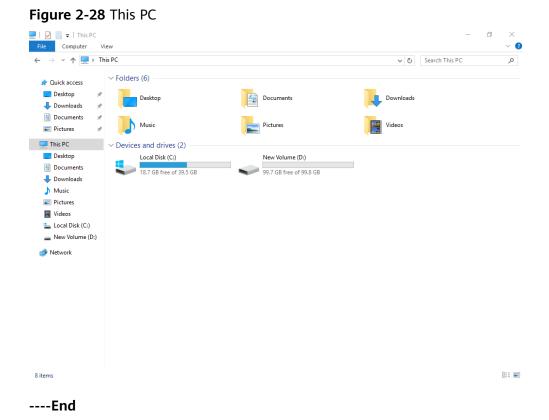


Figure 2-27 Disk initialized

Step 12 After the volume is created, click on the task bar and check whether a new volume appears in **This PC**. In this example, New Volume (D:) is the new volume.

If New Volume (D:) appears, the disk is successfully initialized and no further action is required.



2.4.4 Initializing a Linux Data Disk (fdisk)

Scenarios

This section uses CentOS 7.4 64bit to describe how to initialize a data disk attached to a server running Linux and use fdisk to partition the data disk.

The maximum partition size that MBR supports is 2 TiB and that GPT supports is 18 EiB. If the disk size you need to partition is greater than 2 TiB, partition the disk using GPT.

The fdisk partitioning tool is suitable only for MBR partitions, and the parted partitioning tool is suitable for both MBR and GPT partitions. For more information, see Introduction to Data Disk Initialization Scenarios and Partition Styles.

The method for initializing a disk varies slightly depending on the OS running on the server. This document is used for reference only. For the detailed operations and differences, see the product documents of the corresponding OS.

NOTICE

When using a disk for the first time, if you have not initialized it, including creating partitions and file systems, the additional space added to this disk in an expansion later may not be normally used.

Prerequisites

- A data disk has been attached to a server and has not been initialized.
- You have logged in to the server.
 - For how to log in to an ECS, see the Elastic Cloud Server User Guide.
 - For how to log in to a BMS, see the *Bare Metal Server User Guide*.

Creating and Mounting a Partition

The following example shows you how a new primary partition can be created on a new data disk that has been attached to a server. The primary partition will be created using fdisk, and MBR will be used. Furthermore, the partition will be formatted using the ext4 file system, mounted on /mnt/sdc, and configured to mount automatically at startup.

Step 1 Query what block devices are available on the server.

fdisk -l

Information similar to the following is displayed:

[root@ecs-test-0001 ~]# fdisk -l

Disk /dev/vda: 42.9 GiB, 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: 0x000bcb4e

Device Boot Start End Blocks Id System /dev/vda1 * 2048 83886079 41942016 83 Linux

Disk /dev/vdb: 107.4 GiB, 107374182400 bytes, 209715200 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

In the command output, this server contains two disks. /dev/vda and /dev/vdb. /dev/vda is the system disk, and /dev/vdb is the new data disk.

Step 2 Launch fdisk to partition the new data disk.

fdisk New data disk

In this example, run the following command:

fdisk /dev/vdb

Information similar to the following is displayed:

[root@ecs-test-0001 ~]# fdisk /dev/vdb 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 0x38717fc1.

Command (m for help):

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

Information similar to the following is displayed:

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

There are two types of disk partitions:

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

■ NOTE

If the MBR partition style is used, a maximum of 4 primary partitions, or 3 primary partitions and 1 extended partition can be created. The extended partition cannot be used directly and must be divided into logical partitions before use.

Disk partitions created using GPT are not categorized.

Step 4 Enter **p** and press **Enter** to create a primary partition in this example.

Information similar to the following is displayed:

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

Partition number indicates the serial number of the primary partition. The value ranges from 1 to 4.

Step 5 Enter the serial number of the primary partition and press **Enter**. Primary partition number **1** is used in this example. One usually starts with partition number **1** when partitioning an empty disk.

Information similar to the following is displayed:

```
Partition number (1-4, default 1): 1
First sector (2048-209715199, default 2048):
```

First sector indicates the start sector. The value ranges from **2048** to **209715199**, and the default value is **2048**.

Step 6 Select the default start sector **2048** and press **Enter**.

The system displays the start and end sectors of the partition's available space. You can customize the value within this range or use the default value. The start sector must be smaller than the partition's end sector.

Information similar to the following is displayed:

```
First sector (2048-209715199, default 2048):
Using default value 2048
Last sector, +sectors or +size{K,M,G} (2048-209715199, default 209715199):
```

Last sector indicates the end sector. The value ranges from **2048** to **209715199**, and the default value is **209715199**.

Step 7 Select the default end sector **209715199** and press **Enter**.

The system displays the start and end sectors of the partition's available space. You can customize the value within this range or use the default value. The start sector must be smaller than the partition's end sector.

Information similar to the following is displayed:

```
Last sector, +sectors or +size{K,M,G} (2048-209715199, default 209715199):
Using default value 209715199
Partition 1 of type Linux and of size 100 GiB is set

Command (m for help):
```

A primary partition has been created for the new data disk.

Step 8 Enter **p** and press **Enter** to print the partition details.

Information similar to the following is displayed:

Command (m for help): p

Disk /dev/vdb: 107.4 GiB, 107374182400 bytes, 209715200 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: 0x38717fc1

Device Boot Start End Blocks Id System /dev/vdb1 2048 209715199 104856576 83 Linux

Command (m for help):

Details about the /dev/vdb1 partition are displayed.

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

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.

Syncing disks.

The partition is created.

□ NOTE

In case that you want to discard the changes made before, you can exit fdisk by entering q.

Step 10 Synchronize the new partition table to the OS.

partprobe

Step 11 Format the new partition with a desired file system format.

mkfs -t File system format /dev/vdb1

In this example, the **ext4** format is used for the new partition.

mkfs -t ext4 /dev/vdb1

Information similar to the following is displayed:

[root@ecs-test-0001 ~]# mkfs -t ext4 /dev/vdb1

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

6553600 inodes, 26214144 blocks

1310707 blocks (5.00%) reserved for the super user

First data block=0

Maximum filesystem blocks=2174746624

800 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, 11239424, 20480000, 23887872

```
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.

NOTICE

The partition sizes supported by file systems vary. Choose an appropriate file system format based on your service requirements.

Step 12 Create a mount point.

mkdir Mount point

In this example, the /mnt/sdc mount point is created.

mkdir /mnt/sdc

□ NOTE

The /mnt directory exists on all Linux systems. If the mount point cannot be created, it may be that the /mnt directory has been accidentally deleted. You can run mkdir -p /mnt/sdc to create the mount point.

Step 13 Mount the new partition on the created mount point.

mount Disk partition Mount point

In this example, the /dev/vdb1 partition is mounted on /mnt/sdc.

mount /dev/vdb1 /mnt/sdc

Step 14 Check the mount result.

df -TH

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# df -TH
                  Size Used Avail Use% Mounted on
Filesystem
          Type
/dev/vda1
           ext4
                  43G 1.9G 39G 5%/
           devtmpfs 2.0G 0 2.0G 0% /dev
devtmpfs
          tmpfs 2.0G 0 2.0G 0% /dev/shm
tmpfs
tmpfs
          tmpfs
                 2.0G 9.1M 2.0G 1% /run
                 2.0G 0 2.0G 0% /sys/fs/cgroup
tmpfs
          tmpfs
                 398M 0 398M 0% /run/user/0
tmpfs
          tmpfs
/dev/vdb1 ext4 106G 63M 101G 1% /mnt/sdc
```

You should now see that partition /dev/vdb1 is mounted on /mnt/sdc.

□ NOTE

After the server is restarted, the disk will not be automatically mounted. To configure automount at startup, you will need to modify the /etc/fstab file. For details, see Configuring Automatic Mounting at System Start.

----End

Configuring Automatic Mounting at System Start

The **fstab** file controls what disks are automatically mounted at startup. You can use **fstab** to configure your data disks to mount automatically. This operation will not affect the existing data.

The example here uses UUIDs to identify disks in the **fstab** file. You are advised not to use device names to identify disks in the file because device names are assigned dynamically and may change (for example, from /dev/vdb1 to /dev/vdb2) after a reboot. This can even prevent your disk from booting up.

■ NOTE

UUID is the unique character string for disk partitions in a Linux system.

Step 1 Query the partition UUID.

blkid Disk partition

In this example, the UUID of the /dev/vdb1 partition is queried.

blkid /dev/vdb1

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# blkid /dev/vdb1 /dev/vdb1: UUID="0b3040e2-1367-4abb-841d-ddb0b92693df" TYPE="ext4"
```

Carefully record the UUID, as you will need it for the following step.

Step 2 Open the **fstab** file using the vi editor.

vi /etc/fstab

- **Step 3** Press i to enter editing mode.
- **Step 4** Move the cursor to the end of the file and press **Enter**. Then, add the following information:

```
UUID=0b3040e2-1367-4abb-841d-ddb0b92693df /mnt/sdc ext4 defaults 0
```

In this example, the line starting with "UUID" is the information added. Edit this line to match the following format:

- UUID: The UUID obtained in Step 1.
- Mount point: The directory on which the partition is mounted. You can query the mount point using df -TH.
- Filesystem: The file system format of the partition. You can query the file system format using **df -TH**.
- Mount option: The partition mount option. Usually, this parameter is set to defaults.
- Dump: The Linux dump backup option.
 - **0**: Linux dump backup is not used. Usually, dump backup is not used, and you can set this parameter to **0**.
 - 1: Linux dump backup is used.
- fsck: The fsck option, which means whether to use fsck to check the disk during startup.
 - 0: not use fsck.

- If the mount point is the root partition (/), this parameter must be set to **1**.

If this parameter is set to **1** for the root partition, this parameter for other partitions must start with **2** because the system checks the partitions in the ascending order of the values.

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

The system saves the configurations and exits the vi editor.

Step 6 Verify that the disk is auto-mounted at startup.

1. Unmount the partition.

umount Disk partition

In this example, run the following command:

umount /dev/vdb1

2. Reload all the content in the /etc/fstab file.

mount -a

3. Query the file system mounting information.

mount | grep Mount point

In this example, run the following command:

mount | grep /mnt/sdc

If information similar to the following is displayed, automatic mounting has been configured:

root@ecs-test-0001 ~]# mount | grep /mnt/sdc /dev/vdb1 on /mnt/sdc type ext4 (rw,relatime,data=ordered)

----End

2.4.5 Initializing a Linux Data Disk (parted)

Scenarios

This section uses CentOS 7.4 64bit to describe how to initialize a data disk attached to a server running Linux and use parted to partition the data disk.

The maximum partition size that MBR supports is 2 TiB and that GPT supports is 18 EiB. If the disk size you need to partition is greater than 2 TiB, partition the disk using GPT.

The fdisk partitioning tool is suitable only for MBR partitions, and the parted partitioning tool is suitable for both MBR and GPT partitions. For more information, see Introduction to Data Disk Initialization Scenarios and Partition Styles.

The method for initializing a disk varies slightly depending on the OS running on the server. This document is used for reference only. For the detailed operations and differences, see the product documents of the corresponding OS.

NOTICE

When using a disk for the first time, if you have not initialized it, including creating partitions and file systems, the additional space added to this disk in an expansion later may not be normally used.

Prerequisites

- A data disk has been attached to a server and has not been initialized.
- You have logged in to the server.
 - For how to log in to an ECS, see the *Elastic Cloud Server User Guide*.
 - For how to log in to a BMS, see the *Bare Metal Server User Guide*.

Creating and Mounting a Partition

The following example shows you how a new partition can be created on a new data disk that has been attached to a server. The partition will be created using parted, and GPT will be used. Furthermore, the partition will be formatted using the ext4 file system, mounted on /mnt/sdc, and configured to mount automatically at startup.

Step 1 Query information about the new data disk.

lsblk

Information similar to the following is displayed:

```
root@ecs-test-0001 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda 253:0 0 40G 0 disk
__vda1 253:1 0 40G 0 part /
vdb 253:16 0 100G 0 disk
```

In the command output, this server contains two disks. /dev/vda and /dev/vdb. /dev/vda is the system disk, and /dev/vdb is the new data disk.

Step 2 Launch parted to partition the new data disk.

```
parted New data disk
```

In this example, run the following command:

parted /dev/vdb

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# parted /dev/vdb
GNU Parted 3.1
Using /dev/vdb
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted)
```

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/vdb: unrecognised disk label Model: Virtio Block Device (virtblk) Disk /dev/vdb: 107GiB Sector size (logical/physical): 512B/512B

Partition Table: unknown

Disk Flags: (parted)

In the command output, the **Partition Table** value is **unknown**, indicating that no partition style is set for the new disk.

Step 4 Set the disk partition style.

mklabel Disk partition style

This command lets you control whether to use MBR or GPT for your partition table. In this example, GPT is used.

mklabel gpt

NOTICE

The maximum disk size supported by MBR is 2 TiB, and that supported by GPT is 18 EiB. Because an EVS data disk currently supports up to 32 TiB, use GPT if your disk size is larger than 2 TiB.

If the partition style is changed after the disk has been used, all data on the disk will be lost, so take care to select an appropriate partition style when initializing the disk. If you must change the partition style to GPT after a disk has been used, it is recommended that you back up the disk data before the change.

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: Virtio Block Device (virtblk)

Disk /dev/vdb: 107GiB

Sector size (logical/physical): 512B/512B

Partition Table: gpt

Disk Flags:

Number Start End Size File system Name Flags

(parted)

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

- **Step 6** Enter **unit s** and press **Enter** to set the measurement unit of the disk to sector.
- **Step 7** Create a new partition.

mkpart Partition name Start sector End sector

In this example, run the following command:

mkpart test 2048s 100%

In this example, one partition is created for the new data disk, starting on **2048** and using **100%** of the rest of the disk. The two values are used for reference only. You can determine the number of partitions and the partition size based on your service requirements.

Information similar to the following is displayed:

(parted) mkpart opt 2048s 100% (parted)

Step 8 Enter **p** and press **Enter** to print the partition details.

Information similar to the following is displayed:

```
(parted) p
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 209715200s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:
Number Start End
                        Size
                                 File system Name Flags
     2048s 209713151s 209711104s
                                              test
```

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

Information similar to the following is displayed:

```
(parted) a
Information: You may need to update /etc/fstab.
```

You can configure automatic mounting by updating the /etc/fstab file. Before doing so, format the partition with a desired file system and mount the partition on the mount point.

Step 10 View the disk partition information.

lsblk

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda 253:0 0 40G 0 disk
└─vda1 253:1 0 40G 0 part /
vdb 253:16 0 100G 0 disk
└vdb1 253:17 0 100G 0 part
```

In the command output, /dev/vdb1 is the partition you created.

Step 11 Format the new partition with a desired file system format.

```
mkfs -t File system format /dev/vdb1
```

In this example, the **ext4** format is used for the new partition.

mkfs -t ext4 /dev/vdb1

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# mkfs -t ext4 /dev/vdb1
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
6553600 inodes, 26213888 blocks
1310694 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2174746624
800 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, 11239424, 20480000, 23887872
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.

NOTICE

The partition sizes supported by file systems vary. Choose an appropriate file system format based on your service requirements.

Step 12 Create a mount point.

mkdir Mount point

In this example, the /mnt/sdc mount point is created.

mkdir /mnt/sdc

∩ NOTE

The /mnt directory exists on all Linux systems. If the mount point cannot be created, it may be that the /mnt directory has been accidentally deleted. You can run mkdir -p /mnt/sdc to create the mount point.

Step 13 Mount the new partition on the created mount point.

mount Disk partition Mount point

In this example, the /dev/vdb1 partition is mounted on /mnt/sdc.

mount /dev/vdb1 /mnt/sdc

Step 14 Check the mount result.

df -TH

Information similar to the following is displayed:

You should now see that partition /dev/vdb1 is mounted on /mnt/sdc.

After the server is restarted, the disk will not be automatically mounted. To configure automount at startup, you will need to modify the /etc/fstab file. For details, see Configuring Automatic Mounting at System Start.

----End

Configuring Automatic Mounting at System Start

The **fstab** file controls what disks are automatically mounted at startup. You can use **fstab** to configure your data disks to mount automatically. This operation will not affect the existing data.

The example here uses UUIDs to identify disks in the **fstab** file. You are advised not to use device names to identify disks in the file because device names are assigned dynamically and may change (for example, from /dev/vdb1 to /dev/vdb2) after a reboot. This can even prevent your disk from booting up.

■ NOTE

UUID is the unique character string for disk partitions in a Linux system.

Step 1 Query the partition UUID.

blkid Disk partition

In this example, the UUID of the /dev/vdb1 partition is queried.

blkid /dev/vdb1

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# blkid /dev/vdb1 /dev/vdb1: UUID="0b3040e2-1367-4abb-841d-ddb0b92693df" TYPE="ext4"
```

Carefully record the UUID, as you will need it for the following step.

Step 2 Open the fstab file using the vi editor.

vi /etc/fstab

- **Step 3** Press **i** to enter editing mode.
- **Step 4** Move the cursor to the end of the file and press **Enter**. Then, add the following information:

```
UUID=0b3040e2-1367-4abb-841d-ddb0b92693df /mnt/sdc ext4 defaults 0 2
```

In this example, the line starting with "UUID" is the information added. Edit this line to match the following format:

- UUID: The UUID obtained in Step 1.
- Mount point: The directory on which the partition is mounted. You can query the mount point using **df** -**TH**.
- Filesystem: The file system format of the partition. You can query the file system format using **df -TH**.
- Mount option: The partition mount option. Usually, this parameter is set to **defaults**.
- Dump: The Linux dump backup option.
 - **0**: Linux dump backup is not used. Usually, dump backup is not used, and you can set this parameter to **0**.
 - 1: Linux dump backup is used.
- fsck: The fsck option, which means whether to use fsck to check the disk during startup.
 - **0**: not use fsck.
 - If the mount point is the root partition (/), this parameter must be set to 1

If this parameter is set to **1** for the root partition, this parameter for other partitions must start with **2** because the system checks the partitions in the ascending order of the values.

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

The system saves the configurations and exits the vi editor.

Step 6 Verify that the disk is auto-mounted at startup.

1. Unmount the partition.

umount Disk partition

In this example, run the following command:

umount /dev/vdb1

2. Reload all the content in the /etc/fstab file.

mount -a

3. Query the file system mounting information.

mount | grep Mount point

In this example, run the following command:

mount | grep /mnt/sdc

If information similar to the following is displayed, automatic mounting has been configured:

root@ecs-test-0001 ~]# mount | grep /mnt/sdc /dev/vdb1 on /mnt/sdc type ext4 (rw,relatime,data=ordered)

----End

2.4.6 Initializing a Windows Data Disk Larger Than 2 TiB (Windows Server 2008)

Scenarios

This section uses Windows Server 2008 R2 Standard 64bit to describe how to initialize a data disk whose capacity is larger than 2 TiB. In the following operations, the capacity of the example disk is 3 TiB.

The maximum disk capacity supported by MBR is 2 TiB, and that supported by GPT is 18 EiB. Therefore, use the GPT partition style if your disk capacity is larger than 2 TiB. For details, see Initializing a Windows Data Disk Larger Than 2 TiB (Windows Server 2008). To learn more about disk partition styles, see Introduction to Data Disk Initialization Scenarios and Partition Styles.

The method for initializing a disk varies slightly depending on the OS running on the server. This document is used for reference only. For the detailed operations and differences, see the product documents of the corresponding OS.

NOTICE

When using a disk for the first time, if you have not initialized it, including creating partitions and file systems, the additional space added to this disk in an expansion later may not be normally used.

Prerequisites

- A data disk has been attached to a server and has not been initialized.
- You have logged in to the server.

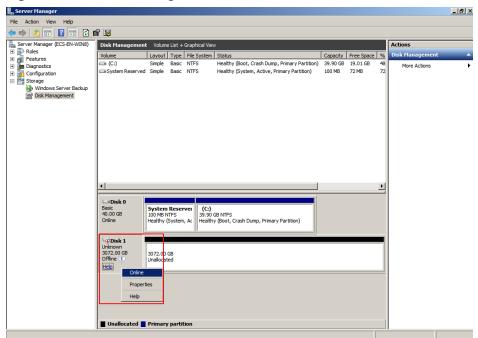
- For how to log in to an ECS, see the *Elastic Cloud Server User Guide*.
- For how to log in to a BMS, see the *Bare Metal Server User Guide*.

Procedure

- **Step 1** On the desktop of the server, click **Start**.
 - The **Start** window is displayed.
- **Step 2** Right-click **Computer** and choose **Manage** from the short-cut menu.

The **Server Manager** window is displayed.

Figure 2-29 Server Manager (Windows Server 2008)



Step 3 Disks are listed in the right pane. If the new disk is in the offline state, bring it online before initializing it.

In the **Disk 1** area, right-click and choose **Online** from the shortcut menu.

When the status of Disk 1 changes from **Offline** to **Not Initialized**, the disk has been brought online.

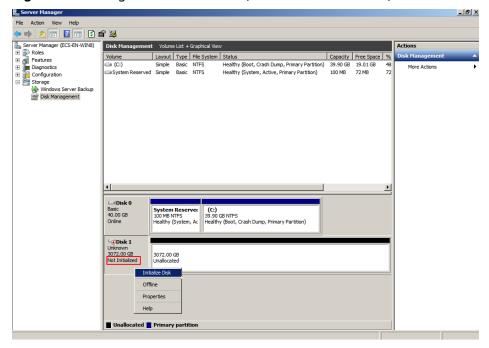
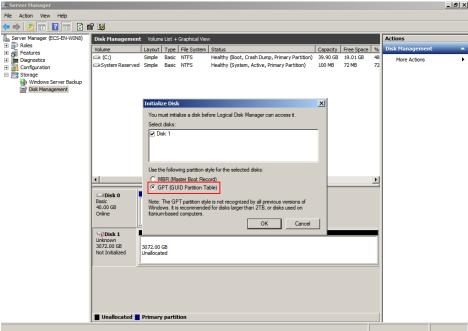


Figure 2-30 Bring online succeeded (Windows Server 2008)

Step 4 In the **Disk 1** area, right-click and choose **Initialize Disk** from the shortcut menu.

The Initialize Disk dialog box is displayed.





Step 5 In the **Initialize Disk** dialog box, the to-be-initialized disk is selected. In this example, the disk capacity is larger than 2 TiB. Therefore, select **GPT (GUID Partition Table)** and click **OK**.

The **Server Manager** window is displayed.

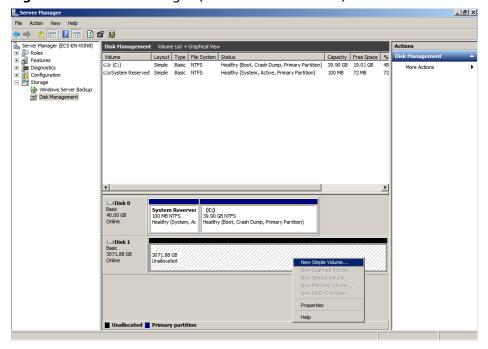


Figure 2-32 Server Manager (Windows Server 2008)

NOTICE

The maximum disk size supported by MBR is 2 TiB, and that supported by GPT is 18 EiB. Because an EVS data disk currently supports up to 32 TiB, use GPT if your disk size is larger than 2 TiB.

If the partition style is changed after the disk has been used, all data on the disk will be lost, so take care to select an appropriate partition style when initializing the disk. If you must change the partition style to GPT after a disk has been used, it is recommended that you back up the disk data before the change.

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

The New Simple Volume Wizard window is displayed.

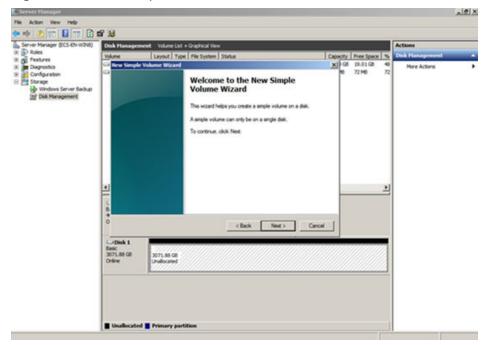


Figure 2-33 New Simple Volume Wizard (Windows Server 2008)

Step 7 Follow the prompts and click Next.

The **Specify Volume Size** page is displayed.

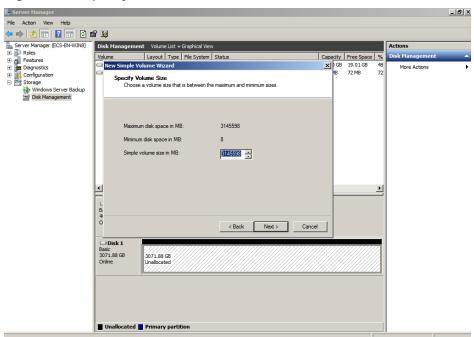


Figure 2-34 Specify Volume Size (Windows Server 2008)

Step 8 Specify the volume size and click **Next**. The system selects the maximum volume size by default. You can specify the volume size as required. In this example, the default setting is used.

The Assign Drive Letter or Path page is displayed.

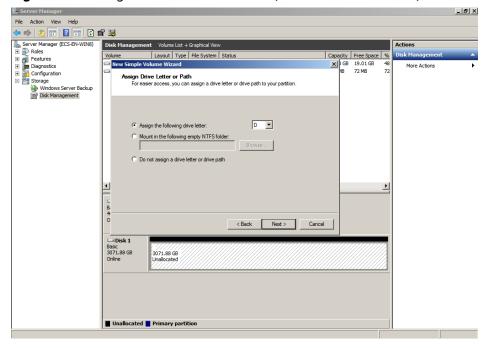
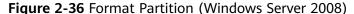
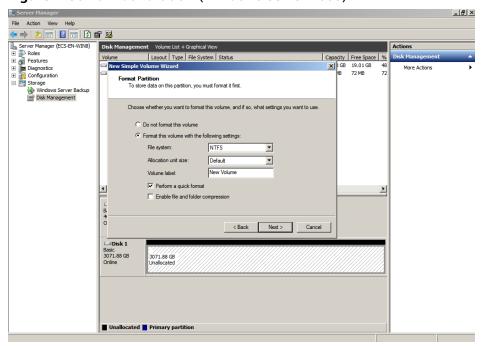


Figure 2-35 Assign Driver Letter or Path (Windows Server 2008)

Step 9 Assign a drive letter or path to your partition and click **Next**. The system assigns drive letter D by default. In this example, the default setting is used.

The Format Partition page is displayed.





Step 10 Specify format settings and click **Next**. The system selects the NTFS file system by default. You can specify the file system type as required. In this example, the default setting is used.

The Completing the New Simple Volume Wizard page is displayed.



Figure 2-37 Completing the New Simple Volume Wizard

NOTICE

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

Step 11 Click Finish.

Wait for the initialization to complete. When the volume status changes to **Healthy**, the initialization has finished successfully, as shown in **Figure 2-38**.

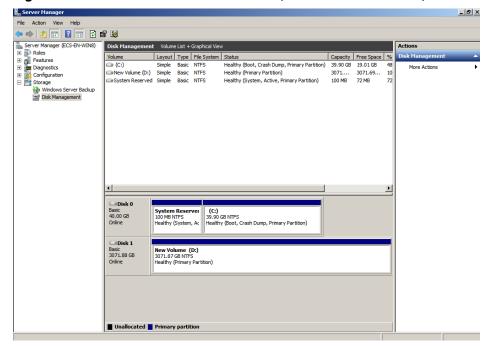


Figure 2-38 Disk initialization succeeded (Windows Server 2008)

Step 12 After the volume is created, click and check whether a new volume appears in **Computer**. In this example, New Volume (D:) is the new volume.

If New Volume (D:) appears, the disk is successfully initialized and no further action is required.

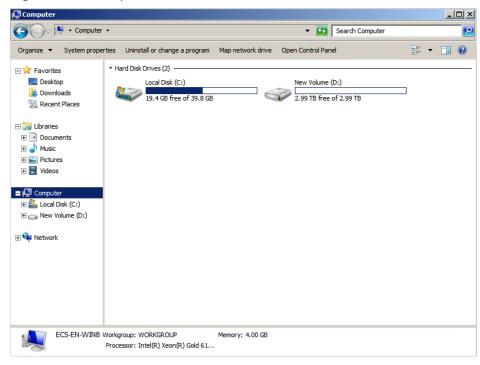


Figure 2-39 Computer (Windows Server 2008)

----End

2.4.7 Initializing a Windows Data Disk Larger Than 2 TiB (Windows Server 2012)

Scenarios

This section uses Windows Server 2012 R2 Standard 64bit to describe how to initialize a data disk whose capacity is larger than 2 TiB. In the following operations, the capacity of the example disk is 3 TiB.

The maximum disk capacity supported by MBR is 2 TiB, and that supported by GPT is 18 EiB. Therefore, use the GPT partition style if your disk capacity is larger than 2 TiB. For details, see Initializing a Windows Data Disk Larger Than 2 TiB (Windows Server 2008). To learn more about disk partition styles, see Introduction to Data Disk Initialization Scenarios and Partition Styles.

The method for initializing a disk varies slightly depending on the OS running on the server. This document is used for reference only. For the detailed operations and differences, see the product documents of the corresponding OS.

NOTICE

When using a disk for the first time, if you have not initialized it, including creating partitions and file systems, the additional space added to this disk in an expansion later may not be normally used.

Prerequisites

- A data disk has been attached to a server and has not been initialized.
- You have logged in to the server.
 - For how to log in to an ECS, see the *Elastic Cloud Server User Guide*.
 - For how to log in to a BMS, see the Bare Metal Server User Guide.

Procedure

Step 1 On the desktop of the server, click in the lower area.

The **Server Manager** window is displayed.

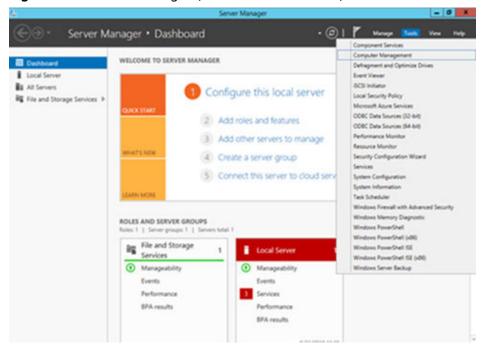
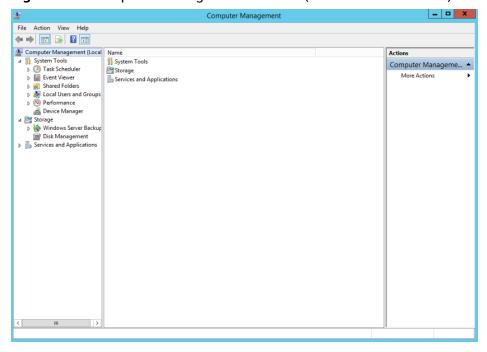


Figure 2-40 Server Manager (Windows Server 2012)

Step 2 In the upper right corner, choose **Tools** > **Computer Management**.

The **Computer Management** window is displayed.

Figure 2-41 Computer Management window (Windows Server 2012)



Step 3 Choose Storage > Disk Management.

Disks are displayed in the right pane.

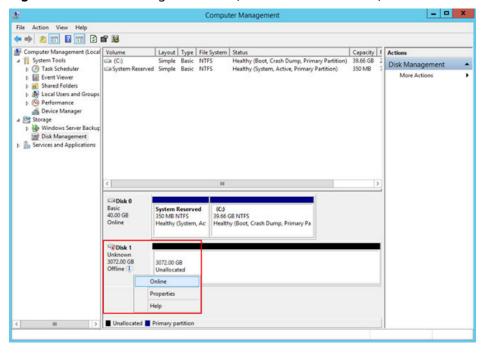


Figure 2-42 Disk Management list (Windows Server 2012)

Step 4 (Optional) If the new disk is in the offline state, bring it online before initializing it.

In the **Disk 1** area, right-click and choose **Online** from the shortcut menu.

When the status of Disk 1 changes from **Offline** to **Not Initialized**, the disk has been brought online.

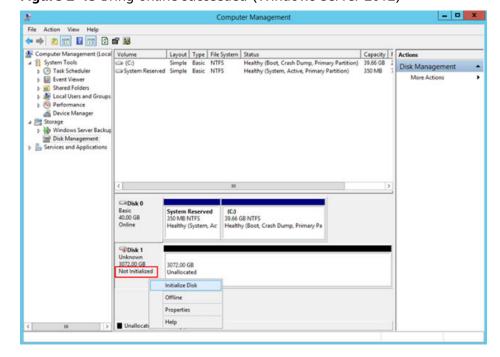
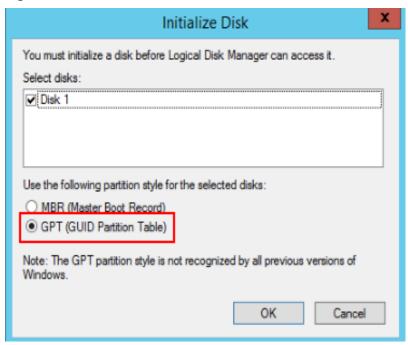


Figure 2-43 Bring online succeeded (Windows Server 2012)

Step 5 In the Disk 1 area, right-click and choose Initialize Disk from the shortcut menu.

The Initialize Disk dialog box is displayed.

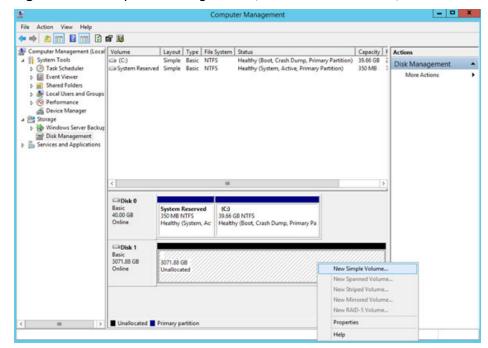
Figure 2-44 Initialize Disk (Windows Server 2012)



Step 6 In the **Initialize Disk** dialog box, the to-be-initialized disk is selected. In this example, the disk capacity is larger than 2 TiB. Therefore, select **GPT (GUID Partition Table)** and click **OK**.

The **Computer Management** window is displayed.

Figure 2-45 Computer Management (Windows Server 2012)



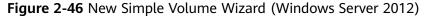
NOTICE

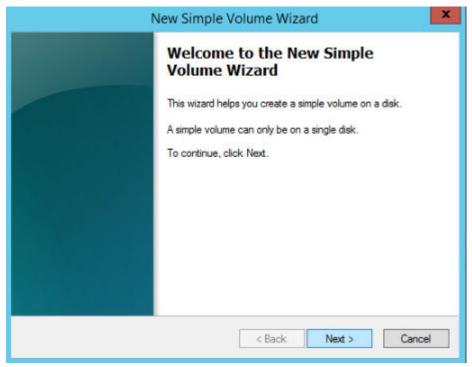
The maximum disk size supported by MBR is 2 TiB, and that supported by GPT is 18 EiB. Because an EVS data disk currently supports up to 32 TiB, use GPT if your disk size is larger than 2 TiB.

If the partition style is changed after the disk has been used, all data on the disk will be lost, so take care to select an appropriate partition style when initializing the disk. If you must change the partition style to GPT after a disk has been used, it is recommended that you back up the disk data before the change.

Step 7 Right-click at the unallocated disk space and choose **New Simple Volume** from the shortcut menu.

The New Simple Volume Wizard window is displayed.





Step 8 Follow the prompts and click **Next**.

The **Specify Volume Size** page is displayed.

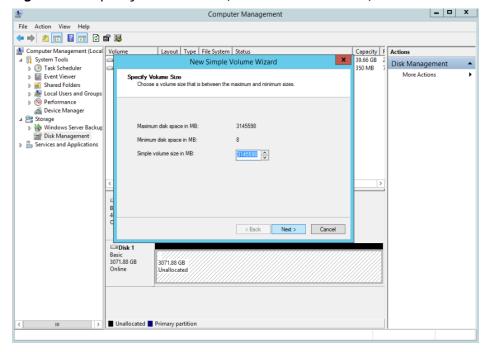
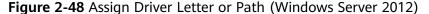
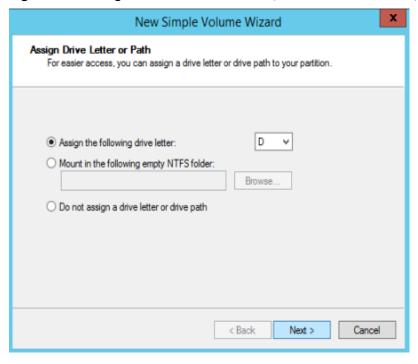


Figure 2-47 Specify Volume Size (Windows Server 2012)

Step 9 Specify the volume size and click **Next**. The system selects the maximum volume size by default. You can specify the volume size as required. In this example, the default setting is used.

The **Assign Drive Letter or Path** page is displayed.

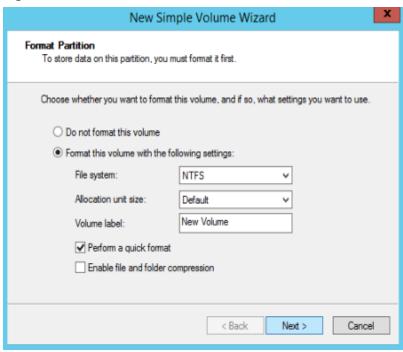




Step 10 Assign a drive letter or path to your partition and click **Next**. The system assigns drive letter D by default. In this example, the default setting is used.

The Format Partition page is displayed.

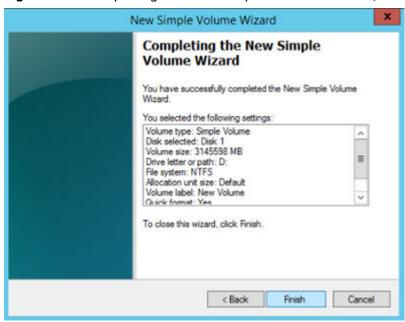
Figure 2-49 Format Partition (Windows Server 2012)



Step 11 Specify format settings and click **Next**. The system selects the NTFS file system by default. You can specify the file system type as required. In this example, the default setting is used.

The Completing the New Simple Volume Wizard page is displayed.

Figure 2-50 Completing the New Simple Volume Wizard (Windows Server 2012)



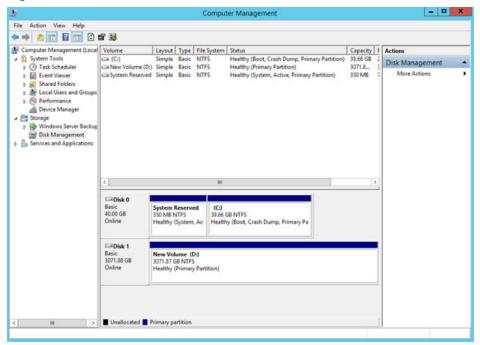
NOTICE

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

Step 12 Click Finish.

Wait for the initialization to complete. When the volume status changes to **Healthy**, the initialization has finished successfully, as shown in **Figure 2-51**.

Figure 2-51 Disk initialization succeeded (Windows Server 2012)



Step 13 After the volume is created, click and check whether a new volume appears in This PC. In this example, New Volume (D:) is the new volume.

If New Volume (D:) appears, the disk is successfully initialized and no further action is required.

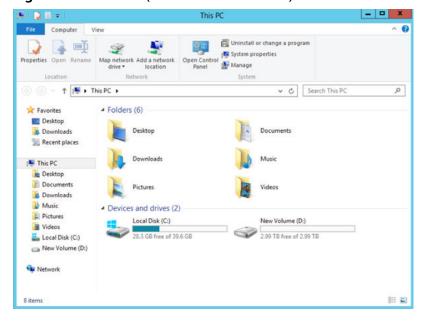


Figure 2-52 This PC (Windows Server 2012)

----End

2.4.8 Initializing a Linux Data Disk Larger Than 2 TiB (parted)

Scenarios

This section uses CentOS 7.4 64bit to describe how to use parted to initialize a data disk whose capacity is larger than 2 TiB. In the following operations, the capacity of the example disk is 3 TiB.

The maximum partition size that MBR supports is 2 TiB and that GPT supports is 18 EiB. If the disk size you need to partition is greater than 2 TiB, partition the disk using GPT.

The fdisk partitioning tool is suitable only for MBR partitions, and the parted partitioning tool is suitable for both MBR and GPT partitions. For more information, see Introduction to Data Disk Initialization Scenarios and Partition Styles.

The method for initializing a disk varies slightly depending on the OS running on the server. This document is used for reference only. For the detailed operations and differences, see the product documents of the corresponding OS.

NOTICE

When using a disk for the first time, if you have not initialized it, including creating partitions and file systems, the additional space added to this disk in an expansion later may not be normally used.

Prerequisites

A data disk has been attached to a server and has not been initialized.

- You have logged in to the server.
 - For how to log in to an ECS, see the Elastic Cloud Server User Guide.
 - For how to log in to a BMS, see the Bare Metal Server User Guide.

Creating and Mounting a Partition

The following example shows you how a new partition can be created on a new data disk that has been attached to a server. The partition will be created using parted, and GPT will be used. Furthermore, the partition will be formatted using the ext4 file system, mounted on /mnt/sdc, and configured to mount automatically at startup.

Step 1 Query information about the new data disk.

lsblk

Information similar to the following is displayed:

```
[root@ecs-centos74 ~]# lsblk

NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT

vda 253:0 0 40G 0 disk

—vda1 253:1 0 1G 0 part /boot

—vda2 253:2 0 39G 0 part /

vdb 253:16 0 3T 0 disk
```

In the command output, this server contains two disks. /dev/vda and /dev/vdb. /dev/vda is the system disk, and /dev/vdb is the new data disk.

Step 2 Launch parted to partition the new data disk.

```
parted New data disk
```

In this example, run the following command:

parted /dev/vdb

Information similar to the following is displayed:

```
[root@ecs-centos74 ~]# parted /dev/vdb
GNU Parted 3.1
Using /dev/vdb
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted)
```

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/vdb: unrecognised disk label
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 3299GiB
Sector size (logical/physical): 512B/512B
Partition Table: unknown
Disk Flags:
(parted)
```

In the command output, the **Partition Table** value is **unknown**, indicating that no partition style is set for the new disk.

Step 4 Set the disk partition style.

mklabel Disk partition style

The disk partition style can be MBR or GPT. If the disk capacity is greater than 2 TiB, use GPT.

mklabel gpt

NOTICE

The maximum disk size supported by MBR is 2 TiB, and that supported by GPT is 18 EiB. Because an EVS data disk currently supports up to 32 TiB, use GPT if your disk size is larger than 2 TiB.

If the partition style is changed after the disk has been used, all data on the disk will be lost, so take care to select an appropriate partition style when initializing the disk. If you must change the partition style to GPT after a disk has been used, it is recommended that you back up the disk data before the change.

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: Virtio Block Device (virtblk) Disk /dev/vdb: 3299GiB

Sector size (logical/physical): 512B/512B

Partition Table: gpt Disk Flags:

Number Start End Size File system Name Flags

(parted)

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

Step 7 Create a new partition.

mkpart Partition name Start sector End sector

In this example, run the following command:

mkpart opt 2048s 100%

In this example, one partition is created for the new data disk, starting on **2048** and using **100%** of the rest of the disk. The two values are used for reference only. You can determine the number of partitions and the partition size based on your 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? Ignore

If the preceding warning message is displayed, enter **Ignore** to ignore the performance warning.

Step 8 Enter **p** and press **Enter** to print the partition details.

Information similar to the following is displayed:

(parted) p Model: Virtio Block Device (virtblk) Disk /dev/vdb: 6442450944s

```
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number Start End Size File system Name Flags
1 2048s 6442448895s 6442446848s opt
```

Details about the dev/vdb1 partition are displayed.

- **Step 9** Enter **q** and press **Enter** to exit parted.
- **Step 10** View the disk partition information.

lsblk

Information similar to the following is displayed:

```
[root@ecs-centos74 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda 253:0 0 40G 0 disk
—vda1 253:1 0 1G 0 part /boot
—vda2 253:2 0 39G 0 part /
vdb 253:16 0 3T 0 disk
—vdb1 253:17 0 3T 0 part
```

In the command output, /dev/vdb1 is the partition you created.

Step 11 Format the new partition with a desired file system format.

```
mkfs -t File system format /dev/vdb1
```

In this example, the ext4 format is used for the new partition.

mkfs -t ext4 /dev/vdb1

Information similar to the following is displayed:

```
[root@ecs-centos74 ~]# mkfs -t ext4 /dev/vdb1
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
201326592 inodes, 805305856 blocks
40265292 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2952790016
24576 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, 11239424, 20480000, 23887872, 71663616, 78675968,
     102400000, 214990848, 512000000, 550731776, 644972544
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.

NOTICE

The partition sizes supported by file systems vary. Choose an appropriate file system format based on your service requirements.

Step 12 Create a mount point.

mkdir Mount point

In this example, the /mnt/sdc mount point is created.

mkdir /mnt/sdc

The /mnt directory exists on all Linux systems. If the mount point cannot be created, it may be that the /mnt directory has been accidentally deleted. You can run mkdir -p /mnt/sdc to create the mount point.

Step 13 Mount the new partition on the created mount point.

mount Disk partition Mount point

In this example, the /dev/vdb1 partition is mounted on /mnt/sdc.

mount /dev/vdb1 /mnt/sdc

Step 14 Check the mount result.

df -TH

Information similar to the following is displayed:

```
[root@ecs-centos74 ~]# df -TH
            Type Size Used Avail Use% Mounted on
Filesystem
             ext4
                     42G 1.5G 38G 4% /
/dev/vda2
             devtmpfs 2.0G 0 2.0G 0% /dev
devtmpfs
            tmpfs 2.0G 0 2.0G 0% /dev/shm
tmpfs
           tmpfs 2.0G 8.9M 2.0G 1% /run
tmpfs 2.0G 0 2.0G 0% /sys/fs/cgro
ext4 1.1G 153M 801M 17% /boot
tmpfs
                    2.0G  0  2.0G  0% /sys/fs/cgroup
tmpfs
/dev/vda1
                    398M 0 398M 0% /run/user/0
tmpfs
            tmpfs
/dev/vdb1
           ext4 3.3T 93M 3.1T 1% /mnt/sdc
```

You should now see that partition /dev/vdb1 is mounted on /mnt/sdc.

----End

Configuring Automatic Mounting at System Start

The **fstab** file controls what disks are automatically mounted at startup. You can use **fstab** to configure your data disks to mount automatically. This operation will not affect the existing data.

The example here uses UUIDs to identify disks in the **fstab** file. You are advised not to use device names to identify disks in the file because device names are assigned dynamically and may change (for example, from /dev/vdb1 to /dev/vdb2) after a reboot. This can even prevent your disk from booting up.

■ NOTE

UUID is the unique character string for disk partitions in a Linux system.

Step 1 Query the partition UUID.

blkid Disk partition

In this example, the UUID of the /dev/vdb1 partition is queried.

blkid /dev/vdb1

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# blkid /dev/vdb1 /dev/vdb1: UUID="0b3040e2-1367-4abb-841d-ddb0b92693df" TYPE="ext4"
```

Carefully record the UUID, as you will need it for the following step.

Step 2 Open the fstab file using the vi editor.

vi /etc/fstab

- **Step 3** Press **i** to enter editing mode.
- **Step 4** Move the cursor to the end of the file and press **Enter**. Then, add the following information:

```
UUID=0b3040e2-1367-4abb-841d-ddb0b92693df /mnt/sdc ext4 defaults 0 2
```

In this example, the line starting with "UUID" is the information added. Edit this line to match the following format:

- UUID: The UUID obtained in Step 1.
- Mount point: The directory on which the partition is mounted. You can query the mount point using **df** -**TH**.
- Filesystem: The file system format of the partition. You can query the file system format using **df -TH**.
- Mount option: The partition mount option. Usually, this parameter is set to defaults.
- Dump: The Linux dump backup option.
 - **0**: Linux dump backup is not used. Usually, dump backup is not used, and you can set this parameter to **0**.
 - 1: Linux dump backup is used.
- fsck: The fsck option, which means whether to use fsck to check the disk during startup.
 - 0: not use fsck.
 - If the mount point is the root partition (/), this parameter must be set to 1.

If this parameter is set to **1** for the root partition, this parameter for other partitions must start with **2** because the system checks the partitions in the ascending order of the values.

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

The system saves the configurations and exits the vi editor.

Step 6 Verify that the disk is auto-mounted at startup.

1. Unmount the partition.

umount Disk partition

In this example, run the following command:

umount /dev/vdb1

2. Reload all the content in the /etc/fstab file.

mount -a

3. Query the file system mounting information.

mount | grep Mount point

In this example, run the following command:

mount | grep /mnt/sdc

If information similar to the following is displayed, automatic mounting has been configured:

root@ecs-test-0001 ~]# mount | grep /mnt/sdc /dev/vdb1 on /mnt/sdc type ext4 (rw,relatime,data=ordered)

----End

3 Permissions Management

3.1 Creating a User and Granting EVS Permissions

This chapter describes how to use IAM to implement fine-grained permissions control for your EVS resources. With IAM, you can:

- Create IAM users for employees based on your enterprise's organizational structure. Each IAM user will have their own security credentials for accessing EVS resources.
- Grant only the permissions required for users to perform a specific task.
- Entrust an account or cloud service to perform professional and efficient O&M on your EVS resources.

If your account does not require individual IAM users, skip this chapter.

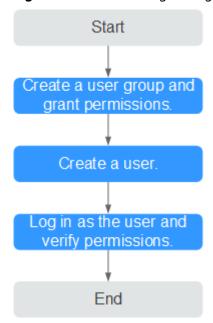
This section describes the procedure for granting permissions (see Figure 3-1).

Prerequisites

Learn about the permissions (see section "Permissions Management" in the *Elastic Volume Service User Guide*) supported by EVS and choose policies or roles according to your requirements.

Process Flow

Figure 3-1 Process for granting EVS permissions



- Create a user group and assign permissions to it.
 - Create a user group on the IAM console, and attach the **EVS ReadOnlyAccess** policy to the group.
- 2. Create a user and add it to a user group.
 - Create a user on the IAM console and add the user to the group created in 1.
- 3. Log in and verify permissions.
 - Log in to the EVS console using the created user and verify that the user only has read permissions for EVS.
 - In Service List, choose Elastic Volume Service. On the EVS console, click Create Disk in the upper right corner. If a message appears indicating that you have insufficient permissions to perform the operation, the EVS ReadOnlyAccess policy has already taken effect.
 - Choose any other service in Service List. If a message appears indicating that you have insufficient permissions to access the service, the EVS ReadOnlyAccess policy has already taken effect.

3.2 EVS Custom Policies

Custom policies can be created to supplement the system-defined policies of EVS. For the actions supported for custom policies, see section "Permissions Policies and Supported Actions" in the *Elastic Volume Service API Reference*.

You can create custom policies in either of the following ways:

• Visual editor: Select cloud services, actions, resources, and request conditions. This does not require knowledge of policy syntax.

• JSON: Edit JSON policies from scratch or based on an existing policy.

For operation details, see section "Creating a Custom Policy" in the *Identity*and Access Management User Guide. The following section contains examples of common EVS custom policies.

Example Custom Policies

• Example 1: Allowing users to create disks.

```
"Version": "1.1",
"Statement": [
     {
           "Action": [
                 "evs:volumes:list",
                 "evs:volumes:get",
                 "evs:quotas:get",
                 "evs:volumeTags:list",
                 "evs:types:get",
                 "evs:volumes:create",
                 "ecs:cloudServerFlavors:get",
                 "ecs:cloudServers:list".
                 "bss:balance:view",
                 "bss:order:pay".
                 "bss:order:update"
           ],
"Effect": "Allow"
     }
]
```

• Example 2: Denying disk deletion

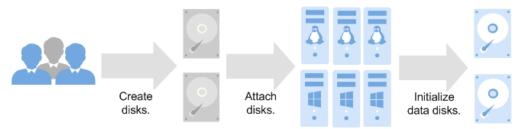
A policy with only "Deny" permissions must be used in conjunction with other policies to take effect. If the permissions assigned to a user contain both "Allow" and "Deny", the "Deny" permissions take precedence over the "Allow" permissions.

The following method can be used if you need to assign permissions of the **EVS FullAccess** policy to a user but you want to prevent the user from deleting EVS disks. Create a custom policy for denying disk deletion, and attach both policies to the group to which the user belongs. Then, the user can perform all operations on disks except deleting disks. The following is an example of a deny policy:

4 Creating and Using an EVS Disk

Figure 4-1 shows the basic EVS operation process.

Figure 4-1 Process overview



EVS disks can be attached to servers and be used as system disks and data disks. For details, see **Table 4-1**.

Table 4-1 Method of creation

Disk	Description	Method
System disk	System disks are created along with servers and cannot be created separately.	 See section "Creating an ECS" in the Elastic Cloud Server User Guide. See section "Creating a BMS" in the Bare Metal Server User Guide.
Data disk	Data disks can be created along with servers or separately.	 See section "Creating an ECS" in the Elastic Cloud Server User Guide. See section "Creating a BMS" in the Bare Metal Server User Guide. Create an EVS Disk

Figure 4-2 shows how to purchase a data disk separately.

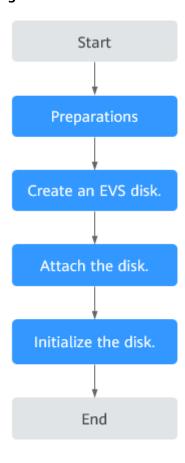


Figure 4-2 Process overview

- 1. **Preparations**: Register an account on the console and obtain the permissions required to create servers and disks.
- 2. **Create a disk.** Configure the disk parameters, including the disk type, size, name, and other information. For details, see **Create an EVS Disk**.
- 3. **Attach the data disk.** Attach the separately created disk to a server. For details, see the following sections:
 - Attaching a Non-Shared Disk
 - Attaching a Shared Disk
- 4. **Initialize the data disk.** Log in to the server and initialize the data disk before using it. For details about how to initialize the disk, see the following sections:
 - Introduction to Data Disk Initialization Scenarios and Partition Styles
 - Windows
 - Initializing a Windows Data Disk (Windows Server 2008)
 - Initializing a Windows Data Disk (Windows Server 2019)
 - Initializing a Windows Data Disk Larger Than 2 TiB (Windows Server 2008)
 - Initializing a Windows Data Disk Larger Than 2 TiB (Windows Server 2012)

- Linux
 - Initializing a Linux Data Disk (fdisk)
 - Initializing a Linux Data Disk (parted)
 - Initializing a Linux Data Disk Larger Than 2 TiB (parted)

5 Disk Capacity Expansion

5.1 Expansion Overview

What Is Capacity Expansion?

If the capacity of an existing disk is insufficient, you can expand the disk capacity to increase the storage space.

Both system disks and data disks can be expanded. Disk capacity can only be expanded, not reduced.

How to Expand the Disk Capacity?

A capacity expansion operation includes the following steps:

- 1. Expand the disk capacity on the management console.
- 2. Log in to the server and extend the disk partition and file system.

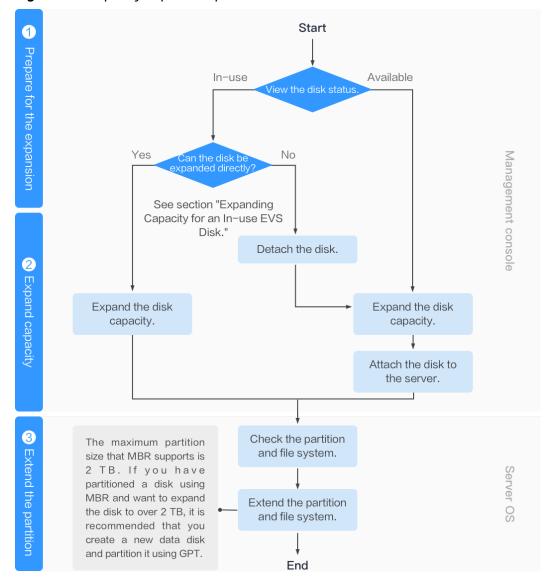


Figure 5-1 Capacity expansion procedure

Expand the Disk Capacity on the Management Console

Choose a proper expansion method based on the disk status.

For an In-use disk:

The disk has been attached to a server. Check whether the disk can be expanded in the In-use state by referring to **Constraints**.

- If yes, expand the disk capacity according to Expanding Capacity for an In-use EVS Disk.
- If no, detach the disk. Then, expand the disk capacity according to Expanding Capacity for an Available EVS Disk.
- For an Available disk:

The disk has not been attached to any server and can be directly expanded by referring to **Expanding Capacity for an Available EVS Disk**.

A shared disk can be expanded only when its status is **Available**.

Log In to the Server and Extend the Disk Partition and File System

After the disk has been expanded on the management console, only the disk storage capacity is enlarged, but its additional space cannot be used directly. You must log in to the server and extend the disk partition and file system. For details, see **Table 5-1**.

Table 5-1 Extending the disk partition and file system

Capacity After Expansion	Extend Disk Partition and File System
Disk capacity ≤ 2 TiB	Windows: Extending Disk Partitions and File Systems (Windows Server 2008)
	• Linux: Partition and File System Extension Preparations (Linux)
Disk capacity > 2 TiB	 GPT partition style: Extending Disk Partitions and File Systems (Windows Server 2008) or Partition and File System Extension Preparations (Linux) MBR partition style: Not supported The maximum disk capacity that MBR supports is 2 TiB, and the disk space exceeding 2 TiB cannot be used. If your disk uses MBR and you need to expand the disk capacity to over 2 TiB, change the partition style from MBR to GPT. Ensure that the disk data has been backed up before changing the partition style because services will be interrupted
	and data on the disk will be deleted during this change.

■ NOTE

If the server is stopped during the expansion, the additional space of a Windows system disk, Windows data disk, and Linux system disk may be automatically added to the end partition after the server is started. In this case, the additional space can be directly used. If the additional space is not automatically added, extend the partition and file system according to the preceding steps.

Related Operations

For more information, see **FAQs** > **Capacity Expansion**.

5.2 Expanding Capacity for an In-use EVS Disk

Scenarios

This section describes how to expand the capacity of an In-use EVS disk on the management console. The In-use status indicates that the disk has been attached to a server. You do not need to detach the disk when expanding an In-use disk.

Constraints

- Disk capacity can be expanded, but cannot be reduced.
- When expanding an In-use disk, the server attached with this disk must be in the Running or Stopped state.
- A shared disk in the In-use state cannot be expanded. To expand such a disk, you must detach it from all its servers, wait until its status changes to
 Available, and then expand its capacity. For more information, see Expanding
 Capacity for an Available EVS Disk.
- Only some server OSs support capacity expansion of In-use disks. For servers
 without such support, detach the disk and then expand its capacity.
 Otherwise, you may need to stop and then start the server after the
 expansion to make the additional space available.
 - Perform the following operations to check whether your server OS supports capacity expansion of In-use disks:
 - a. Check your server image. Both public images and private images listed on the console support the capacity expansion of In-use disks.
 - How to view: Log in to the management console. In the navigation pane
 - on the left, click and choose **Compute** > **Image Management Service**. On the **Public Images** tab, view the images of the **ECS image** type.
 - b. If your server OS does not appear in the image list, check whether it is included in **Table 5-2**.
 - If yes, you can expand the disk. Otherwise, you must detach the disk and then expand its capacity. For details, see **Expanding Capacity for an Available EVS Disk**.

Table 5-2 OSs that support the capacity expansion of In-use disks

OS	Version
CentOS 8	8.0 64-bit or later
CentOS 7	7.2 64-bit or later
CentOS 6	6.5 64-bit or later
Debian	8.5.0 64-bit or later
Fedora	24 64-bit or later

os	Version
SUSE 12	SUSE Linux Enterprise Server 12 64-bit or later
SUSE 11	SUSE Linux Enterprise Server 11 SP4 64-bit
OpenSUSE	42.1 64-bit or later
Oracle Linux Server release 7	7.2 64-bit or later
Oracle Linux Server release 6	6.7 64-bit or later
Ubuntu Server	14.04 64-bit or later
Red Hat Enterprise Linux 7	7.3 64-bit
Red Hat Enterprise Linux 6	6.8 64-bit
EulerOS	2.2 64-bit or later
Windows Server 2016	Windows Server 2016 R2 Enterprise 64- bit
Windows Server 2012	Windows Server 2012 R2 Standard 64- bit
Windows Server 2008	Windows Server 2008 R2 Enterprise 64- bit

Prerequisites

Disk data has been backed up using CBR or snapshots. For details about backups, see **Managing EVS Backups**. For details about snapshots, see **Creating a Snapshot**.

Procedure

- **Step 1** Log in to the management console.
- **Step 2** Under **Storage**, click **Elastic Volume Service**.

The disk list page is displayed.

- **Step 3** Choose a way to expand the disk by determining whether you want to check server information first.
 - If yes, perform the following procedure:
 - a. In the disk list, click the name of the to-be-expanded disk. The disk details page is displayed.
 - b. Click the **Servers** tab to view the server where the target disk has been attached.
 - c. Click **Expand Capacity** in the upper right corner of the page. The expansion page is displayed.

- If no, perform the following procedure:
 - a. In the disk list, locate the row that contains the target disk and click **Expand Capacity** in the **Operation** column.

The expansion page is displayed.

- **Step 4** Set the **Add Capacity (GiB)** parameter and click **Next**.
- **Step 5** On the **Details** page, check the disk details.
 - Click **Submit** to start the expansion.
 - Click **Previous** to change the settings.

After the configuration is submitted, go back to the disk list page.

Step 6 In the disk list, view the capacity of the target disk.

When the disk status changes from **Expanding** to **In-use** and the disk capacity increases, the expansion has succeeded.

■ NOTE

When the status of a disk is **Expanding**, you are not allowed to modify the specifications of the ECS where the disk is attached.

Step 7 Log in to the server and extend the partition and file system after the disk has been expanded on the console, because previous steps only enlarge the disk space.

The operations vary depending on the server OS.

- In Windows, see Extending Disk Partitions and File Systems (Windows Server 2008).
- In Linux, see Partition and File System Extension Preparations (Linux).

----End

5.3 Expanding Capacity for an Available EVS Disk

Scenarios

This section describes how to expand the capacity of an Available EVS disk on the management console. The Available status indicates that the disk has not been attached to any server.

Constraints

- Disk capacity can be expanded, but cannot be reduced.
- A shared disk in the In-use state cannot be expanded. To expand such a disk, you must detach it from all its servers, wait until its status changes to
 Available, and then expand its capacity.

Prerequisites

Disk data has been backed up using CBR or snapshots. For details about backups, see **Managing EVS Backups**. For details about snapshots, see **Creating a Snapshot**.

Procedure

- **Step 1** Log in to the management console.
- Step 2 Under Storage, click Elastic Volume Service.

The disk list page is displayed.

Step 3 In the disk list, locate the row that contains the target disk and click **Expand Capacity** in the **Operation** column.

The expansion page is displayed.

- **Step 4** Set the **Add Capacity (GiB)** parameter and click **Next**.
- **Step 5** On the **Details** page, check the disk details.
 - Click **Submit** to start the expansion.
 - Click **Previous** to change the settings.

After the configuration is submitted, go back to the disk list page.

Step 6 In the disk list, view the capacity of the target disk.

When the disk status changes from **Expanding** to **Available** and the disk capacity increases, the expansion has succeeded.

□ NOTE

When the status of a disk is **Expanding**, you are not allowed to modify the specifications of the ECS where the disk is attached.

- **Step 7** Attach the disk to the server. For details, see **Attach an EVS Disk**.
- **Step 8** Log in to the server and extend the partition and file system after the disk has been expanded on the console, because previous steps only enlarge the disk space.

The operations vary depending on the server OS.

- In Windows, see Extending Disk Partitions and File Systems (Windows Server 2008).
- In Linux, see Partition and File System Extension Preparations (Linux).

----End

5.4 Extending Disk Partitions and File Systems (Windows Server 2008)

Scenarios

After a disk is expanded on the management console, the disk size is enlarged, but the additional space cannot be used directly.

In Windows, you must allocate the additional space to an existing partition or a new partition.

If the disk capacity is expanded on a stopped server, the additional space of a Windows system disk or Windows data disk will be automatically added to the

partition at the end of the disk upon the server startup. In this case, the additional space can be used directly.

This section uses Windows Server 2008 R2 Enterprise 64bit as the sample OS to describe the expansion methods:

- For a system disk:
 - If volume (C:) already exists, you can add the additional space to volume
 (C:) and use it as a system volume. For details, see System Disk: Add Additional Space to Volume (C:).
 - If volume (C:) already exists, you can create a new volume such as volume (F:) with the additional space and use the new volume as a data volume. For details, see System Disk: Create New Volume (F:) with Additional Space.
 - If the additional space has been added to volume (C:), you can shrink volume (C:), create a new volume with the available space, and use the new volume as a data volume. Only the available space can be shrunk and used to create the new volume. The additional space cannot be shrunk if it has already been used. This section uses a system disk to describe how to perform extension operations for a Windows disk. These operations are also suitable for data disks. For details, see System Disk: Add Additional Space to Volume (C:) and Allocate Available Space to New Volume (D:) via the Shrink Function.
- For a data disk:
 - If volume (D:) already exists, you can add the additional space to volume (D:) and use it as a data volume. For details, see Data Disk: Add Additional Space to Volume (D:).
 - If volume (D:) already exists, you can create a new volume such as volume (E:) with the additional space and use the new volume as a data volume. For details, see Data Disk: Create New Volume (E:) with Additional Space.

The method for allocating the additional space varies with the server OS. This section is used for reference only. For detailed operations and differences, see the corresponding OS documents.

NOTICE

Performing the expansion operations with caution. Incorrect operations may lead to data loss or exceptions. So you are advised to back up the disk data using CBR or snapshots before expansion. For details about using CBR, see **Managing EVS Backups**. For details about using snapshots, see **Creating a Snapshot**.

Prerequisites

- You have expanded the disk capacity and attached the disk to a server on the management console. For details, see Expanding Capacity for an In-use EVS Disk or Expanding Capacity for an Available EVS Disk.
- You have logged in to the server.
 - For how to log in to an ECS, see the Elastic Cloud Server User Guide.

For how to log in to a BMS, see the Bare Metal Server User Guide.

System Disk: Add Additional Space to Volume (C:)

In this example, the system disk has 50 GiB originally, and 22 GiB is added on the management console. The following procedure describes how to add this 22 GiB to volume (C:) on the server. After the operation is complete, volume (C:) will have 72 GiB of capacity and can be used as a system volume.

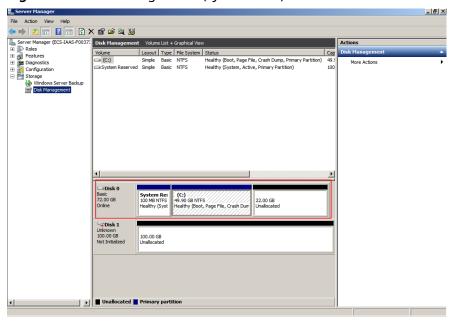
Step 1 On the desktop of the server, right-click **Computer** and choose **Manage** from the shortcut menu.

The **Server Manager** window is displayed.

Step 2 In the navigation tree, choose **Storage** > **Disk Management**.

The **Disk Management** window is displayed.

Figure 5-2 Disk Management (system disk)



If you cannot view the additional space, right-click **Disk Management** and choose **Refresh** from the shortcut menu.

- **Step 3** On the **Disk Management** page, select the disk and volume that you want to extend. The current volume size and unallocated space are displayed.
- **Step 4** Right-click the target volume and choose **Extend Volume**.

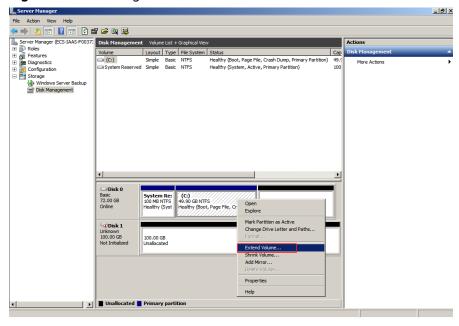


Figure 5-3 Choosing Extend Volume

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

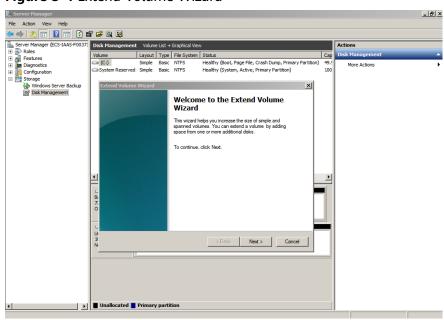


Figure 5-4 Extend Volume Wizard

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

File Action Vew Help

Server Manager (ECS-IAAS-F0037)

Disk Planagement Volume List + Craphical Vew

Volume Layout Type File System Status

Cop Features

Cop Fine System Reserved Simple Basic NTPS Healthy (System, Active, Primary Partition)

More Actions

Prince Planagement

Volume Layout Type File System Status

Cop Cop File System Reserved Simple Basic NTPS Healthy (System, Active, Primary Partition)

Actions

Disk Planagement

Volume Layout Type File System Status

Cop Cop Simple Basic NTPS Healthy (System, Active, Primary Partition)

More Actions

Actions

More Actions

Actions

Disk Planagement

Volume List + Craphical Vew

Volume Layout Type File System Status

Cop Cop Simple Basic NTPS Healthy (System, Active, Primary Partition)

Actions

Act

Figure 5-5 Selecting space

Step 7 Click Finish.

After the expansion succeeded, the partition size is larger than the original size.

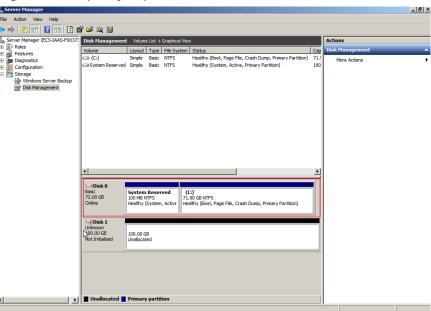


Figure 5-6 Capacity expansion succeeded

----End

System Disk: Create New Volume (F:) with Additional Space

In this example, the system disk has 40 GiB originally, and 60 GiB is added on the management console. The following procedure describes how to use this 60 GiB to create a new volume, for example volume (F:), on the server. After the operation is complete, new volume (F:) has 60 GiB of capacity and can be used as a data volume.

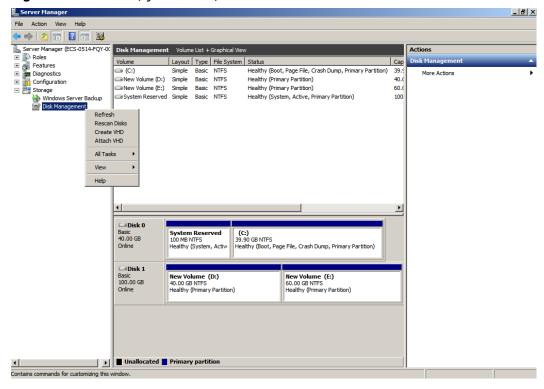
Step 1 On the desktop of the server, right-click **Computer** and choose **Manage** from the shortcut menu.

The **Server Manager** window is displayed.

Step 2 In the navigation tree, choose **Storage** > **Disk Management**.

The **Disk Management** window is displayed.

Figure 5-7 Refresh (system disk)



Step 3 If you cannot view the additional space, right-click **Disk Management** and choose **Refresh** from the shortcut menu.

After the refresh, the additional space is displayed in the right area and is unallocated.

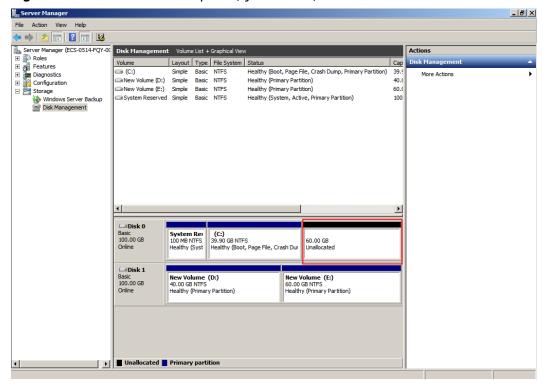


Figure 5-8 Unallocated disk space (system disk)

Step 4 In the **Unallocated** area of **Disk 0**, right-click the blank area and choose **New Simple Volume**.

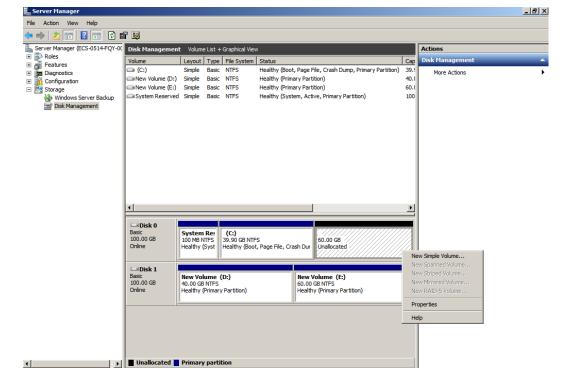


Figure 5-9 New Simple Volume (system disk)

Step 5 On the displayed New Simple Volume Wizard window, click Next.

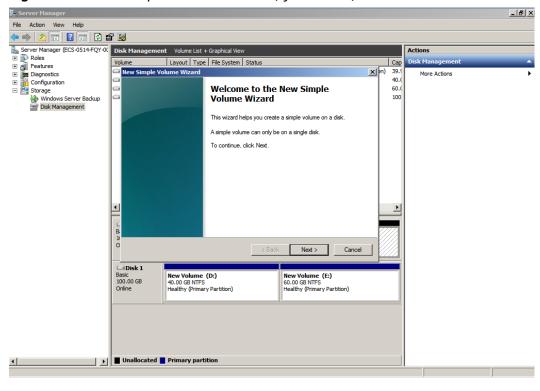
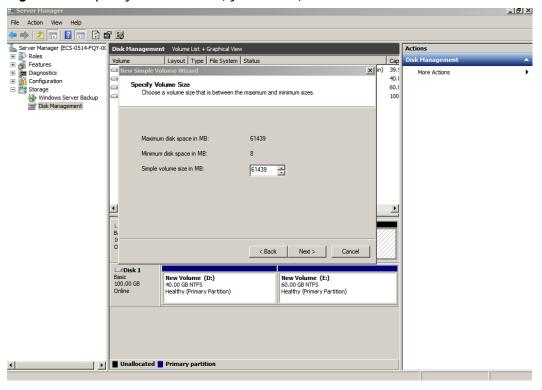


Figure 5-10 New Simple Volume Wizard (system disk)

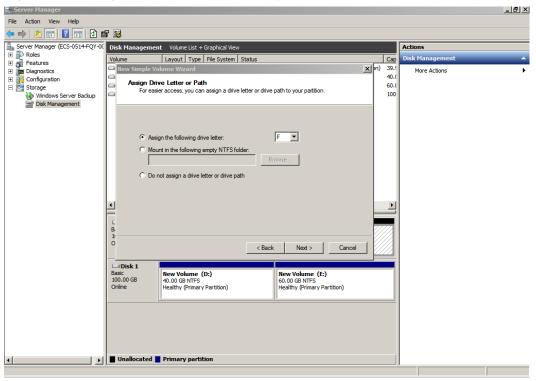
Step 6 On the displayed **Specify Volume Size** page, set **Simple volume size in MB** and click **Next**. In this example, the default size is used.





Step 7 On the displayed **Assign Drive Letter and Path** page, click **Assign the following drive letter**, select a drive letter, and click **Next**. In this example, drive letter **F** is selected.





Step 8 On the displayed **Format Partition** page, click **Format this volume with the following settings**, set parameters based on the requirements, and select **Perform a quick format**. Then, click **Next**.

_ B ×

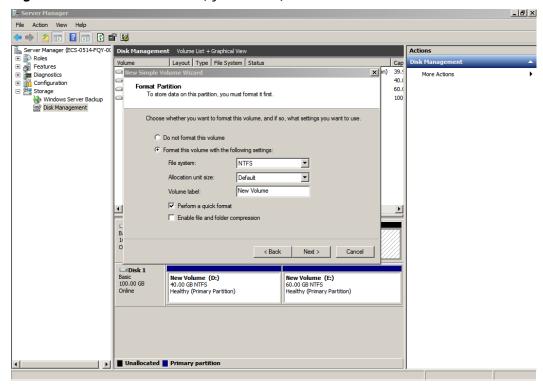


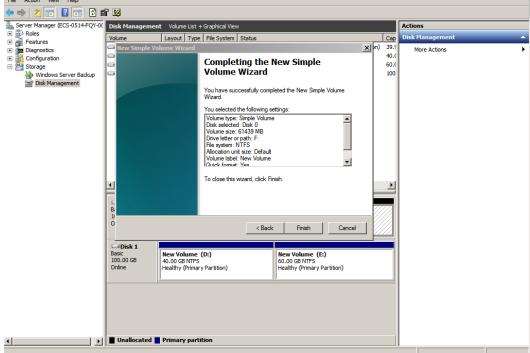
Figure 5-13 Format Partition (system disk)

Step 9 Click Finish.

After the expansion succeeded, new volume (F:) is displayed.

Server Manager File Action View Help Actions

Figure 5-14 Completing the New Simple Volume Wizard (new volume F:)



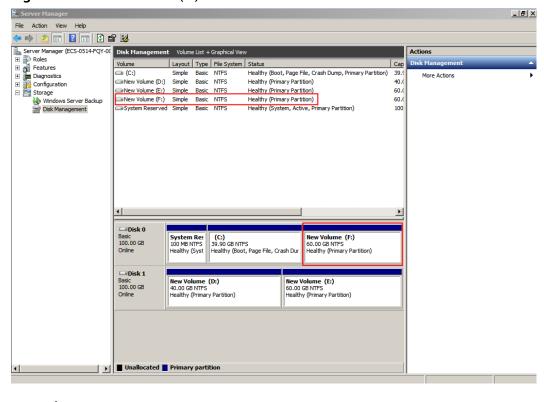


Figure 5-15 New Volume (F:)

----End

System Disk: Add Additional Space to Volume (C:) and Allocate Available Space to New Volume (D:) via the Shrink Function

In this example, the system disk has 40 GiB originally, and 60 GiB is added on the management console and then formatted and added to volume (C:). This 60 GiB has not been used.

The following procedure describes how to use the shrink function to create new volume (D:) with this 60 GiB. After the operation is complete, new volume (D:) can be used as a data volume.

Step 1 On the desktop of the server, 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** window is displayed.

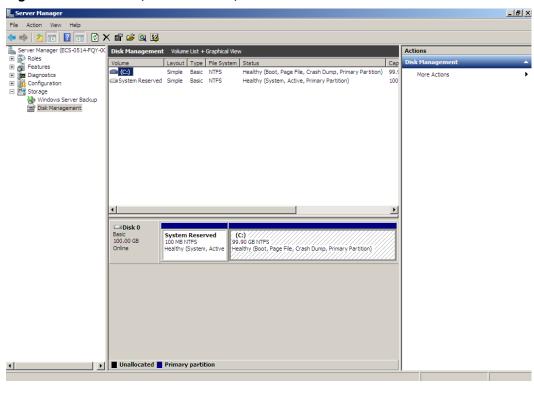


Figure 5-16 Refresh (shrink volume)

Step 3 In the (C:) area of Disk 0, right-click the blank area and choose Shrink Volume.

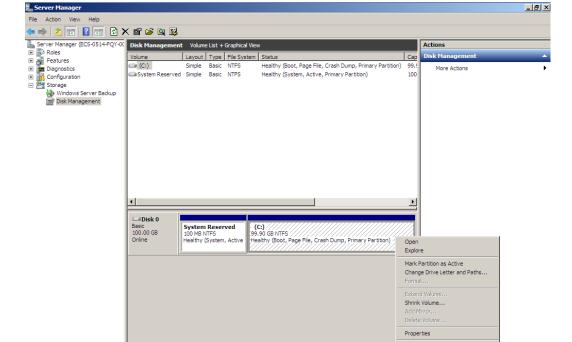


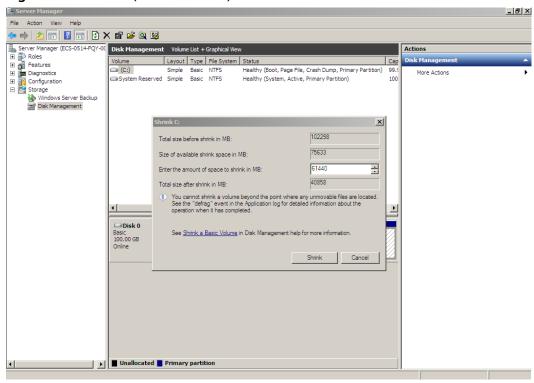
Figure 5-17 Shrink Volume

Step 4 The system automatically queries the available shrink space. In the displayed dialog box, enter the available space and click **Shrink**.

■ Unallocated ■ Primary partition

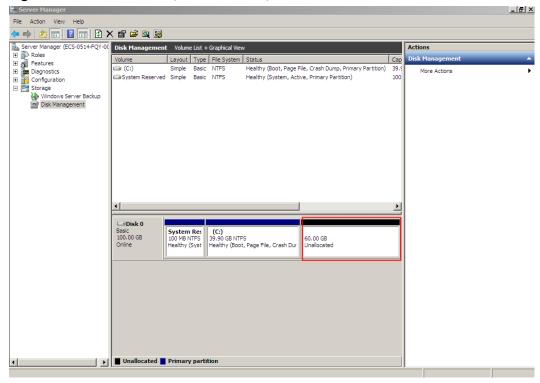
In this example, the volume available space is 60 GiB. Therefore, enter 61440 (60 \times 1024 MiB).

Figure 5-18 Shrink (shrink volume)



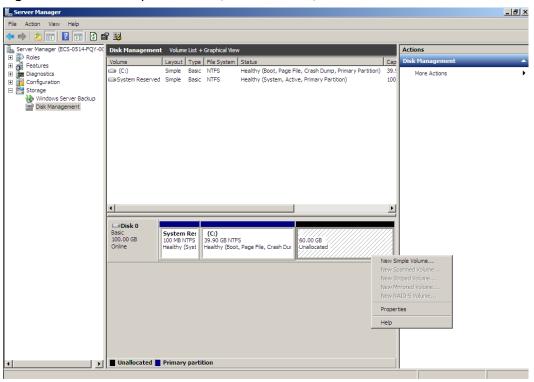
After the operation is complete, **Disk 0** has 60 GiB unallocated space.

Figure 5-19 Unallocated (shrink volume)



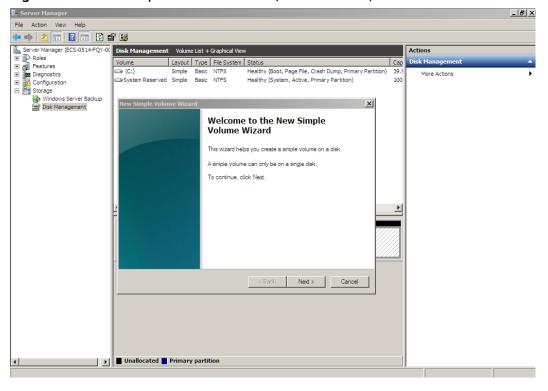
Step 5 In the **Unallocated** area of **Disk 0**, right-click the blank area and choose **New Simple Volume**.

Figure 5-20 New Simple Volume (shrink volume)



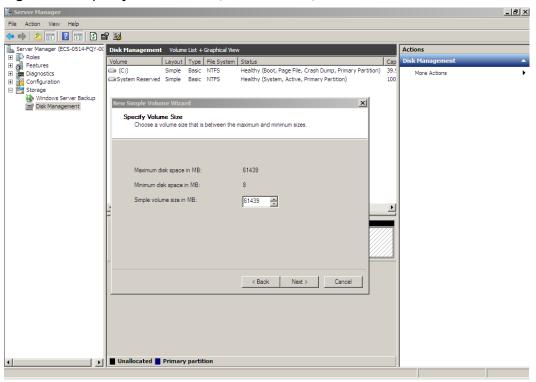
Step 6 On the displayed New Simple Volume Wizard window, click Next.

Figure 5-21 New Simple Volume Wizard (shrink volume)



Step 7 On the displayed **Specify Volume Size** page, set **Simple volume size in MB** and click **Next**. In this example, the default size is used.





Step 8 On the displayed **Assign Drive Letter and Path** page, click **Assign the following drive letter**, select a drive letter, and click **Next**. In this example, drive letter **D** is selected.

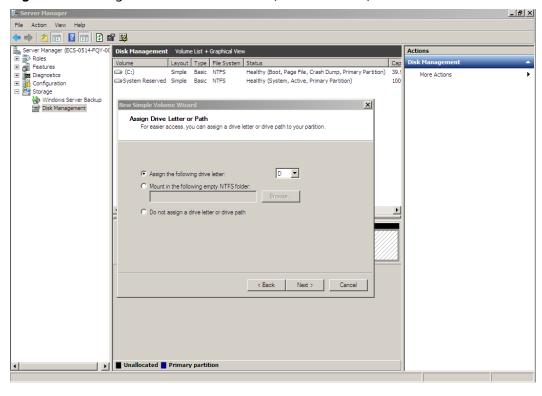
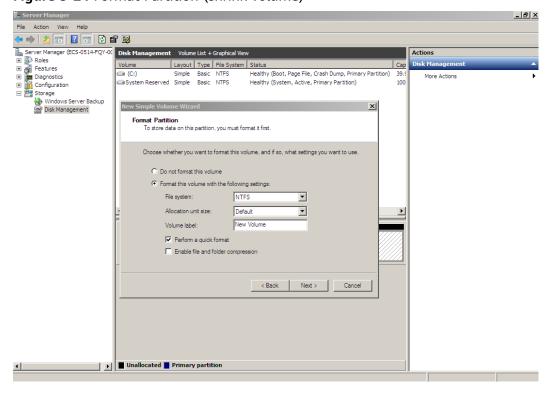


Figure 5-23 Assign Driver Letter or Path (shrink volume)

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

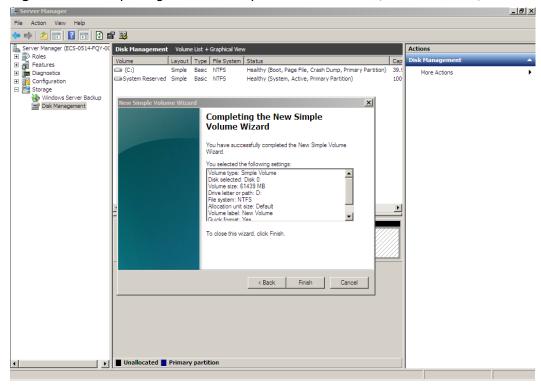




Step 10 Click Finish.

After the expansion succeeded, new volume (D:) is displayed.

Figure 5-25 Completing the New Simple Volume Wizard (new volume D:)



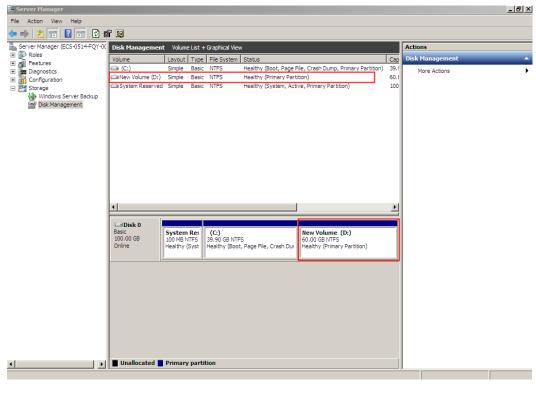


Figure 5-26 New Volume (D:)

----End

Data Disk: Add Additional Space to Volume (D:)

In this example, the data disk has 100 GiB originally, and 50 GiB is added on the management console. The following procedure describes how to add this 50 GiB to volume (D:) on the server. After the operation is complete, volume (D:) has 150 GiB of capacity and can be used as a data volume.

Step 1 On the desktop of the server, 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** window is displayed.

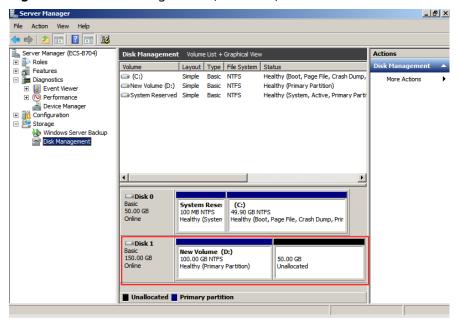


Figure 5-27 Disk Management (data disk)

□ NOTE

If you cannot view the additional space, right-click **Disk Management** and choose **Refresh** from the shortcut menu.

- **Step 3** On the **Disk Management** page, select the disk and volume that you want to extend. The current volume size and unallocated space are displayed.
- **Step 4** Right-click the target volume and choose **Extend Volume**.

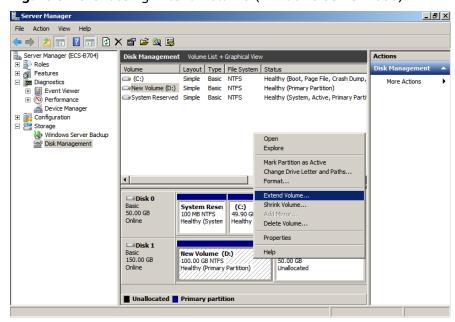


Figure 5-28 Choosing Extend Volume (Windows Server 2008)

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



Figure 5-29 Extend Volume Wizard (Windows Server 2008)

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

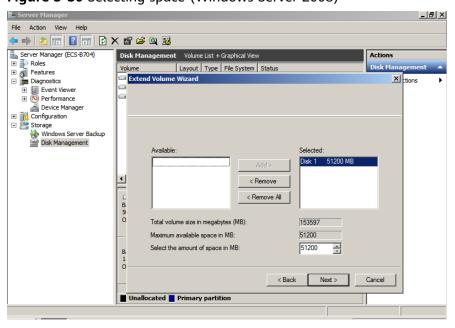


Figure 5-30 Selecting space (Windows Server 2008)

Step 7 Click Finish.

After the expansion succeeded, the partition size is larger than the original size.

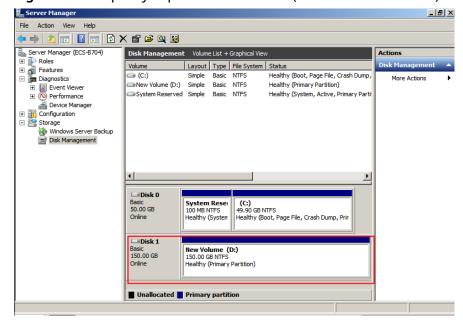


Figure 5-31 Capacity expansion succeeded (Windows Server 2008)

----End

Data Disk: Create New Volume (E:) with Additional Space

In this example, the data disk has 40 GiB originally, and 60 GiB is added on the management console. The following procedure describes how to use this 60 GiB to create a new volume, for example volume (E:), on the server. After the operation is complete, new volume (E:) has 60 GiB of capacity and can be used as a data volume.

Step 1 On the desktop of the server, 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** window is displayed.

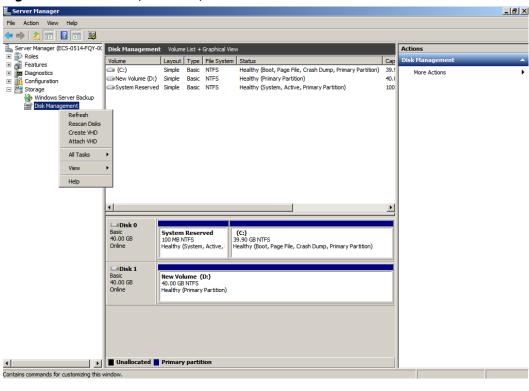


Figure 5-32 Refresh (data disk)

Step 3 If you cannot view the additional space, right-click **Disk Management** and choose **Refresh** from the shortcut menu.

After the refresh, the additional space is displayed in the right area and is unallocated.

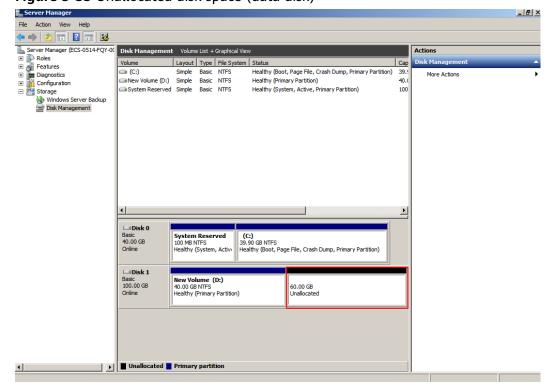
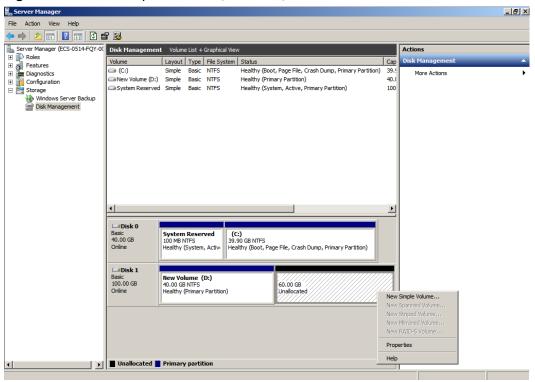


Figure 5-33 Unallocated disk space (data disk)

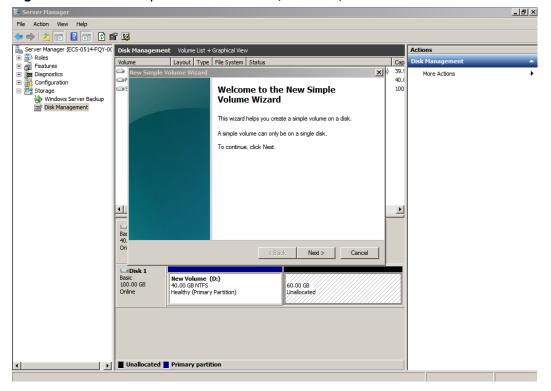
Step 4 In the **Unallocated** area of **Disk 1**, right-click the blank area and choose **New Simple Volume**.

Figure 5-34 New Simple Volume (data disk)

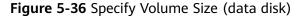


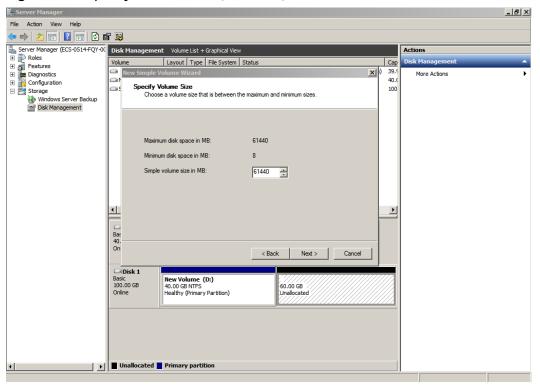
Step 5 On the displayed New Simple Volume Wizard window, click Next.

Figure 5-35 New Simple Volume Wizard (data disk)



Step 6 On the displayed **Specify Volume Size** page, set **Simple volume size in MB** and click **Next**. In this example, the default size is used.





Step 7 On the displayed **Assign Drive Letter and Path** page, click **Assign the following drive letter**, select a drive letter, and click **Next**. In this example, drive letter **E** is selected.

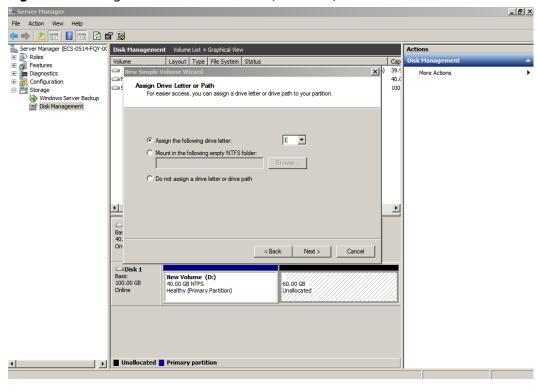
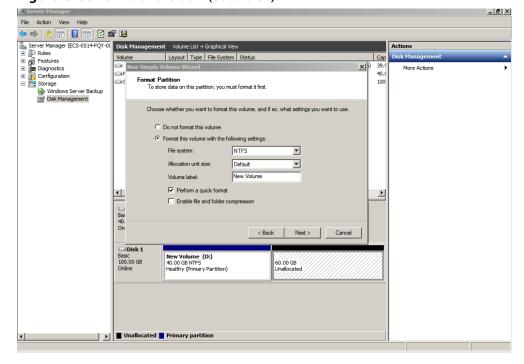


Figure 5-37 Assign Driver Letter or Path (data disk)

Step 8 On the displayed **Format Partition** page, click **Format this volume with the following settings**, set parameters based on the requirements, and select **Perform a quick format**. Then, click **Next**.

Figure 5-38 Format Partition (data disk)



Step 9 Click Finish.

After the expansion succeeded, new volume (E:) is displayed.

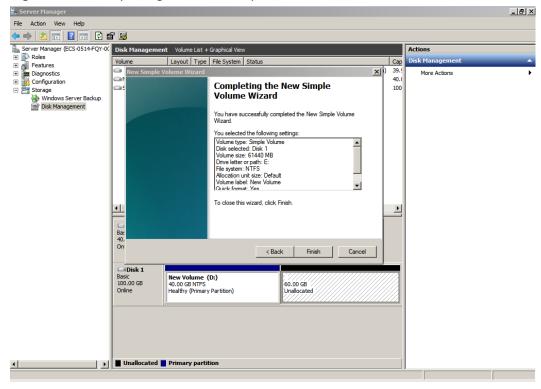
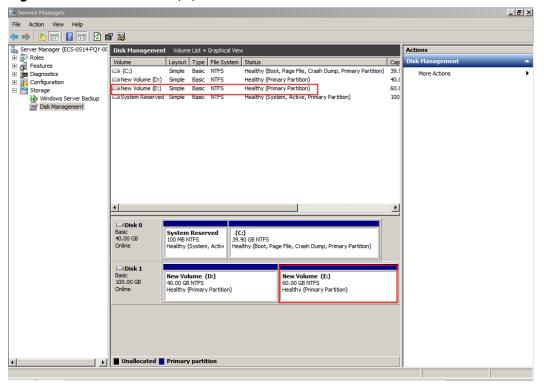


Figure 5-39 Completing the New Simple Volume Wizard (new volume E:)

Figure 5-40 New Volume (E:)



----End

5.5 Extending Disk Partitions and File Systems (Windows Server 2016)

Scenarios

After a disk is expanded on the management console, the disk size is enlarged, but the additional space cannot be used directly.

In Windows, you must allocate the additional space to an existing partition or a new partition.

If the disk capacity is expanded on a stopped server, the additional space of a Windows system disk or Windows data disk will be automatically added to the partition at the end of the disk upon the server startup. In this case, the additional space can be used directly.

This section uses Windows Server 2016 Standard 64bit as the sample OS to describe the expansion methods:

- For a system disk:
 - If volume (C:) already exists, you can add the additional space to volume
 (C:) and use it as a system volume. For details, see System Disk: Add Additional Space to Volume (C:).
 - If volume (C:) already exists, you can create a new volume such as volume (F:) with the additional space and use the new volume as a data volume. For details, see System Disk: Create New Volume (F:) with Additional Space.
- For a data disk:
 - If volume (D:) already exists, you can add the additional space to volume (D:) and use it as a data volume. For details, see Data Disk: Add Additional Space to Volume (D:).
 - If volume (D:) already exists, you can create a new volume such as volume (E:) with the additional space and use the new volume as a data volume. For details, see Data Disk: Create New Volume (E:) with Additional Space.

The method for allocating the additional space varies with the server OS. This section is used for reference only. For detailed operations and differences, see the corresponding OS documents.

NOTICE

Performing the expansion operations with caution. Incorrect operations may lead to data loss or exceptions. So you are advised to back up the disk data using CBR or snapshots before expansion. For details about using CBR, see Managing EVS Backups. For details about using snapshots, see Creating a Snapshot.

Prerequisites

- You have expanded the disk capacity and attached the disk to a server on the management console. For details, see Expanding Capacity for an In-use EVS Disk or Expanding Capacity for an Available EVS Disk.
- You have logged in to the server.
 - For how to log in to an ECS, see the Elastic Cloud Server User Guide.
 - For how to log in to a BMS, see the Bare Metal Server User Guide.

System Disk: Add Additional Space to Volume (C:)

In this example, the system disk has 40 GiB originally, and 30 GiB is added on the management console. The following procedure describes how to add this 30 GiB to volume (C:) on the server. After the operation is complete, volume (C:) will have 70 GiB of capacity and can be used as a system volume.

Step 1 On the desktop of the server, right-click the start icon in lower left corner and choose **Disk Management**.

The **Disk Management** window is displayed.

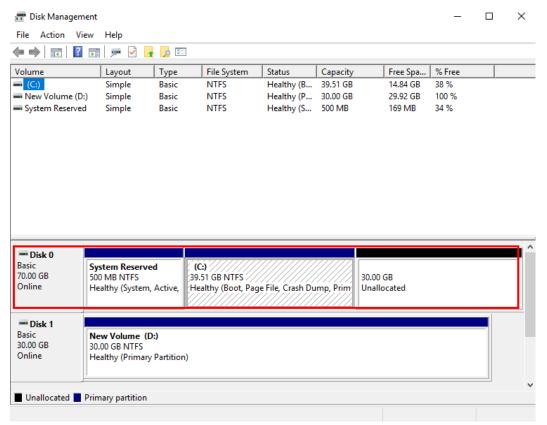


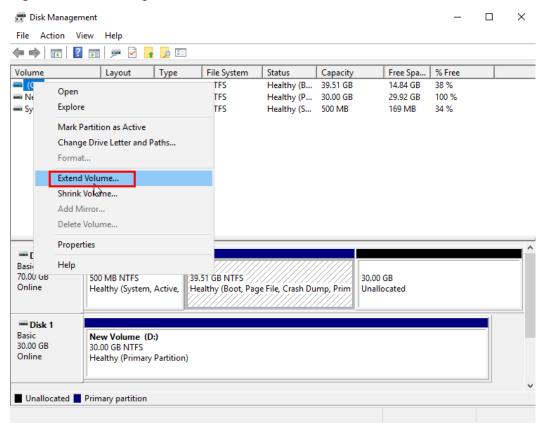
Figure 5-41 Disk Management (Windows Server 2016)

◯ NOTE

If you cannot view the additional space, right-click **Disk Management** and choose **Refresh** from the shortcut menu.

- **Step 2** On the **Disk Management** page, select the disk and volume that you want to extend. The current volume size and unallocated space are displayed.
- **Step 3** Right-click the target volume and choose **Extend Volume**.

Figure 5-42 Choosing Extend Volume (Windows Server 2016)



Step 4 On the displayed Extend Volume Wizard window, click Next.

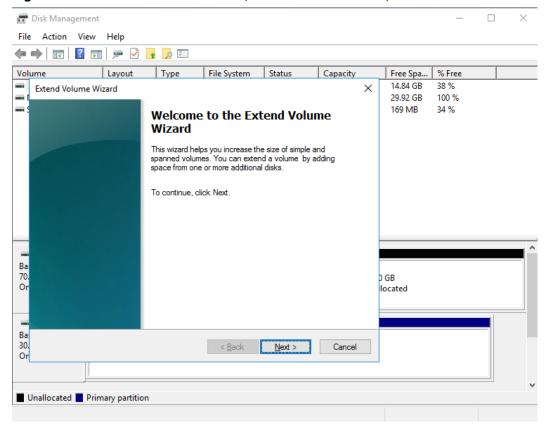


Figure 5-43 Extend Volume Wizard (Windows Server 2016)

Step 5 In the text box to the right of **Select the amount of space in MB**, enter the amount of the additional space and click **Next**.

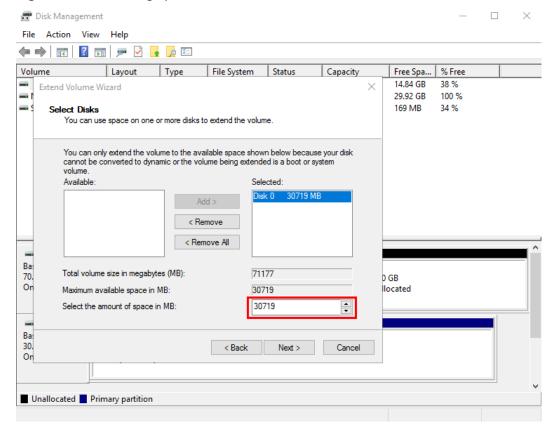


Figure 5-44 Selecting space (Windows Server 2016)

Step 6 Click Finish.

After the expansion succeeded, the partition size is larger than the original size.

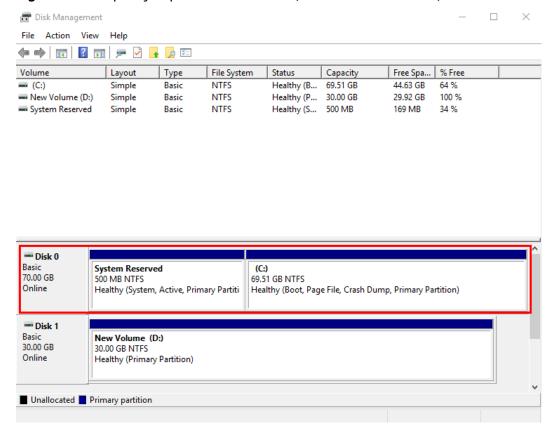


Figure 5-45 Capacity expansion succeeded (Windows Server 2016)

----End

System Disk: Create New Volume (F:) with Additional Space

In this example, the system disk has 40 GiB originally, and 60 GiB is added on the management console. The following procedure describes how to use this 60 GiB to create a new volume, for example volume (F:), on the server. After the operation is complete, new volume (F:) has 60 GiB of capacity and can be used as a data volume.

Step 1 On the desktop of the server, right-click the start icon in lower left corner and choose **Disk Management**.

The **Disk Management** window is displayed.

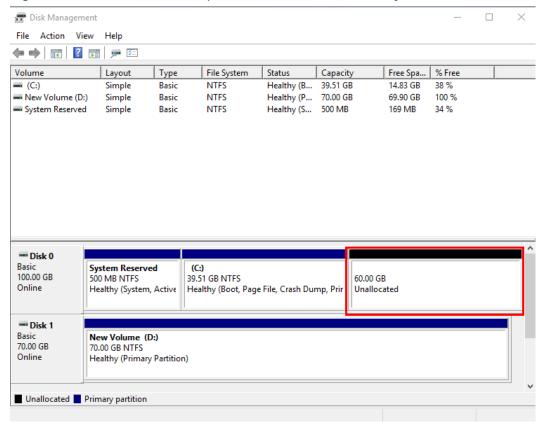


Figure 5-46 Unallocated disk space (Windows Server 2016 system disk)

□ NOTE

If you cannot view the additional space, right-click **Disk Management** and choose **Refresh** from the shortcut menu.

Step 2 In the **Unallocated** area of **Disk 0**, right-click the blank area and choose **New Simple Volume**.

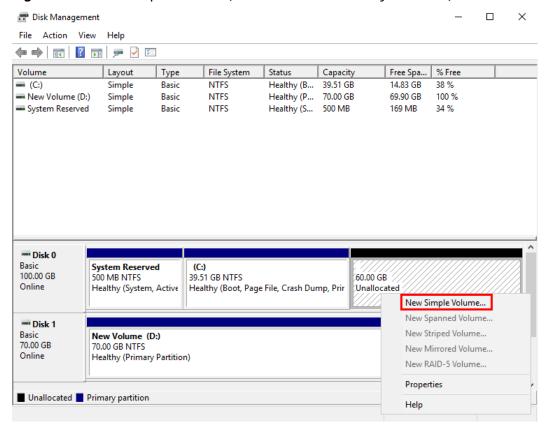


Figure 5-47 New Simple Volume (Windows Server 2016 system disk)

Step 3 On the displayed New Simple Volume Wizard window, click Next.

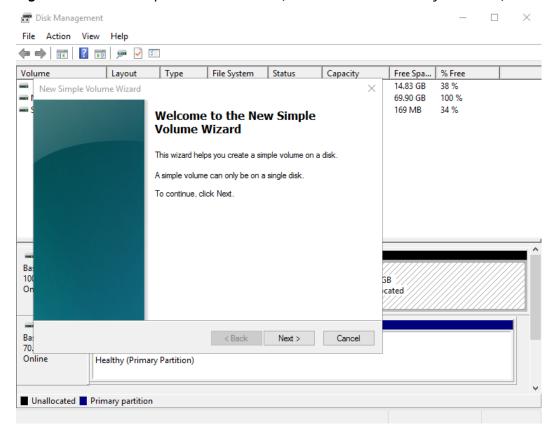


Figure 5-48 New Simple Volume Wizard (Windows Server 2016 system disk)

Step 4 On the displayed **Specify Volume Size** page, set **Simple volume size in MB** and click **Next**. In this example, the default size is used.

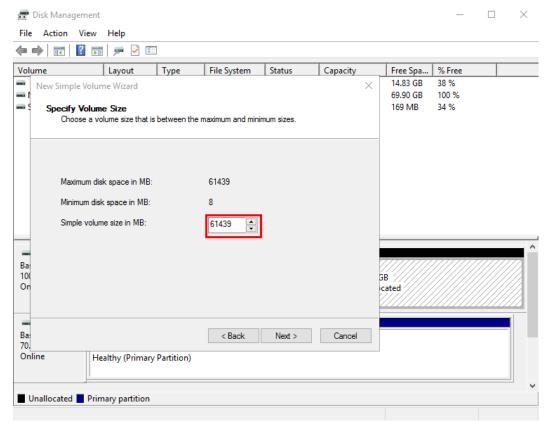


Figure 5-49 Specify Volume Size (Windows Server 2016 system disk)

Step 5 On the displayed **Assign Drive Letter and Path** page, click **Assign the following drive letter**, select a drive letter, and click **Next**. In this example, drive letter **F** is selected.

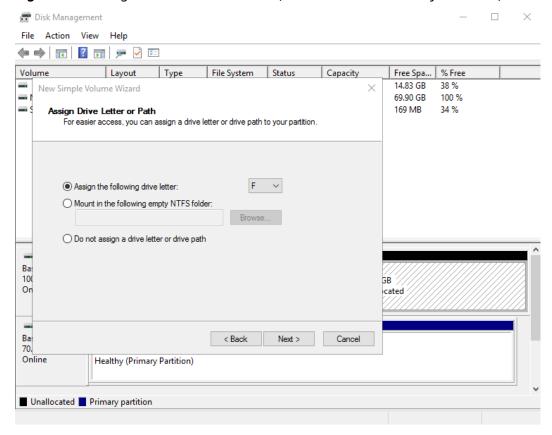


Figure 5-50 Assign Driver Letter or Path (Windows Server 2016 system disk)

Step 6 On the displayed Format Partition page, click Format this volume with the following settings, set parameters based on the requirements, and select Perform a quick format. Then, click Next.

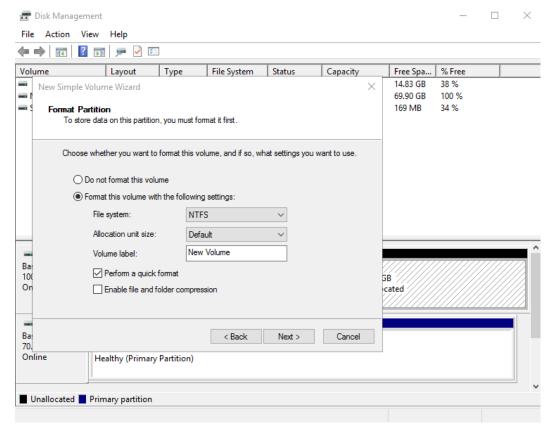


Figure 5-51 Format Partition (Windows Server 2016 system disk)

Step 7 Click Finish.

After the expansion succeeded, new volume (F:) is displayed.

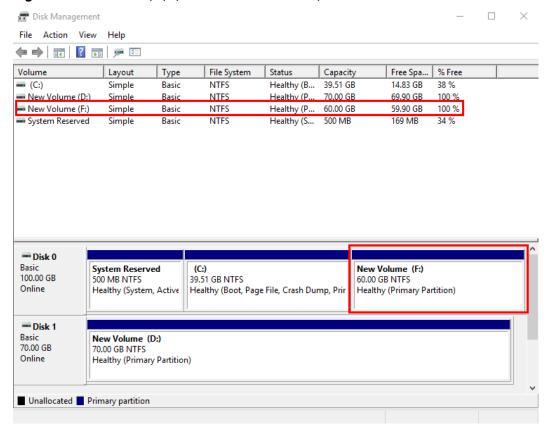


Figure 5-52 Volume (F:) (Windows Server 2016)

----End

Data Disk: Add Additional Space to Volume (D:)

In this example, the data disk has 30 GiB originally, and 50 GiB is added on the management console. The following procedure describes how to add this 50 GiB to volume (D:) on the server. After the operation is complete, volume (D:) has 80 GiB of capacity and can be used as a data volume.

Step 1 On the desktop of the server, right-click the start icon in lower left corner and choose **Disk Management**.

The **Disk Management** window is displayed.

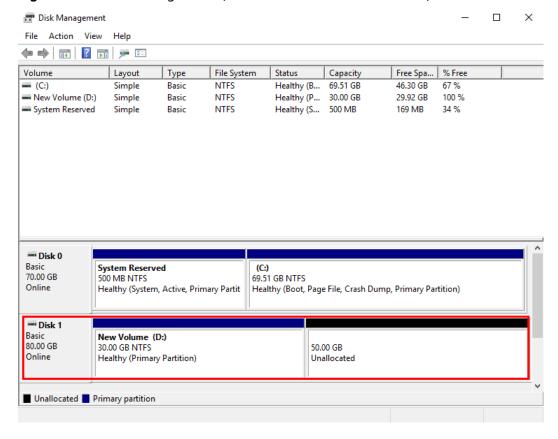


Figure 5-53 Disk Management (Windows Server 2016 data disk)

□ NOTE

If you cannot view the additional space, right-click **Disk Management** and choose **Refresh** from the shortcut menu.

- **Step 2** On the **Disk Management** page, select the disk and volume that you want to extend. The current volume size and unallocated space are displayed.
- Step 3 Right-click the target volume and choose Extend Volume.

2023-06-12

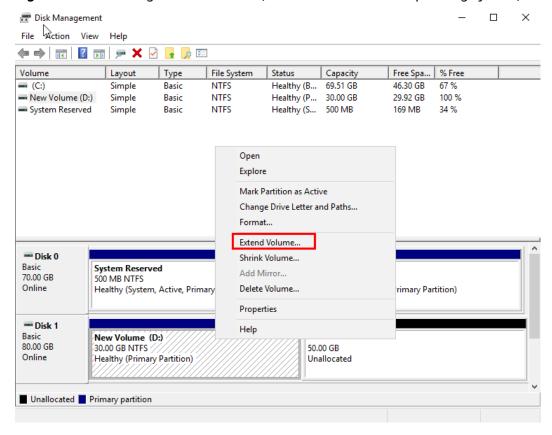


Figure 5-54 Choosing Extend Volume (Windows Server 2016 operating system)

Step 4 On the displayed **Extend Volume Wizard** window, click **Next**.

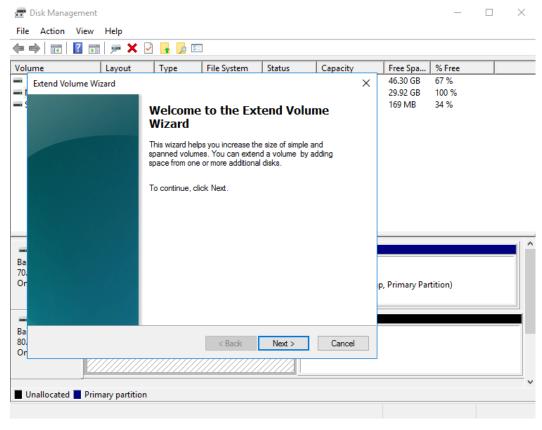


Figure 5-55 Extend Volume Wizard (Windows Server 2016 operating system)

Step 5 In the text box to the right of **Select the amount of space in MB**, enter the amount of the additional space and click **Next**.

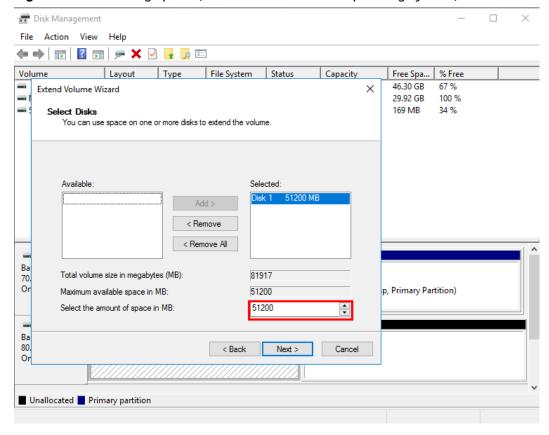


Figure 5-56 Selecting space (Windows Server 2016 operating system)

Step 6 Click Finish.

After the expansion succeeded, the partition size is larger than the original size.

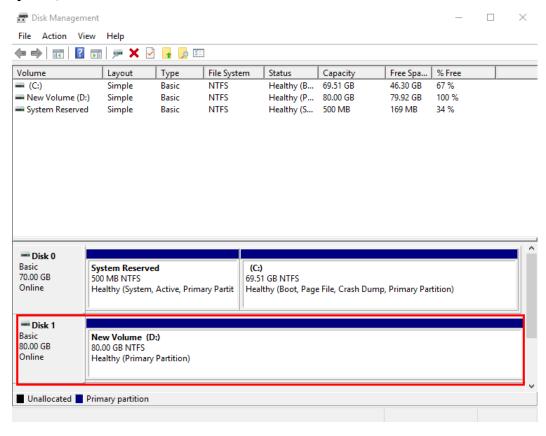


Figure 5-57 Capacity expansion succeeded (Windows Server 2016 operating system)

----End

Data Disk: Create New Volume (E:) with Additional Space

In this example, the data disk has 80 GiB originally, and 50 GiB is added on the management console. The following procedure describes how to use this 50 GiB to create a new volume, for example volume (E:), on the server. After the operation is complete, new volume (E:) has 50 GiB of capacity and can be used as a data volume.

Step 1 On the desktop of the server, right-click the start icon in lower left corner and choose **Disk Management**.

The **Disk Management** window is displayed.

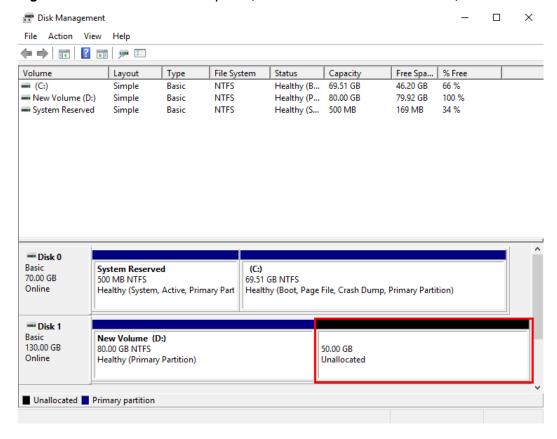


Figure 5-58 Unallocated disk space (Windows Server 2016 data disk)

Ⅲ NOTE

If you cannot view the additional space, right-click **Disk Management** and choose **Refresh** from the shortcut menu.

Step 2 In the **Unallocated** area of **Disk 1**, right-click the blank area and choose **New Simple Volume**.

2023-06-12

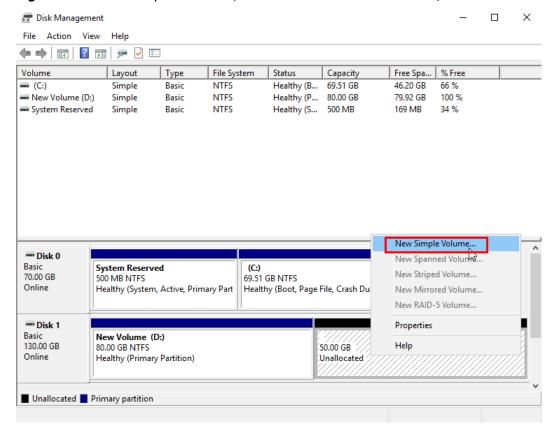


Figure 5-59 New Simple Volume (Windows Server 2016 data disk)

Step 3 On the displayed New Simple Volume Wizard window, click Next.

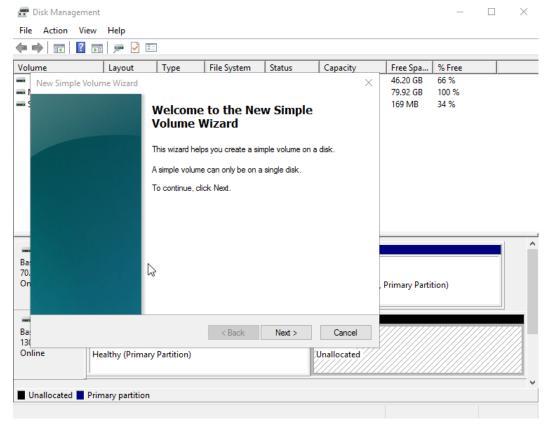


Figure 5-60 New Simple Volume Wizard (Windows Server 2016 data disk)

Step 4 On the displayed **Specify Volume Size** page, set **Simple volume size in MB** and click **Next**. In this example, the default size is used.

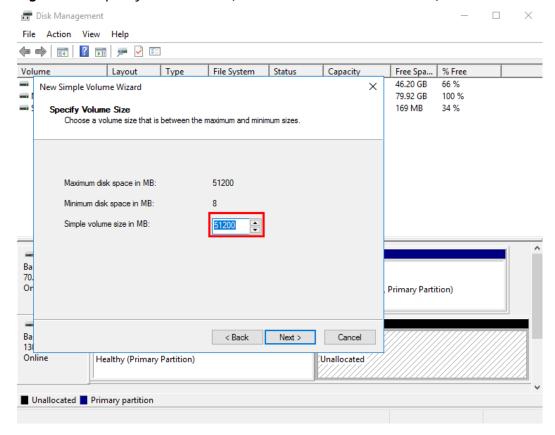


Figure 5-61 Specify Volume Size (Windows Server 2016 data disk)

Step 5 On the displayed **Assign Drive Letter and Path** page, click **Assign the following drive letter**, select a drive letter, and click **Next**. In this example, drive letter **E** is selected.

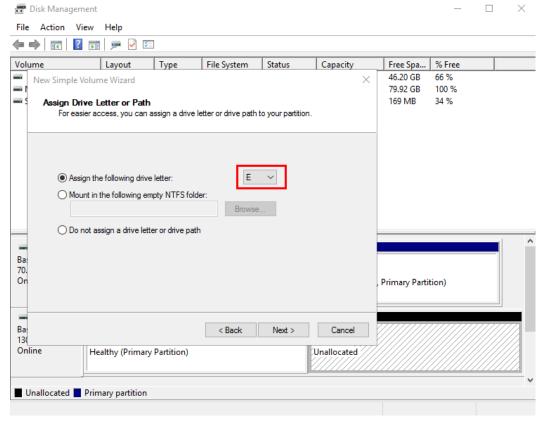


Figure 5-62 Assign Driver Letter or Path (Windows Server 2016 data disk)

Step 6 On the displayed Format Partition page, click Format this volume with the following settings, set parameters based on the requirements, and select Perform a quick format. Then, click Next.

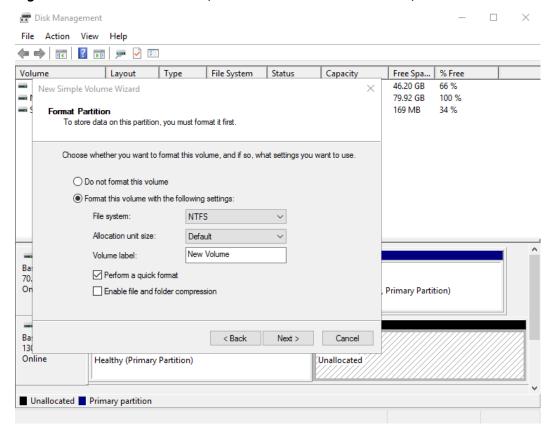
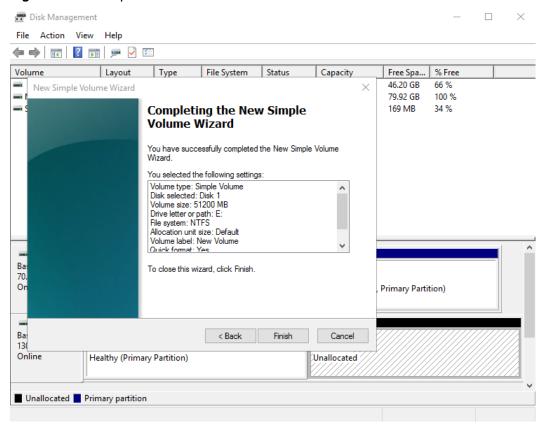


Figure 5-63 Format Partition (Windows Server 2016 data disk)

Step 7 Click Finish.

After the expansion succeeded, new volume (E:) is displayed.

Figure 5-64 Completed



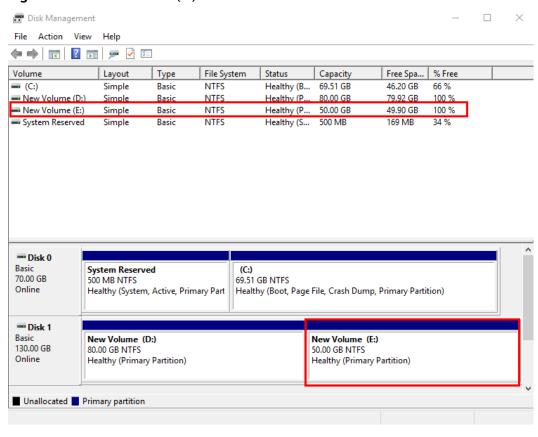


Figure 5-65 New Volume (E:)

----End

5.6 Extending Disk Partitions and File Systems (Linux)

5.6.1 Partition and File System Extension Preparations (Linux)

Before extending the disk partition and file system, you must check the disk partition style and file system format, and then select the appropriate operation accordingly.

- 1. To view the disk partition style, see the following methods:
 - Method 1: Check Partition Style and File System Format Using fdisk
 - Method 2: Check Partition Style and File System Format Using parted
- To select the appropriate operation, see Table 5-3.

Table 5-3 Disk partition and file system extension scenarios

Disk	Scenario	Method
System disk	Create a new MBR partition with the additional space.	Creating a New MBR Partition

Disk	Scenario	Method
	Allocate the additional space to an existing MBR partition.	• Extending an Existing MBR Partition (Kernel Version Later Than 3.6.0)
		 Extending an Existing MBR Partition (Kernel Version Earlier Than 3.6.0)
Data disk	Create a new MBR partition with the additional space.	Creating a New MBR Partition
	Allocate the additional space to an existing MBR partition.	Extending an Existing MBR Partition
	Create a new GPT partition with the additional space.	Creating a New GPT Partition
	Allocate the additional space to an existing GPT partition.	Extending an Existing GPT Partition
SCSI data disk	Create a new MBR partition with the additional space.	Creating a New MBR Partition
	Allocate the additional space to an existing MBR partition.	Extending an Existing MBR Partition

Ⅲ NOTE

The maximum disk capacity that MBR supports is 2 TiB, and the disk space exceeding 2 TiB cannot be used.

If your disk uses MBR and you need to expand the disk capacity to over 2 TiB, change the partition style from MBR to GPT. Ensure that the disk data has been backed up before changing the partition style because services will be interrupted and data on the disk will be cleared during this change.

Method 1: Check Partition Style and File System Format Using fdisk

Step 1 Run the following command to view all the disks attached to the server:

lsblk

Information similar to the following is displayed:

[root@ecs-test-0001 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda 253:0 0 40G 0 disk
—vda1 253:1 0 40G 0 part /
vdb 253:16 0 150G 0 disk
—vdb1 253:17 0 100G 0 part /mnt/sdc

In this example, data disk /dev/vdb already has partition /dev/vdb1 before capacity expansion, and the additional 50 GiB added has not been allocated yet. Therefore, /dev/vdb has 150 GiB, and /dev/vdb1 has 100 GiB.

If you run **lsblk** and find out that disk **/dev/vdb** has no partitions, format the disk by referring to **How Do I Extend the File System of an Unpartitioned Data Disk in Linux?** and expand the capacity. Otherwise, the additional space cannot be used after expansion.

Step 2 Run the following command to view the current disk partition style:

fdisk -l

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# fdisk -l
Disk /dev/vda: 42.9 GiB, 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: 0x000bcb4e
 Device Boot
                 Start
                           End
                                   Blocks Id System
                 2048 83886079 41942016 83 Linux
/dev/vda1 *
Disk /dev/vdb: 161.1 GiB, 161061273600 bytes, 314572800 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: 0x38717fc1
 Device Boot
                 Start
                           End
                                   Blocks Id System
                 2048 209715199 104856576 83 Linux
/dev/vdb1
```

The value in the **System** column indicates the disk partition style. Value **Linux** indicates the MBR partition style. Value **GPT** indicates the GPT partition style.

- If the disk partitions displayed are inconsistent with those obtained in Step 1, the partition that is not displayed uses the GPT partition style and has unallocated space. In this case, you cannot query all the partition information using the fdisk -l command. Go to Method 2: Check Partition Style and File System Format Using parted.
- If the disk partitions displayed are consistent with those obtained in Step 1, continue with the following operations.

Step 3 Run the following command to view the partition's file system format:

blkid Disk partition

In this example, run the following command:

blkid /dev/vdb1

```
[root@ecs-test-0001 ~]# blkid /dev/vdb1
/dev/vdb1: UUID="0b3040e2-1367-4abb-841d-ddb0b92693df" TYPE="ext4"
```

In the command output, the **TYPE** value is **ext4**, indicating that **/dev/vdb1**'s file system format is **ext4**.

Step 4 Run the following command to view the file system status:

```
ext*: e2fsck -n Disk partition
xfs: xfs_repair -n Disk partition
```

In this example, the ext4 file system is used. Therefore, run the following command:

e2fsck -n /dev/vdb1

Information similar to the following is displayed:

[root@ecs-test-0001 ~]# e2fsck -n /dev/vdb1 e2fsck 1.42.9 (28-Dec-2013) Warning! /dev/vdb1 is mounted. Warning: skipping journal recovery because doing a read-only filesystem check. /dev/vdb1: clean, 11/6553600 files, 459544/26214144 blocks

If the file system status is **clean**, the file system status is normal. Otherwise, rectify the faulty and then perform the capacity expansion.

----End

Method 2: Check Partition Style and File System Format Using parted

Step 1 Run the following command to view all the disks attached to the server:

lsblk

Information similar to the following is displayed:

In this example, data disk /dev/vdb already has partition /dev/vdb1 before capacity expansion, and the additional 50 GiB added has not been allocated yet. Therefore, /dev/vdb has 150 GiB, and /dev/vdb1 has 100 GiB.

If you run **lsblk** and find out that disk **/dev/vdb** has no partitions, format the disk by referring to **How Do I Extend the File System of an Unpartitioned Data Disk in Linux?** and expand the capacity. Otherwise, the additional space cannot be used after expansion.

Step 2 Run the following command and enter **p** to view the disk partition style:

parted Disk

For example, run the following command to view /dev/vdb's partition style:

parted /dev/vdb

Information similar to the following is displayed:

[root@ecs-test-0001 ~]# parted /dev/vdb GNU Parted 3.1

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

(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/vdb 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: Virtio Block Device (virtblk)

```
Disk /dev/vdb: 161GiB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number Start End Size File system Name Flags
1 1049kB 107GiB 107GiB ext4 test

(parted)
```

Partition Table indicates the disk partition style. **Partition Table: msdos** means MBR, **Partition Table: gpt** means GPT, and **Partition Table: loop** means that the whole disk is partitioned.

• If the following error information is displayed, enter **Fix**.

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)?

The GPT partition table information is stored at the start of the disk. To reduce the risk of damage, a backup of the information is saved at the end of the disk. When you expand the disk capacity, the end of the disk changes accordingly. In this case, enter **Fix** to move the backup file of the information to new disk end.

• If the following warning information is displayed, enter Fix.

Warning: Not all of the space available to /dev/vdb 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

Enter **Fix** as prompted. The system automatically sets the GPT partition style for the additional space.

Step 3 Enter **q** and press **Enter** to exit parted.

----End

5.6.2 Extending Partitions and File Systems for System Disks (Linux)

Scenarios

After a disk is expanded on the management console, the disk size is enlarged, but the additional space cannot be used directly.

In Linux, you must allocate the additional space to an existing partition or a new partition.

If the disk capacity is expanded when its server is stopped, the additional space of a Linux system disk will be automatically added to the partition at the disk end upon the server startup. In this case, the additional space can be used directly.

This section uses CentOS 7.4 64bit and CentOS 6.5 64bit as the sample OSs to describe how to extend the disk partition using growpart and fdisk. The method for allocating the additional space varies with the server OS. This section is used for reference only. For detailed operations and differences, see the corresponding OS documents.

For how to query the Linux kernel version, see **Querying the Linux Kernel Version**.

- Extending an Existing MBR Partition (Kernel Version Later Than 3.6.0)
- Extending an Existing MBR Partition (Kernel Version Earlier Than 3.6.0)
- Creating a New MBR Partition

NOTICE

Performing the expansion operations with caution. Incorrect operations may lead to data loss or exceptions. So you are advised to back up the disk data using CBR or snapshots before expansion. For details about using CBR, see **Managing EVS Backups**. For details about using snapshots, see **Creating a Snapshot**.

Prerequisites

- You have expanded the disk capacity and attached the disk to a server on the management console. For details, see Expanding Capacity for an In-use EVS Disk or Expanding Capacity for an Available EVS Disk.
- You have logged in to the server.
 - For how to log in to an ECS, see the *Elastic Cloud Server User Guide*.
 - For how to log in to a BMS, see the Bare Metal Server User Guide.

Constraints

The additional space can only be added to the last partition of the disk.

Querying the Linux Kernel Version

Run the following command to query the Linux kernel version:

uname -a

Then, perform corresponding operations depending on whether the Linux kernel version is later than 3.6.0.

- For CentOS 7.4 64bit, information similar to the following is displayed: Linux ecs-test-0001 3.10.0-957.5.1.el7.x86_64 #1 SMP Fri Feb 1 14:54:57 UTC 2019 x86_64 x86_64
 - The kernel version is 3.10.0, which is later than 3.6.0. For subsequent operations, see Extending an Existing MBR Partition (Kernel Version Later Than 3.6.0).
- For CentOS 6.5 64bit, information similar to the following is displayed: Linux ecs-test-0002 2.6.32-754.10.1.el6.x86_64 #1 SMP Tue Jan 15 17:07:28 UTC 2019 x86_64
 - The kernel version is 2.6.32, which is earlier than 3.6.0. In this case, the disk partition and file system extension take effect only after a server reboot. For subsequent operations, see **Extending an Existing MBR Partition (Kernel Version Earlier Than 3.6.0)**.

Extending an Existing MBR Partition (Kernel Version Later Than 3.6.0)

CentOS 7.4 64bit is used as the sample OS. Originally, system disk /dev/vda has 40 GiB and one partition (/dev/vda1), and then 60 GiB is added to the disk. The following procedure shows you how to allocate the additional 60 GiB to the existing MBR partition /dev/vda1.

Step 1 (Optional) Run the following command to install the growpart tool:

yum install cloud-utils-growpart

□ NOTE

You can run the **growpart** command to check whether the growpart tool has been installed. If the command output displays the tool usage instructions, the tool has been installed and you do not need to install it separately.

Step 2 Run the following command to view the total capacity of the **/dev/vda** system disk:

fdisk -l

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# fdisk -l

Disk /dev/vda: 107.4 GiB, 107374182400 bytes, 209715200 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: 0x000bcb4e

Device Boot Start End Blocks Id System

/dev/vda1 * 2048 83886079 41942016 83 Linux
```

Step 3 Run the following command to view the capacity of the /dev/vda1 partition:

df -TH

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# df -TH
Filesystem
           Type
                  Size Used Avail Use% Mounted on
                  43G 2.0G 39G 5%/
/dev/vda1
           ext4
           devtmpfs 2.0G 0 2.0G 0% /dev
devtmpfs
tmpfs
          tmpfs 2.0G 0 2.0G 0% /dev/shm
tmpfs
                 2.0G 9.0M 2.0G 1% /run
          tmpfs
                  2.0G  0  2.0G  0% /sys/fs/cgroup
tmpfs
          tmpfs
          tmpfs 398M 0 398M 0% /run/user/0
tmpfs
```

Step 4 Run the following command to extend the partition using growpart:

growpart System disk Partition number

In this example, run the following command:

growpart /dev/vda 1

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# growpart /dev/vda 1
CHANGED: partition=1 start=2048 old: size=83884032 end=83886080 new: size=209713119,end=209715167
```

Step 5 Run the following command to extend the file system of the partition:

resize2fs Disk partition

In this example, run the following command:

resize2fs /dev/vda1

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# resize2fs /dev/vda1
resize2fs 1.42.9 (28-Dec-2013)
Filesystem at /dev/vda1 is mounted on /; on-line resizing required
old_desc_blocks = 5, new_desc_blocks = 13
The filesystem on /dev/vda1 is now 26214139 blocks long.
```


If the error message "open: No such file or directory while opening /dev/vdb1" is returned, an incorrect partition is specified. Run **df -TH** to view the disk partitions.

Step 6 Run the following command to view the new capacity of the /dev/vda1 partition:

df -TH

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# df -TH
           Type Size Used Avail Use% Mounted on
Filesystem
/dev/vda1
            ext4
                   106G 2.0G 99G 2%/
            devtmpfs 2.0G 0 2.0G 0% /dev
devtmpfs
           tmpfs 2.0G 0 2.0G 0% /dev/shm
tmpfs
           tmpfs 2.0G 9.0M 2.0G 1% /run
tmpfs 2.0G 0 2.0G 0% /svs/fs
tmpfs
                   2.0G  0  2.0G  0% /sys/fs/cgroup
tmpfs
           tmpfs 398M 0 398M 0% /run/user/0
tmpfs
```

----End

Extending an Existing MBR Partition (Kernel Version Earlier Than 3.6.0)

NOTICE

If the OS kernel version is earlier than 3.6.0, the extension of an existing MBR disk partition and file system takes effect only after a server reboot, and services will be interrupted.

CentOS 6.5 64bit is used as the sample OS. Originally, system disk /dev/vda has 40 GiB and one partition (/dev/vda1), and then 60 GiB is added to the disk. The following procedure shows you how to allocate the additional 60 GiB to the existing MBR partition /dev/vda1.

Step 1 (Optional) Run the following command to install the growpart tool:

yum install cloud-utils-growpart

◯ NOTE

You can run the **growpart** command to check whether the growpart tool has been installed. If the command output displays the tool usage instructions, the tool has been installed and you do not need to install it separately.

Step 2 Run the following command to install the dracut-modules-growroot tool:

yum install cloud-utils-growpart

Information similar to the following is displayed:

[root@ecs-test-0002 ~]# yum install cloud-utils-growpart Loaded plugins: fastestmirror, security Setting up Install Process Determining fastest mirrors

Package cloud-utils-growpart-0.27-10.el6.x86_64 already installed and latest version Nothing to do

Step 3 Run the following command to regenerate the initramfs file:

dracut -f

Step 4 Run the following command to view the total capacity of the **/dev/vda** system disk:

fdisk -l

Information similar to the following is displayed:

[root@ecs-test-0002 ~]# fdisk -l

Disk /dev/vda: 107.4 GiB, 107374182400 bytes
255 heads, 63 sectors/track, 13054 cylinders

Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x0004e0be

Device Boot Start End Blocks Id System
/dev/vda1 * 1 5222 41942016 83 Linux

Step 5 Run the following command to view the capacity of the /dev/vda1 partition:

df -TH

Information similar to the following is displayed:

[root@ecs-test-0002 ~]# df -TH
Filesystem Type Size Used Avail Use% Mounted on /dev/vda1 ext4 43G 1.7G 39G 5% /
tmpfs tmpfs 2.1G 0 2.1G 0% /dev/shm

Step 6 Run the following command to extend the partition using growpart:

growpart System disk Partition number

In this example, run the following command:

growpart /dev/vda 1

Information similar to the following is displayed:

[root@ecs-test-0002 ~]# growpart /dev/vda 1 CHANGED: partition=1 start=2048 old: size=83884032 end=83886080 new: size=209710462,end=209712510

Step 7 Run the following command to restart the server:

reboot

After the server is restarted, reconnect to the server and perform the following steps.

Step 8 Run the following command to extend the file system of the partition:

resize2fs Disk partition

In this example, run the following command:

resize2fs /dev/vda1

Information similar to the following is displayed:

[root@ecs-test-0002 ~]# resize2fs /dev/vda1 resize2fs 1.41.12 (17-May-2010)
The filesystem is already 26213807 blocks long. Nothing to do!

Ⅲ NOTE

If the error message "open: No such file or directory while opening /dev/vdb1" is returned, an incorrect partition is specified. Run **df** -**TH** to view the disk partitions.

Step 9 Run the following command to view the new capacity of the /dev/vda1 partition:

df -TH

Information similar to the following is displayed:

```
[root@ecs-test-0002 ~]# df -TH
Filesystem Type Size Used Avail Use% Mounted on
/dev/vda1 ext4 106G 1.7G 99G 2% /
tmpfs tmpfs 2.1G 0 2.1G 0% /dev/shm
```

----End

Creating a New MBR Partition

Originally, system disk /dev/vda has 40 GiB and one partition (/dev/vda1), and then 40 GiB is added to the disk. The following procedure shows you how to create a new MBR partition /dev/vda2 with this 40 GiB.

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

fdisk -l

Information similar to the following is displayed:

```
[root@ecs-2220 ~]# fdisk -l

Disk /dev/vda: 85.9 GiB, 85899345920 bytes, 167772160 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: 0x0008d18f

Device Boot Start End Blocks Id System
/dev/vda1 * 2048 83886079 41942016 83 Linux
```

In the command output, the capacity of the **dev/vda** system disk is 80 GiB, in which the in-use **dev/vda1** partition takes 40 GiB and the additional 40 GiB has not been allocated.

Step 2 Run the following command to enter fdisk:

fdisk /dev/vda

Information similar to the following is displayed:

```
[root@ecs-2220 ~]# fdisk /dev/vda
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 new partition.

Information similar to the following is displayed:

```
Command (m for help): n
Partition type:
    p primary (1 primary, 0 extended, 3 free)
    e extended
```

There are two types of disk partitions:

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

□ NOTE

If the MBR partition style is used, a maximum of 4 primary partitions, or 3 primary partitions and 1 extended partition can be created. The extended partition cannot be used directly and must be divided into logical partitions before use.

Disk partitions created using GPT are not categorized.

Step 4 In this example, a primary partition is created. Therefore, enter **p** and press **Enter** to create a primary partition.

Information similar to the following is displayed:

Select (default p): p Partition number (2-4, default 2):

Step 5 Partition number **2** is used in this example. Therefore, enter **2** and press **Enter**.

Information similar to the following is displayed:

Partition number (2-4, default 2): 2 First sector (83886080-167772159, default 83886080):

Step 6 Enter the new partition's start sector and press Enter. In this example, the default start sector is used.

The system displays the start and end sectors of the partition's available space. You can customize the value within this range or use the default value. The start sector must be smaller than the partition's end sector.

Information similar to the following is displayed:

First sector (83886080-167772159, default 83886080): Using default value 83886080 Last sector, +sectors or +size{K,M,G} (83886080-167772159,default 167772159):

Step 7 Enter the new partition's end sector and press **Enter**. In this example, the default end sector is used.

The system displays the start and end sectors of the partition's available space. You can customize the value within this range or use the default value. The start sector must be smaller than the partition's end sector.

Information similar to the following is displayed:

Last sector, +sectors or +size{K,M,G} (83886080-167772159, default 167772159): Using default value 167772159 Partition 2 of type Linux and of size 40 GiB is set

Command (m for help):

Step 8 Enter **p** and press **Enter** to view the new partition.

Information similar to the following is displayed:

Command (m for help): p

Disk /dev/vda: 85.9 GiB, 85899345920 bytes, 167772160 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: 0x0008d18f

2023-06-12 172

```
Device Boot Start End Blocks Id System
/dev/vda1 * 2048 83886079 41942016 83 Linux
/dev/vda2 83886080 167772159 41943040 83 Linux
Command (m for help):
```

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

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.

The partition is created.

Ⅲ NOTE

In case that you want to discard the changes made before, you can exit fdisk by entering \mathbf{q} .

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

partprobe

Step 11 Run the following command to set the file system format for the new partition:

mkfs -t File system Disk partition

• Sample command of the ext* file system:

(The ext4 file system is used in this example.)

mkfs -t ext4 /dev/vda2

Information similar to the following is displayed:

[root@ecs-2220 ~]# mkfs -t ext4 /dev/vda2 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

Sample command of the xfs file system:

mkfs -t xfs /dev/vda2

Information similar to the following is displayed:

```
[root@ecs-2220 ~]# mkfs -t xfs /dev/vda2
meta-data=/dev/vda2 isize=512 agcount=4, agsize=2621440 blks
```

```
= sectsz=512 attr=2, projid32bit=1
= crc=1 finobt=0, sparse=0
data = bsize=4096 blocks=10485760, imaxpct=25
= sunit=0 swidth=0 blks
naming =version2 bsize=4096 ascii-ci=0 ftype=1
log =internal log bsize=4096 blocks=5120, version=2
= sectsz=512 sunit=0 blks, lazy-count=1
realtime =none extsz=4096 blocks=0, rtextents=0
```

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 12 (Optional) Run the following command to create a mount point:

Perform this step if you want to mount the partition on a new mount point.

mkdir Mount point

In this example, run the following command to create the **/opt** mount point:

mkdir /opt

Step 13 Run the following command to mount the new partition:

mount Disk partition Mount point

In this example, run the following command to mount the new partition /dev/vda2 on /opt:

mount /dev/vda2 /opt

□ NOTE

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

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

df -TH

Information similar to the following is displayed:

```
[root@ecs-2220 ~]# df -TH
Filesystem Type
               Size Used Avail Use% Mounted on
/dev/vda1
         ext4
                43G 2.0G 39G 5% /
         devtmpfs 509M 0 509M 0% /dev
devtmpfs
        tmpfs 520M 0 520M 0% /dev/shm
tmpfs
               520M 7.2M 513M 2% /run
tmpfs
        tmpfs
               tmpfs
        tmpfs
        tmpfs
tmpfs
/dev/vda2 ext4 43G 51M 40G 1% /opt
```

∩ NOTE

If the server is restarted, the mounting will become invalid. You can set automatic mounting for partitions at system start by modifying the /etc/fstab file. For details, see Configuring Automatic Mounting at System Start.

----End

Configuring Automatic Mounting at System Start

The **fstab** file controls what disks are automatically mounted at startup. You can use **fstab** to configure your data disks to mount automatically. This operation will not affect the existing data.

The example here uses UUIDs to identify disks in the **fstab** file. You are advised not to use device names to identify disks in the file because device names are assigned dynamically and may change (for example, from /dev/vdb1 to /dev/vdb2) after a reboot. This can even prevent your disk from booting up.

■ NOTE

UUID is the unique character string for disk partitions in a Linux system.

Step 1 Query the partition UUID.

blkid Disk partition

In this example, the UUID of the /dev/vdb1 partition is queried.

blkid /dev/vdb1

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# blkid /dev/vdb1 /dev/vdb1: UUID="0b3040e2-1367-4abb-841d-ddb0b92693df" TYPE="ext4"
```

Carefully record the UUID, as you will need it for the following step.

Step 2 Open the **fstab** file using the vi editor.

vi /etc/fstab

- **Step 3** Press i to enter editing mode.
- **Step 4** Move the cursor to the end of the file and press **Enter**. Then, add the following information:

```
UUID=0b3040e2-1367-4abb-841d-ddb0b92693df /mnt/sdc ext4 defaults 0
```

In this example, the line starting with "UUID" is the information added. Edit this line to match the following format:

- UUID: The UUID obtained in Step 1.
- Mount point: The directory on which the partition is mounted. You can query the mount point using df -TH.
- Filesystem: The file system format of the partition. You can query the file system format using **df -TH**.
- Mount option: The partition mount option. Usually, this parameter is set to defaults.
- Dump: The Linux dump backup option.
 - **0**: Linux dump backup is not used. Usually, dump backup is not used, and you can set this parameter to **0**.
 - 1: Linux dump backup is used.
- fsck: The fsck option, which means whether to use fsck to check the disk during startup.
 - 0: not use fsck.

If the mount point is the root partition (/), this parameter must be set to

If this parameter is set to **1** for the root partition, this parameter for other partitions must start with **2** because the system checks the partitions in the ascending order of the values.

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

The system saves the configurations and exits the vi editor.

Step 6 Verify that the disk is auto-mounted at startup.

1. Unmount the partition.

umount Disk partition

In this example, run the following command:

umount /dev/vdb1

2. Reload all the content in the /etc/fstab file.

mount -a

3. Query the file system mounting information.

mount | grep Mount point

In this example, run the following command:

mount | grep /mnt/sdc

If information similar to the following is displayed, automatic mounting has been configured:

```
root@ecs-test-0001 ~]# mount | grep /mnt/sdc
/dev/vdb1 on /mnt/sdc type ext4 (rw,relatime,data=ordered)
```

----End

5.6.3 Extending Partitions and File Systems for Data Disks (Linux)

Scenarios

After a disk is expanded on the management console, the disk size is enlarged, but the additional space cannot be used directly.

In Linux, you must allocate the additional space to an existing partition or a new partition.

This section uses CentOS 7.4 64bit as the sample OS to describe how to extend an MBR or GPT partition. The method for allocating the additional space varies with the server OS. This section is used for reference only. For detailed operations and differences, see the corresponding OS documents.

- Creating a New MBR Partition
- Extending an Existing MBR Partition
- Creating a New GPT Partition
- Extending an Existing GPT Partition

NOTICE

Performing the expansion operations with caution. Incorrect operations may lead to data loss or exceptions. So you are advised to back up the disk data using CBR or snapshots before expansion. For details about using CBR, see **Managing EVS Backups**. For details about using snapshots, see **Creating a Snapshot**.

Prerequisites

- You have expanded the disk capacity and attached the disk to a server on the management console. For details, see Expanding Capacity for an In-use EVS Disk or Expanding Capacity for an Available EVS Disk.
- You have logged in to the server.
 - For how to log in to an ECS, see the *Elastic Cloud Server User Guide*.
 - For how to log in to a BMS, see the Bare Metal Server User Guide.

Constraints

- The additional space of a data disk cannot be added to the root partition. To extend the root partition, expand the system disk instead.
- The additional space can only be added to the last partition of the disk.

Creating a New MBR Partition

Originally, data disk /dev/vdb has 100 GiB and one partition (/dev/vdb1), and then 50 GiB is added to the disk. The following procedure shows you how to create a new MBR partition /dev/vdb2 with this 50 GiB.

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

fdisk -l

```
Information similar to the following is displayed:
```

```
[root@ecs-test-0001 ~]# fdisk -l
Disk /dev/vda: 42.9 GiB, 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: 0x000bcb4e
 Device Boot
                 Start
                           End
                                  Blocks Id System
                 2048 83886079 41942016 83 Linux
/dev/vda1
Disk /dev/vdb: 161.1 GiB, 161061273600 bytes, 314572800 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: 0x38717fc1
  Device Boot
                 Start
                           End
                                  Blocks Id System
           2048 209715199 104856576 83 Linux
/dev/vdb1
```

Step 2 Run the following command to enter fdisk:

fdisk Disk

In this example, run the following command:

fdisk /dev/vdb

Information similar to the following is displayed:

[root@ecs-test-0001 ~]# fdisk /dev/vdb 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 new partition.

Information similar to the following is displayed:

Command (m for help): n

Partition type:

- p primary (1 primary, 0 extended, 3 free)
- e extended

Select (default p):

There are two types of disk partitions:

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

If the MBR partition style is used, a maximum of 4 primary partitions, or 3 primary partitions and 1 extended partition can be created. The extended partition cannot be used directly and must be divided into logical partitions before use.

Disk partitions created using GPT are not categorized.

Step 4 In this example, a primary partition is created. Therefore, enter **p** and press **Enter** to create a primary partition.

Information similar to the following is displayed:

Select (default p): p

Partition number (2-4, default 2):

Partition number indicates the serial number of the primary partition. Because partition number 1 has been used, the value ranges from **2** to **4**.

Step 5 Enter the serial number of the primary partition and press Enter. Partition number2 is used in this example. Therefore, enter 2 and press Enter.

Information similar to the following is displayed:

Partition number (2-4, default 2): 2

First sector (209715200-314572799, default 209715200):

First sector indicates the start sector. The value ranges from **209715200** to **314572799**, and the default value is **209715200**.

Step 6 Enter the new partition's start sector and press **Enter**. In this example, the default start sector is used.

The system displays the start and end sectors of the partition's available space. You can customize the value within this range or use the default value. The start sector must be smaller than the partition's end sector.

Information similar to the following is displayed:

```
First sector (209715200-314572799, default 209715200):
Using default value 209715200
Last sector, +sectors or +size{K,M,G} (209715200-314572799, default 314572799):
```

Last sector indicates the end sector. The value ranges from **209715200** to **314572799**, and the default value is **314572799**.

Step 7 Enter the new partition's end sector and press **Enter**. In this example, the default end sector is used.

The system displays the start and end sectors of the partition's available space. You can customize the value within this range or use the default value. The start sector must be smaller than the partition's end sector.

Information similar to the following is displayed:

Last sector, +sectors or +size{K,M,G} (209715200-314572799, default 314572799):
Using default value 314572799
Partition 2 of type Linux and of size 50 GiB is set

Command (m for help):

Step 8 Enter **p** and press **Enter** to view the new partition.

Information similar to the following is displayed:

Command (m for help): p

Disk /dev/vdb: 161.1 GiB, 161061273600 bytes, 314572800 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: 0x38717fc1

Device Boot Start End Blocks Id System

/dev/vdb1 2048 209715199 104856576 83 Linux
/dev/vdb2 209715200 314572799 52428800 83 Linux

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

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.

□ NOTE

Command (m for help):

In case that you want to discard the changes made before, you can exit fdisk by entering \mathbf{q} .

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

partprobe

Step 11 Run the following command to set the file system format for the new partition:

mkfs -t File system Disk partition

• Sample command of the ext* file system:

mkfs -t ext4 /dev/vdb2

```
Information similar to the following is displayed:
```

```
[root@ecs-test-0001 ~]# mkfs -t ext4 /dev/vdb2
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
3276800 inodes, 13107200 blocks
655360 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2162163712
400 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, 11239424
Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

• Sample command of the xfs file system:

mkfs -t xfs /dev/vdb2

Information similar to the following is displayed:

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 12 (Optional) Run the following command to create a mount point:

Perform this step if you want to mount the partition on a new mount point.

mkdir Mount point

In this example, run the following command to create the /mnt/test mount point:

mkdir /mnt/test

Step 13 Run the following command to mount the new partition:

mount Disk partition Mount point

In this example, run the following command to mount the new partition /dev/vdb2 on /mnt/test:

mount /dev/vdb2 /mnt/test

□ NOTE

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

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

df -TH

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# df -TH
Filesystem
                  Size Used Avail Use% Mounted on
          Type
                  43G 1.9G 39G 5% /
/dev/vda1
           ext4
           devtmpfs 2.0G 0 2.0G 0% /dev
devtmpfs
                 2.0G 0 2.0G 0% /dev/shm
          tmpfs
tmpfs
tmpfs
          tmpfs
                  2.0G 9.1M 2.0G 1% /run
                 2.0G 0 2.0G 0% /sys/fs/cgroup
tmpfs
          tmpfs
tmpfs
          tmpfs
                 398M
                        0 398M 0% /run/user/0
/dev/vdb1
           ext4
                  106G 63M 101G 1% /mnt/sdc
/dev/vdb2
                 53G 55M 50G 1% /mnt/test
           ext4
```


If the server is restarted, the mounting will become invalid. You can set automatic mounting for partitions at system start by modifying the /etc/fstab file. For details, see Configuring Automatic Mounting at System Start.

----End

Extending an Existing MBR Partition

NOTICE

If the additional space is allocated to an existing partition, data on the disk will not be cleared but you must use **umount** to unmount the existing partition. In this case, services will be affected.

Originally, data disk /dev/vdb has 150 GiB and two partitions (/dev/vdb1 and /dev/vdb2), and then 80 GiB is added to the disk. The following procedure shows you how to add this 80 GiB to the existing MBR partition /dev/vdb2.

NOTICE

During an expansion, the additional space is added to the end of the disk. Therefore, if the disk has multiple partitions, the additional space can only be allocated to the partition at the disk end.

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

fdisk -l

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# fdisk -l
Disk /dev/vda: 42.9 GiB, 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: 0x000bcb4e
 Device Boot
                           End
                                  Blocks Id System
                Start
/dev/vda1 *
                2048 83886079 41942016 83 Linux
Disk /dev/vdb: 247.0 GiB, 246960619520 bytes, 482344960 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: 0x38717fc1
 Device Boot
                Start
                           End
                                  Blocks Id System
                2048 209715199 104856576 83 Linux
/dev/vdb1
/dev/vdb2
             209715200 314572799 52428800 83 Linux
```

In the command output, take note of the partition's start and end sectors. In this example, /dev/vdb2's start sector is 209715200, and its end sector is 314572799.

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

- If the additional space is not included, refresh the capacity according to **Extending Partitions and File Systems for SCSI Disks (Linux)**.
- If the additional space is included, take note of the start and end sectors of the target partition and then go to **Step 2**. These values will be used in the subsequent operations.

Step 2 Run the following command to unmount the partition:

```
umount Disk partition
```

In this example, run the following command:

umount /dev/vdb2

Step 3 Run the following command to enter fdisk:

fdisk Disk

In this example, run the following command:

fdisk /dev/vdb

Information similar to the following is displayed:

[root@ecs-test-0001 ~]# fdisk /dev/vdb 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 4 Run the following command to delete the partition to be extended:

1. Enter **d** and press **Enter** to delete the partition.

Information similar to the following is displayed:

Command (m for help): d Partition number (1,2, default 2):

2. Enter the partition number and press **Enter** to delete the partition. In this example, enter **2**.

Information similar to the following is displayed:

Partition number (1,2, default 2): 2 Partition 2 is deleted

Command (m for help):

◯ NOTE

After deleting the partition, recreate the partition according to the following steps, and data on this disk will not be lost.

Step 5 Enter **n** and press **Enter** to create a new partition.

Information similar to the following is displayed:

Command (m for help): n Partition type:

p primary (1 primary, 0 extended, 3 free)

e extended

Select (default p):

There are two types of disk partitions:

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

If the MBR partition style is used, a maximum of 4 primary partitions, or 3 primary partitions and 1 extended partition can be created. The extended partition cannot be used directly and must be divided into logical partitions before use.

Disk partitions created using GPT are not categorized.

Step 6 Ensure that the entered partition type is the same as the partition had before. In this example, a primary partition is used. Therefore, enter **p** and press **Enter** to create a primary partition.

Information similar to the following is displayed:

Select (default p): p

Partition number (2-4, default 2):

Partition number indicates the serial number of the primary partition.

Step 7 Ensure that entered partition number is the same as the partition had before. In this example, partition number **2** is used. Therefore, enter **2** and press **Enter**.

Information similar to the following is displayed:

Partition number (2-4, default 2): 2

First sector (209715200-482344959, default 209715200):

In the command output, **First sector** specifies the start sector.

□ NOTE

Data will be lost if the following operations are performed:

- Select a start sector other than the partition had before.
- Select an end sector smaller than the partition had before.
- **Step 8** Ensure that the entered start sector is the same as the partition had before. In this example, start sector **209715200** is recorded in **Step 1**. Therefore, enter **209715200** and press **Enter**.

Information similar to the following is displayed:

First sector (209715200-482344959, default 209715200):

Using default value 209715200

Last sector, +sectors or +size{K,M,G} (209715200-482344959, default 482344959):

In the command output, **Last sector** specifies the end sector.

Step 9 Ensure that the entered end sector is larger than or equal to the end sector recorded in **Step 1**. In this example, the recorded end sector is **314572799**, and the default end sector is used. Therefore, enter **482344959** and press **Enter**.

Information similar to the following is displayed:

Using default value 209715200

Last sector, +sectors or +size{K,M,G} (209715200-482344959, default 482344959):

Using default value 482344959

Partition 2 of type Linux and of size 130 GiB is set

Command (m for help):

The partition is created.

Step 10 Enter **p** and press **Enter** to view the partition details.

Information similar to the following is displayed:

Command (m for help): p

Disk /dev/vdb: 247.0 GiB, 246960619520 bytes, 482344960 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: 0x38717fc1

 Device Boot
 Start
 End
 Blocks
 Id
 System

 /dev/vdb1
 2048
 209715199
 104856576
 83
 Linux

/dev/vdb2 209715200 482344959 136314880 83 Linux

Command (m for help):

Step 11 Enter w and press **Enter** to write the changes to the partition table.

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.

In case that you want to discard the changes made before, you can exit fdisk by entering **q**.

Step 12 Run the following command to synchronize the new partition table to the OS:

partprobe

- **Step 13** Perform the following operations based on the file system of the disk:
 - For the ext*file system
 - a. Run the following command to check the correctness of the file system on the partition:

e2fsck -f Disk partition

In this example, run the following command:

e2fsck -f /dev/vdb2

Information similar to the following is displayed:

[root@ecs-test-0001 ~]# e2fsck -f /dev/vdb2

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/vdb2: 11/3276800 files (0.0% non-contiguous), 251790/13107200 blocks

b. Run the following command to extend the file system of the partition:

resize2fs Disk partition

In this example, run the following command:

resize2fs /dev/vdb2

Information similar to the following is displayed:

[root@ecs-test-0001 ~]# resize2fs /dev/vdb2

resize2fs 1.42.9 (28-Dec-2013)

Resizing the filesystem on /dev/vdb2 to 34078720 (4k) blocks.

The filesystem on /dev/vdb2 is now 34078720 blocks long.

If the error message "open: No such file or directory while opening /dev/vdb1" is returned, an incorrect partition is specified. Run **df** -**TH** to view the disk partitions.

c. (Optional) Run the following command to create a mount point:

Perform this step if you want to mount the partition on a new mount point.

mkdir Mount point

In this example, run the following command to create the /mnt/test mount point:

mkdir /mnt/test

d. Run the following command to mount the partition:

mount Disk partition Mount point

In this example, run the following command to mount partition /dev/vdb2 on /mnt/test:

mount /dev/vdb2 /mnt/test

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

- For the **xfs** file system
 - a. (Optional) Run the following command to create a mount point:
 Perform this step if you want to mount the partition on a new mount point.

mkdir Mount point

In this example, run the following command to create the /mnt/test mount point:

mkdir /mnt/test

b. Run the following command to mount the partition:

mount Disk partition Mount point

In this example, run the following command to mount partition /dev/vdb2 on /mnt/test:

mount /dev/vdb2 /mnt/test

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

c. Run the following command to extend the file system of the partition:

sudo xfs_growfs Disk partition

In this example, run the following command:

sudo xfs_growfs /dev/vdb2

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# sudo xfs growfs /dev/vdb2
                                      agcount=4, agsize=3276800 blks
meta-data=/dev/vdb2
                           isize=512
                    sectsz=512 attr=2, projid32bit=1
                    crc=1
     =
                              finobt=0, spinodes=0
                     bsize=4096 blocks=13107200, imaxpct=25
data
    =
                    sunit=0 swidth=0 blks
naming =version2
                         bsize=4096 ascii-ci=0 ftype=1
    =internal
                       bsize=4096 blocks=6400, version=2
log
                    sectsz=512 sunit=0 blks, lazy-count=1
                        extsz=4096 blocks=0, rtextents=0
realtime =none
data blocks changed from 13107200 to 34078720.
```

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

df -TH

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# df -TH
          Type Size Used Avail Use% Mounted on
Filesvstem
                  43G 1.9G 39G 5%/
/dev/vda1
           ext4
           devtmpfs 2.0G 0 2.0G 0% /dev
devtmpfs
          tmpfs 2.0G 0 2.0G 0% /dev/shm
tmpfs
          tmpfs
tmpfs
tmpfs
                 2.0G 9.1M 2.0G 1% /run
                 2.0G 0 2.0G 0% /sys/fs/cgroup
tmpfs
          tmpfs 398M 0 398M 0% /run/user/0
tmpfs
/dev/vdb1
                 106G 63M 101G 1% /mnt/sdc
           ext4
/dev/vdb2
           ext4 138G 63M 131G 1% /mnt/test
```

----End

Creating a New GPT Partition

Originally, data disk /dev/vdb has 100 GiB and one partition (/dev/vdb1), and then 50 GiB is added to the disk. The following procedure shows you how to create a new GPT partition /dev/vdb2 with this 50 GiB.

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

lsblk

Information similar to the following is displayed:

[root@ecs-test-0001 ~]# lsblk

NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda 253:0 0 40G 0 disk

__vda1 253:1 0 40G 0 part /
vdb 253:16 0 150G 0 disk

__vdb1 253:17 0 100G 0 part /mnt/sdc

Step 2 Run the following command to enter parted:

parted Disk

In this example, run the following command:

parted /dev/vdb

Information similar to the following is displayed:

[root@ecs-test-0001 ~]# parted /dev/vdb GNU Parted 3.1 Using /dev/vdb Welcome to GNU Parted! Type 'help' to view a list of commands. (parted)

- **Step 3** Enter **unit s** and press **Enter** to set the measurement unit of the disk to sector.
- **Step 4** Enter **p** and press **Enter** to view the disk partition information.

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/vdb 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: Virtio Block Device (virtblk)
Disk /dev/vdb: 314572800s
Sector size (logical/physical): 512B/512B
Partition Table: qpt
Disk Flags:
Number Start End
                        Size
                                  File system Name Flags
     2048s 209713151s 209711104s ext4
(parted)
```

In the command output, take note of the partition's end sector. In this example, the end sector of the /dev/vdb1 partition is 209713151s.

• If the following error information is displayed, enter **Fix**.

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)?

The GPT partition table information is stored at the start of the disk. To reduce the risk of damage, a backup of the information is saved at the end of the disk. When you expand the disk capacity, the end of the disk changes accordingly. In this case, enter **Fix** to move the backup file of the information to new disk end.

• If the following warning information is displayed, enter **Fix**.

Warning: Not all of the space available to /dev/vdb 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

Enter **Fix** as prompted. The system automatically sets the GPT partition style for the additional space.

Step 5 Run the following command and press **Enter**:

mkpart Partition name Start sector End sector

In this example, run the following command:

mkpart data 209713152s 100%

In this example, the additional space is used to create a new partition. In **Step 4**, the end sector of partition **dev/vdb1** is **209713151s**. Therefore, the start sector of the new partition **dev/vdb2** is set to **209713152s** and the end sector **100%**. This start and end sectors are for reference only. You can plan the number of partitions and partition size based on service requirements.

Information similar to the following is displayed:

(parted) mkpart data 209713152s 100% (parted)

The maximum sector can be obtained in either of the following ways:

- Query the disk's maximum end sector. For details, see **Step 2** to **Step 4**.
- Enter -1s or 100%, and the value displayed is the maximum end sector.

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

Information similar to the following is displayed:

(parted) p Model: Virtio Block Device (virtblk) Disk /dev/vdb: 314572800s Sector size (logical/physical): 512B/512B Partition Table: gpt Disk Flags: Number Start End Size File system Name Flags 209713151s 209711104s ext4 2048s test 209713152s 314570751s 104857600s data (parted)

Step 7 Enter **q** and press **Enter** to exit parted.

Information similar to the following is displayed:

(parted) q Information: You may need to update /etc/fstab.

You can configure automatic mounting by updating the /etc/fstab file. Before doing so, format the partition with a desired file system and mount the partition on the mount point.

Step 8 Run the following command to set the file system format for the new partition:

mkfs -t File system Disk partition

Sample command of the ext* file system:

mkfs -t ext4 /dev/vdb2

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# mkfs -t ext4 /dev/vdb2
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
3276800 inodes, 13107200 blocks
655360 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2162163712
400 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, 11239424
Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

• Sample command of the xfs file system:

mkfs -t xfs /dev/vdb2

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# mkfs -t xfs /dev/vdb2
meta-data=/dev/vdb2 isize=512 agcount=4, agsize=3276800 blks

= sectsz=512 attr=2, projid32bit=1
= crc=1 finobt=0, sparse=0
data = bsize=4096 blocks=13107200, imaxpct=25
= sunit=0 swidth=0 blks
naming =version2 bsize=4096 ascii-ci=0 ftype=1
log =internal log bsize=4096 blocks=6400, version=2
= sectsz=512 sunit=0 blks, lazy-count=1
realtime =none extsz=4096 blocks=0, rtextents=0
```

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 9 (Optional) Run the following command to create a mount point:

Perform this step if you want to mount the partition on a new mount point.

mkdir Mount point

In this example, run the following command to create the /mnt/test mount point:

mkdir /mnt/test

Step 10 Run the following command to mount the new partition:

mount Disk partition Mount point

In this example, run the following command to mount the new partition /dev/vdb2 on /mnt/test:

mount /dev/vdb2 /mnt/test

If the new partition is mounted on a directory that is not empty, the subdirectories and files in the directory will be hidden. Therefore, you are advised to mount the new partition on an empty directory or a new directory. If the new partition must be mounted on a directory that is not empty, move the subdirectories and files in this directory to another directory temporarily. After the partition is successfully 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-test-0001 ~]# df -TH
Filesystem
          Type
                 Size Used Avail Use% Mounted on
                  43G 1.9G 39G 5%/
/dev/vda1
           ext4
           devtmpfs 2.0G 0 2.0G 0% /dev
devtmpfs
          tmpfs 2.0G 0 2.0G 0% /dev/shm
tmpfs
tmpfs
          tmpfs
                 2.0G 9.1M 2.0G 1% /run
                 2.0G 0 2.0G 0% /sys/fs/cgroup
tmpfs
          tmpfs
tmpfs
          tmpfs
                 398M 0 398M 0% /run/user/0
/dev/vdb1
           ext4
                 106G 63M 101G 1% /mnt/sdc
                53G 55M 50G 1% /mnt/test
/dev/vdb2
           ext4
```

□ NOTE

If the server is restarted, the mounting will become invalid. You can set automatic mounting for partitions at system start by modifying the /etc/fstab file. For details, see Configuring Automatic Mounting at System Start.

----End

Extending an Existing GPT Partition

NOTICE

If the additional space is allocated to an existing partition, data on the disk will not be cleared but you must use **umount** to unmount the existing partition. In this case, services will be affected.

Originally, data disk /dev/vdb has 150 GiB and two partitions (/dev/vdb1 and /dev/vdb2), and then 80 GiB is added to the disk. The following procedure shows you how to add this 80 GiB to the existing GPT partition /dev/vdb2.

During an expansion, the additional space is added to the end of the disk. Therefore, if the disk has multiple partitions, the additional space can only be allocated to the partition at the disk end.

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

lsblk

Information similar to the following is displayed:

[root@ecs-test-0001 ~]# lsblk
NAME MAI:MIN RM SIZE RO TYPE MOUNTPOINT
vda 253:0 0 40G 0 disk
—vda1 253:1 0 40G 0 part /
vdb 253:16 0 230G 0 disk

```
-vdb1 253:17 0 100G 0 part /mnt/sdc
-vdb2 253:18 0 50G 0 part /mnt/test
```

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

- If the additional space is not included, refresh the capacity according to Extending Partitions and File Systems for SCSI Disks (Linux).
- If the additional space is included, go to **Step 2**.
- **Step 2** Run the following command to unmount the partition:

```
umount Disk partition
```

In this example, run the following command:

umount /dev/vdb2

Step 3 Run the following command to view the unmount result:

lsblk

```
Information similar to the following is displayed:
```

```
[root@ecs-test-0001 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda 253:0 0 40G 0 disk
-vda1 253:1 0 40G 0 part /
vdb 253:16 0 230G 0 disk
```

Step 4 Run the following command to enter parted:

parted Disk

In this example, run the following command:

parted /dev/vdb

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# parted /dev/vdb
GNU Parted 3.1
Using /dev/vdb
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.
- **Step 6** Enter **p** and press **Enter** to view the disk partition information.

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/vdb appears to be used, you can fix the GPT to use all of
the space (an extra 167772160
blocks) or continue with the current setting?
Fix/Ignore? Fix
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 482344960s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:
Number Start
                  End
                            Size
                                  File system Name Flags
```

2023-06-12 191

```
1 2048s 209713151s 209711104s ext4 test
2 209713152s 314570751s 104857600s ext4 data
(parted)
```

Take note of the start and end sectors of the /dev/vdb2 partition. These values will be used during the partition recreation. In this example, the partition's start sector is 209713152s. and its end sector is 314570751s.

• If the following error information is displayed, enter **Fix**.

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)?

The GPT partition table information is stored at the start of the disk. To reduce the risk of damage, a backup of the information is saved at the end of the disk. When you expand the disk capacity, the end of the disk changes accordingly. In this case, enter **Fix** to move the backup file of the information to new disk end.

If the following warning information is displayed, enter Fix.

Warning: Not all of the space available to /dev/vdb 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

Enter **Fix** as prompted. The system automatically sets the GPT partition style for the additional space.

Step 7 Enter **rm** and the partition number, and then press **Enter**. In this example, partition number **2** is used.

Information similar to the following is displayed:

(parted) rm Partition number? 2 (parted)

Step 8 Run the following command to recreate the partition and press **Enter**:

mkpart Partition name Start sector End sector

In this example, run the following command:

mkpart data 209713152s 100%

- Ensure that the entered start sector is the same as the partition had before. In this example, start sector 209713152s is recorded in Step 6. Therefore, enter 209713152s.
- Ensure that the entered end sector is greater than the partition had before. In this example, the end sector recorded in **Step 6** is **314570751s**, and all the additional space needs to be allocated to **dev/vdb2**. Therefore, enter **100%**.

Information similar to the following is displayed: (parted) mkpart data 209713152s 100%

MOTE

Data will be lost if the following operations are performed:

- Select a start sector other than the partition had before.
- Select an end sector smaller than the partition had before.

Step 9 Enter **p** and press **Enter** to view the partition information.

Information similar to the following is displayed:

```
(parted) p
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 482344960s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number Start End Size File system Name Flags
1 2048s 209713151s 209711104s ext4 test
2 209713152s 482342911s 272629760s ext4 data

(parted)
```

Step 10 Enter **q** and press **Enter** to exit parted.

```
Information similar to the following is displayed:
```

(parted) q Information: You may need to update /etc/fstab.

You can configure automatic mounting by updating the /etc/fstab file. Before doing so, format the partition with a desired file system and mount the partition on the mount point.

Step 11 Perform the following operations based on the file system of the disk:

- For the ext*file system
 - a. Run the following command to check the correctness of the file system on the partition:

e2fsck -f Disk partition

In this example, run the following command:

e2fsck -f /dev/vdb2

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# e2fsck -f /dev/vdb2
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/vdb2: 11/3276800 files (0.0% non-contiguous), 251790/13107200 blocks
```

b. Run the following command to extend the file system of the partition:

resize2fs Disk partition

In this example, run the following command:

resize2fs /dev/vdb2

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# resize2fs /dev/vdb2
resize2fs 1.42.9 (28-Dec-2013)
Resizing the filesystem on /dev/vdb2 to 34078720 (4k) blocks.
The filesystem on /dev/vdb2 is now 34078720 blocks long.
```


If the error message "open: No such file or directory while opening /dev/vdb1" is returned, an incorrect partition is specified. Run **df -TH** to view the disk partitions.

c. (Optional) Run the following command to create a mount point:
 Perform this step if you want to mount the partition on a new mount point.

mkdir Mount point

In this example, run the following command to create the /mnt/test mount point:

mkdir /mnt/test

d. Run the following command to mount the partition:

mount Disk partition Mount point

In this example, run the following command to mount partition /dev/vdb2 on /mnt/test:

mount /dev/vdb2 /mnt/test

∩ NOTE

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

• For the **xfs** file system

point.

a. (Optional) Run the following command to create a mount point:
 Perform this step if you want to mount the partition on a new mount

mkdir Mount point

In this example, run the following command to create the /mnt/test mount point:

mkdir /mnt/test

b. Run the following command to mount the partition:

mount Disk partition Mount point

In this example, run the following command to mount partition /dev/vdb2 on /mnt/test:

mount /dev/vdb2 /mnt/test

MOTE

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

c. Run the following command to extend the file system of the partition:

sudo xfs_growfs Disk partition

In this example, run the following command:

sudo xfs_growfs /dev/vdb2

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# sudo xfs_growfs /dev/vdb2
meta-data=/dev/vdb2 isize=512 agcount=4, agsize=3276800 blks
= sectsz=512 attr=2, projid32bit=1
= crc=1 finobt=0, spinodes=0
data = bsize=4096 blocks=13107200, imaxpct=25
```

```
= sunit=0 swidth=0 blks

naming =version2 bsize=4096 ascii-ci=0 ftype=1
log =internal bsize=4096 blocks=6400, version=2
= sectsz=512 sunit=0 blks, lazy-count=1
realtime =none extsz=4096 blocks=0, rtextents=0
data blocks changed from 13107200 to 34078720.
```

Step 12 Run the following command to view the mount result:

df -TH

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# df -TH
                   Size Used Avail Use% Mounted on
Filesystem
           Type
                   43G 1.9G 39G 5%/
/dev/vda1
           ext4
devtmpfs
            devtmpfs 2.0G 0 2.0G 0% /dev
          tmpfs 2.0G 0 2.0G 0% /dev/shm
tmpfs 2.0G 9.1M 2.0G 1% /run
tmpfs
tmpfs
tmpfs
          tmpfs 2.0G 0 2.0G 0% /sys/fs/cgroup
          tmpfs 398M 0 398M 0% /run/user/0
tmpfs
                   106G 63M 101G 1% /mnt/sdc
/dev/vdb1
            ext4
/dev/vdb2 ext4 138G 63M 131G 1% /mnt/test
```

----End

Configuring Automatic Mounting at System Start

The **fstab** file controls what disks are automatically mounted at startup. You can use **fstab** to configure your data disks to mount automatically. This operation will not affect the existing data.

The example here uses UUIDs to identify disks in the **fstab** file. You are advised not to use device names to identify disks in the file because device names are assigned dynamically and may change (for example, from /dev/vdb1 to /dev/vdb2) after a reboot. This can even prevent your disk from booting up.

MOTE

UUID is the unique character string for disk partitions in a Linux system.

Step 1 Query the partition UUID.

blkid Disk partition

In this example, the UUID of the /dev/vdb1 partition is queried.

blkid /dev/vdb1

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# blkid /dev/vdb1
/dev/vdb1: UUID="0b3040e2-1367-4abb-841d-ddb0b92693df" TYPE="ext4"
```

Carefully record the UUID, as you will need it for the following step.

Step 2 Open the **fstab** file using the vi editor.

vi /etc/fstab

- **Step 3** Press i to enter editing mode.
- **Step 4** Move the cursor to the end of the file and press **Enter**. Then, add the following information:

```
UUID=0b3040e2-1367-4abb-841d-ddb0b92693df /mnt/sdc ext4 defaults 0 2
```

In this example, the line starting with "UUID" is the information added. Edit this line to match the following format:

- UUID: The UUID obtained in **Step 1**.
- Mount point: The directory on which the partition is mounted. You can query the mount point using **df -TH**.
- Filesystem: The file system format of the partition. You can query the file system format using **df -TH**.
- Mount option: The partition mount option. Usually, this parameter is set to defaults.
- Dump: The Linux dump backup option.
 - **0**: Linux dump backup is not used. Usually, dump backup is not used, and you can set this parameter to **0**.
 - 1: Linux dump backup is used.
- fsck: The fsck option, which means whether to use fsck to check the disk during startup.
 - 0: not use fsck.
 - If the mount point is the root partition (/), this parameter must be set to 1.

If this parameter is set to **1** for the root partition, this parameter for other partitions must start with **2** because the system checks the partitions in the ascending order of the values.

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

The system saves the configurations and exits the vi editor.

Step 6 Verify that the disk is auto-mounted at startup.

1. Unmount the partition.

umount Disk partition

In this example, run the following command:

umount /dev/vdb1

2. Reload all the content in the /etc/fstab file.

mount -a

3. Query the file system mounting information.

mount | grep Mount point

In this example, run the following command:

mount | grep /mnt/sdc

If information similar to the following is displayed, automatic mounting has been configured:

root@ecs-test-0001 ~]# mount | grep /mnt/sdc /dev/vdb1 on /mnt/sdc type ext4 (rw,relatime,data=ordered)

----End

5.6.4 Extending Partitions and File Systems for SCSI Disks (Linux)

Scenarios

After a disk is expanded on the management console, the disk size is enlarged, but the additional space cannot be used directly.

In Linux, you must allocate the additional space to an existing partition or a new partition.

This section uses CentOS 7.4 64bit as the sample OS to describe how to extend an MBR partition of a SCSI data disk. The method for allocating the additional space varies with the server OS. This section is used for reference only. For detailed operations and differences, see the corresponding OS documents.

- Creating a New MBR Partition
- Extending an Existing MBR Partition

NOTICE

Performing the expansion operations with caution. Incorrect operations may lead to data loss or exceptions. So you are advised to back up the disk data using CBR or snapshots before expansion. For details about using CBR, see **Managing EVS Backups**. For details about using snapshots, see **Creating a Snapshot**.

Prerequisites

- You have expanded the disk capacity and attached the disk to a server on the management console. For details, see Expanding Capacity for an In-use EVS Disk or Expanding Capacity for an Available EVS Disk.
- You have logged in to the server.
 - For how to log in to an ECS, see the *Elastic Cloud Server User Guide*.
 - For how to log in to a BMS, see the Bare Metal Server User Guide.

Creating a New MBR Partition

Originally, data disk /dev/sda has 50 GiB and one partition (/dev/sda1), and then 50 GiB is added to the disk. The following procedure shows you how to create a new MBR partition /dev/sda2 with this 50 GiB.

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

fdisk -l

Information similar to the following is displayed:

[root@ecs-scsi ~]# fdisk -l

Disk /dev/vda: 42.9 GiB, 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: 0x000bcb4e

2023-06-12

```
Blocks Id System
 Device Boot
                Start
                          End
                2048 83886079 41942016 83 Linux
/dev/vda1 *
Disk /dev/sda: 107.4 GiB, 107374182400 bytes, 209715200 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: 0x915ffe6a
 Device Boot
                Start
                          End
                                 Blocks Id System
                2048 104857599 52427776 83 Linux
/dev/sda1
```

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

- If the additional space is not included, refresh the capacity according to Step
 2.
- If the additional space is included, go to **Step 3**.
- **Step 2** (Optional) Run the following command to update the capacity of the SCSI data disk:
 - 1. Run the following command to update the disk capacity on the server:

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

In the command, **%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

In this example, run the following command:

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

2. After the disk capacity is updated, run the following command to view the disk partition information again:

fdisk -l

If the additional space is included, go to **Step 3**.

Step 3 Run the following command to enter fdisk:

fdisk Disk

In this example, run the following command:

fdisk /dev/sda

Information similar to the following is displayed:

[root@ecs-scsi ~]# fdisk /dev/sda 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 4 Enter **n** and press **Enter** to create a new partition.

Information similar to the following is displayed:

```
Command (m for help): n
Partition type:
p primary (1 primary, 0 extended, 3 free)
e extended
Select (default p):
```

There are two types of disk partitions:

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

If the MBR partition style is used, a maximum of 4 primary partitions, or 3 primary partitions and 1 extended partition can be created. The extended partition cannot be used directly and must be divided into logical partitions before use.

Disk partitions created using GPT are not categorized.

Step 5 In this example, a primary partition is created. Therefore, enter **p** and press **Enter** to create a primary partition.

Information similar to the following is displayed:

Select (default p): p Partition number (2-4, default 2):

Partition number indicates the serial number of the primary partition. Because partition number 1 has been used, the value ranges from **2** to **4**.

Step 6 Enter the serial number of the primary partition and press Enter. Partition number2 is used in this example. Therefore, enter 2 and press Enter.

Information similar to the following is displayed:

Partition number (2-4, default 2): 2 First sector (104857600-209715199, default 104857600):

First sector indicates the start sector. The value ranges from **104857600** to **209715199**, and the default value is **104857600**.

Step 7 Enter the new partition's start sector and press **Enter**. In this example, the default start sector is used.

The system displays the start and end sectors of the partition's available space. You can customize the value within this range or use the default value. The start sector must be smaller than the partition's end sector.

Information similar to the following is displayed:

First sector (104857600-209715199, default 104857600):

Using default value 104857600

Last sector, +sectors or +size{K,M,G} (104857600-209715199, default 209715199):

Last sector indicates the end sector. The value ranges from **104857600** to **209715199**, and the default value is **209715199**.

Step 8 Enter the new partition's end sector and press **Enter**. In this example, the default end sector is used.

The system displays the start and end sectors of the partition's available space. You can customize the value within this range or use the default value. The start sector must be smaller than the partition's end sector.

Information similar to the following is displayed:

Last sector, +sectors or +size{K,M,G} (104857600-209715199, default 209715199): Using default value 209715199

2023-06-12

Partition 2 of type Linux and of size 50 GiB is set

Command (m for help):

Step 9 Enter **p** and press **Enter** to view the new partition.

Information similar to the following is displayed:

Command (m for help): p

Disk /dev/sda: 107.4 GiB, 107374182400 bytes, 209715200 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: 0x915ffe6a

Device Boot Start End Blocks Id System /dev/sda1 2048 104857599 52427776 83 Linux /dev/sda2 104857600 209715199 52428800 83 Linux

Command (m for help):

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

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.

In case that you want to discard the changes made before, you can exit fdisk by entering q.

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

partprobe

Step 12 Run the following command to set the file system format for the new partition:

mkfs -t File system Disk partition

Sample command of the ext* file system:

mkfs -t ext4 /dev/sda2

Information similar to the following is displayed:

[root@ecs-scsi ~]# mkfs -t ext4 /dev/sda2 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 3276800 inodes, 13107200 blocks

655360 blocks (5.00%) reserved for the super user

First data block=0

Maximum filesystem blocks=2162163712

400 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, 11239424

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

• Sample command of the xfs file system:

mkfs -t xfs /dev/sda2

Information similar to the following is displayed:

```
[root@ecs-scsi ~]# mkfs -t xfs /dev/sda2
meta-data=/dev/sda2
                           isize=512
                                      agcount=4, agsize=3276800 blks
                    sectsz=512 attr=2, projid32bit=1
     =
                    crc=1
                            finobt=0, sparse=0
                    bsize=4096 blocks=13107200, imaxpct=25
data =
                    sunit=0 swidth=0 blks
naming =version2
                        bsize=4096 ascii-ci=0 ftype=1
                       bsize=4096 blocks=6400, version=2
    =internal log
                    sectsz=512 sunit=0 blks, lazy-count=1
                       extsz=4096 blocks=0, rtextents=0
realtime =none
```

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 13 (Optional) Run the following command to create a mount point:

Perform this step if you want to mount the partition on a new mount point.

mkdir Mount point

In this example, run the following command to create the /mnt/test mount point:

mkdir /mnt/test

Step 14 Run the following command to mount the new partition:

mount Disk partition Mount point

In this example, run the following command to mount the new partition /dev/sda2 on /mnt/test:

mount /dev/sda2 /mnt/test

■ NOTE

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

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

df -TH

Information similar to the following is displayed:

```
[root@ecs-scsi ~]# df -TH
          Type Size Used Avail Use% Mounted on
Filesystem
/dev/vda1
           ext4
                 43G 2.0G 39G 5% /
           devtmpfs 509M 0 509M 0% /dev
devtmpfs
tmpfs
          tmpfs 520M 0 520M 0% /dev/shm
          tmpfs
                 520M 7.2M 513M 2% /run
tmpfs
tmpfs
          tmpfs
                 520M 0 520M 0% /sys/fs/cgroup
tmpfs
          tmpfs 104M 0 104M 0% /run/user/0
```

```
/dev/sda1 ext4 53G 55M 50G 1% /mnt/sdc
/dev/sda2 ext4 53G 55M 50G 1% /mnt/test
```


If the server is restarted, the mounting will become invalid. You can set automatic mounting for partitions at system start by modifying the /etc/fstab file. For details, see Configuring Automatic Mounting at System Start.

----End

Extending an Existing MBR Partition

NOTICE

If the additional space is allocated to an existing partition, data on the disk will not be cleared but you must use **umount** to unmount the existing partition. In this case, services will be affected.

Originally, SCSI data disk /dev/sda has 100 GiB and two partitions (/dev/sda1 and /dev/sda2), and then 50 GiB is added to the disk. The following procedure shows you how to add this 50 GiB to the existing MBR partition /dev/sda2.

During an expansion, the additional space is added to the end of the disk. Therefore, if the disk has multiple partitions, the additional space can only be allocated to the partition at the disk end.

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

fdisk -l

Information similar to the following is displayed:

```
[root@ecs-scsi ~]# fdisk -l
Disk /dev/vda: 42.9 GiB, 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: 0x000bcb4e
 Device Boot
                Start
                           End
                                  Blocks Id System
/dev/vda1 *
                2048 83886079 41942016 83 Linux
Disk /dev/sda: 161.1 GiB, 161061273600 bytes, 314572800 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: 0x915ffe6a
 Device Boot
                Start
                           End
                                 Blocks Id System
                2048 104857599 52427776 83 Linux
/dev/sda1
/dev/sda2
             104857600 209715199 52428800 83 Linux
```

In the command output, take note of the partition's start and end sectors. In this example, /dev/sda2's start sector is 104857600, and its end sector is 209715199.

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

If the additional space is not included, refresh the capacity according to Step
 2.

- If the additional space is included, take note of the start and end sectors of the target partition and then go to **Step 3**. These values will be used in the subsequent operations.
- **Step 2** (Optional) Run the following command to update the capacity of the SCSI data disk:
 - 1. Run the following command to update the disk capacity on the server:

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

In the command, %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

In this example, run the following command:

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

2. After the disk capacity is updated, run the following command to view the disk partition information again:

fdisk -l

If the additional space is included, take note of the start and end sectors of the target partition and then go to **Step 3**. These values will be used in the subsequent operations.

Step 3 Run the following command to unmount the partition:

umount Disk partition

In this example, run the following command:

umount /dev/sda2

Step 4 Run the following command to enter fdisk:

fdisk Disk

In this example, run the following command:

fdisk /dev/sda

Information similar to the following is displayed:

[root@ecs-scsi ~]# fdisk /dev/sda

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 5 Run the following command to delete the partition to be extended:

1. Enter **d** and press **Enter** to delete the partition.

Information similar to the following is displayed:

Command (m for help): d

Partition number (1,2, default 2):

2. Enter the partition number and press **Enter** to delete the partition. In this example, enter **2**.

Information similar to the following is displayed:

Partition number (1,2, default 2): 2 Partition 2 is deleted

Command (m for help):

□ NOTE

After deleting the partition, recreate the partition according to the following steps, and data on this disk will not be lost.

Step 6 Enter **n** and press **Enter** to create a new partition.

Information similar to the following is displayed:

Command (m for help): n
Partition type:
 p primary (1 primary, 0 extended, 3 free)
 e extended

Select (default p):

There are two types of disk partitions:

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

□ NOTE

If the MBR partition style is used, a maximum of 4 primary partitions, or 3 primary partitions and 1 extended partition can be created. The extended partition cannot be used directly and must be divided into logical partitions before use.

Disk partitions created using GPT are not categorized.

Step 7 Ensure that the entered partition type is the same as the partition had before. In this example, a primary partition is used. Therefore, enter **p** and press **Enter** to create a primary partition.

Information similar to the following is displayed:

Select (default p): p

Partition number (2-4, default 2):

Partition number indicates the serial number of the primary partition.

Step 8 Ensure that entered partition number is the same as the partition had before. In this example, partition number **2** is used. Therefore, enter **2** and press **Enter**.

Information similar to the following is displayed:

Partition number (2-4, default 2): 2

First sector (104857600-314572799, default 104857600):

In the command output, **First sector** specifies the start sector.

Ⅲ NOTE

Data will be lost if the following operations are performed:

- Select a start sector other than the partition had before.
- Select an end sector smaller than the partition had before.
- **Step 9** Ensure that the entered start sector is the same as the partition had before. In this example, start sector **104857600** is recorded in **Step 1** or **Step 2**. Therefore, enter **104857600** and press **Enter**.

Information similar to the following is displayed:

First sector (104857600-314572799, default 104857600):

Using default value 104857600

Last sector, +sectors or +size{K,M,G} (104857600-314572799, default 314572799):

In the command output, **Last sector** specifies the end sector.

Step 10 Ensure that the entered end sector is larger than or equal to the end sector recorded in **Step 1** or **Step 2**. In this example, the recorded end sector is **209715199**, and the default end sector is used. Therefore, enter **314572799** and press **Enter**.

Information similar to the following is displayed:

Last sector, +sectors or +size{K,M,G} (104857600-314572799, default 314572799):

Using default value 314572799

Partition 2 of type Linux and of size 100 GiB is set

Command (m for help):

The partition is created.

Step 11 Enter **p** and press **Enter** to view the partition details.

Information similar to the following is displayed:

Command (m for help): p

Disk /dev/sda: 161.1 GiB, 161061273600 bytes, 314572800 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: 0x915ffe6a

Device Boot Start End Blocks Id System

/dev/sda1 2048 104857599 52427776 83 Linux /dev/sda2 104857600 314572799 104857600 83 Linux

Command (m for help):

Step 12 Enter **w** and press **Enter** to write the changes to 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.

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.

□ NOTE

In case that you want to discard the changes made before, you can exit fdisk by entering \mathbf{q} .

Step 13 Run the following command to synchronize the new partition table to the OS:

partprobe

Step 14 Perform the following operations based on the file system of the disk:

For the ext*file system

a. Run the following command to check the correctness of the file system on the partition:

e2fsck -f Disk partition

In this example, run the following command:

e2fsck -f /dev/sda2

Information similar to the following is displayed:

[root@ecs-scsi ~]# e2fsck -f /dev/sda2

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/sda2: 11/3276800 files (0.0% non-contiguous), 251790/13107200 blocks

b. Run the following command to extend the file system of the partition:

resize2fs Disk partition

In this example, run the following command:

resize2fs /dev/sda2

Information similar to the following is displayed:

[root@ecs-scsi ~]# resize2fs /dev/sda2

resize2fs 1.42.9 (28-Dec-2013)

Resizing the filesystem on /dev/sda2 to 26214400 (4k) blocks.

The filesystem on /dev/sda2 is now 26214400 blocks long.

c. (Optional) Run the following command to create a mount point:

Perform this step if you want to mount the partition on a new mount point.

mkdir Mount point

In this example, run the following command to create the /mnt/test mount point:

mkdir /mnt/test

d. Run the following command to mount the partition:

mount Disk partition Mount point

In this example, run the following command to mount partition /dev/sda2 on /mnt/test:

mount /dev/sda2 /mnt/test

□ NOTE

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

For the xfs file system

a. (Optional) Run the following command to create a mount point:
 Perform this step if you want to mount the partition on a new mount point.

mkdir Mount point

In this example, run the following command to create the /mnt/test mount point:

mkdir /mnt/test

b. Run the following command to mount the partition:

mount Disk partition Mount point

In this example, run the following command to mount partition /dev/sda2 on /mnt/test:

mount /dev/sda2 /mnt/test

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

c. Run the following command to extend the file system of the partition:

sudo xfs_growfs Disk partition

In this example, run the following command:

sudo xfs_growfs /dev/sda2

Information similar to the following is displayed:

```
[root@ecs-scsi ~]# sudo xfs_growfs /dev/sda2
meta-data=/dev/sda2
                           isize=512 agcount=4, agsize=3276800 blks
                    sectsz=512 attr=2, projid32bit=1
                             finobt=0, spinodes=0
                    crc=1
data =
                    bsize=4096 blocks=13107200, imaxpct=25
                    sunit=0 swidth=0 blks
naming =version2
                        bsize=4096 ascii-ci=0 ftype=1
log =internal
                       bsize=4096 blocks=6400, version=2
                    sectsz=512 sunit=0 blks, lazy-count=1
realtime =none
                        extsz=4096 blocks=0, rtextents=0
data blocks changed from 13107200 to 26214400df.
```

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

df -TH

Information similar to the following is displayed:

```
[root@ecs-scsi ~]# df -TH
Filesystem
                 Size Used Avail Use% Mounted on
          Type
/dev/vda1
           ext4
                  43G 2.0G 39G 5% /
          devtmpfs 509M 0 509M 0% /dev
devtmpfs
tmpfs
         tmpfs 520M 0 520M 0% /dev/shm
tmpfs
         tmpfs
                 520M 7.2M 513M 2% /run
         tmpfs
                 520M 0 520M 0% /sys/fs/cgroup
tmpfs
                 104M 0 104M 0% /run/user/0
tmpfs
         tmpfs
                 53G 55M 50G 1% /mnt/sdc
/dev/sda1
          ext4
          ext4 106G 63M 101G 1% /mnt/test
/dev/sda2
```

----End

Configuring Automatic Mounting at System Start

The **fstab** file controls what disks are automatically mounted at startup. You can use **fstab** to configure your data disks to mount automatically. This operation will not affect the existing data.

The example here uses UUIDs to identify disks in the **fstab** file. You are advised not to use device names to identify disks in the file because device names are

assigned dynamically and may change (for example, from /dev/vdb1 to /dev/vdb2) after a reboot. This can even prevent your disk from booting up.

◯ NOTE

UUID is the unique character string for disk partitions in a Linux system.

Step 1 Query the partition UUID.

blkid Disk partition

In this example, the UUID of the /dev/vdb1 partition is queried.

blkid /dev/vdb1

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# blkid /dev/vdb1 /dev/vdb1: UUID="0b3040e2-1367-4abb-841d-ddb0b92693df" TYPE="ext4"
```

Carefully record the UUID, as you will need it for the following step.

Step 2 Open the **fstab** file using the vi editor.

vi /etc/fstab

- **Step 3** Press **i** to enter editing mode.
- **Step 4** Move the cursor to the end of the file and press **Enter**. Then, add the following information:

```
UUID=0b3040e2-1367-4abb-841d-ddb0b92693df /mnt/sdc ext4 defaults 0 2
```

In this example, the line starting with "UUID" is the information added. Edit this line to match the following format:

- UUID: The UUID obtained in **Step 1**.
- Mount point: The directory on which the partition is mounted. You can query the mount point using **df-TH**.
- Filesystem: The file system format of the partition. You can query the file system format using **df -TH**.
- Mount option: The partition mount option. Usually, this parameter is set to defaults.
- Dump: The Linux dump backup option.
 - **0**: Linux dump backup is not used. Usually, dump backup is not used, and you can set this parameter to **0**.
 - 1: Linux dump backup is used.
- fsck: The fsck option, which means whether to use fsck to check the disk during startup.
 - 0: not use fsck.
 - If the mount point is the root partition (/), this parameter must be set to **1**.

If this parameter is set to **1** for the root partition, this parameter for other partitions must start with **2** because the system checks the partitions in the ascending order of the values.

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

The system saves the configurations and exits the vi editor.

Step 6 Verify that the disk is auto-mounted at startup.

1. Unmount the partition.

umount Disk partition

In this example, run the following command:

umount /dev/vdb1

2. Reload all the content in the /etc/fstab file.

mount -a

3. Query the file system mounting information.

mount | grep Mount point

In this example, run the following command:

mount | grep /mnt/sdc

If information similar to the following is displayed, automatic mounting has been configured:

root@ecs-test-0001 ~]# mount | grep /mnt/sdc /dev/vdb1 on /mnt/sdc type ext4 (rw,relatime,data=ordered)

----End

2023-06-12

6 Detaching an EVS Disk

6.1 Detaching a System Disk

Scenarios

A system disk can only be detached offline, which means that its server must be in the **Stopped** state before the system disk is detached. To detach a system disk from a running server, you must first stop the server and then detach the disk.

□ NOTE

- For an attached system disk, the disk function is displayed as **System disk**, and the disk status is displayed as **In-use** in the disk list. After a system disk is detached from the server, the disk function changes to **Bootable disk**, and the status changes to **Available**.
- Bootable disks are the system disks detached from servers. A bootable disk can be reattached to a server and be used as a system disk or data disk depending on the disk function selected.

Procedure

- **Step 1** Log in to the management console.
- Step 2 Under Compute, click Elastic Cloud Server.

The **Elastic Cloud Server** page is displayed.

Step 3 In the server list, locate the row that contains the server whose system disk is to be detached, click **More** in the **Operation** column, and choose **Stop**.

When the server status changes to **Stopped**, the server has been stopped.

Step 4 Click the name of this server.

The server details page is displayed.

- **Step 5** Click the **Disks** tab to view the system disk attached to the server.
- **Step 6** Locate the row that contains the system disk and click **Detach**.

The **Detach Disk** dialog box is displayed.

Step 7 Click **Yes** to detach the disk.

After the operation had succeeded, the detached system disk is no longer displayed under the **Disks** tab.

- **Step 8** (Optional) Re-attach the bootable disk to a server. You can use it as a system disk or data disk depending on the disk function you select.
 - To re-attach and use it as a system disk, see Attaching an Existing System
 Disk.
 - To re-attach and use it as a data disk, see Attaching an Existing Non-Shared Disk

----End

Related Operations

For more detachment FAQs, see **Detachment**.

6.2 Detaching a Data Disk

Scenarios

Data disks can be detached online or offline, meaning that the server using the to-be-detached data disk can either be in the **Running** or **Stopped** state.

ECS

Detach a disk from a running server. For details, see **Storage** > **Detaching an EVS Disk from a Running ECS** in the *Elastic Cloud Server User Guide*.

BMS

SCSI disks can be attached to BMSs and used as data disks. You can detach a data disk either from a running or stopped BMS.

Ⅲ NOTE

For an attached data disk, the disk function is displayed as **Data disk**, and the disk status is displayed as **In-use** in the disk list. After the data disk has been detached from the server, the disk function remains unchanged, the disk status changes to **Available** for a non-shared data disk, and the disk status changes to **Available** for a shared data disk after it is detached from all its servers.

Precautions

Data may be lost after you detach an encrypted disk. For more information, see If I Detach a Disk, Will I Lose the Data on My Disk?.

Prerequisites

- Before detaching an EVS disk from a running Windows ECS, ensure that no programs are reading data from or writing data to the disk. Otherwise, data will be lost.
- Before detaching an EVS disk from a running Linux ECS, you must log in to the ECS and run the umount command to cancel the association between the

disk and the file system. In addition, ensure that no programs are reading data from or writing data to the disk. Otherwise, detaching the disk will fail.

Detaching a Non-shared Disk

- **Step 1** Log in to the management console.
- Step 2 Under Storage, click Elastic Volume Service.

The disk list page is displayed.

- **Step 3** Choose a way to detach the disk by determining whether you want to check server information first.
 - If yes, perform the following procedure:
 - a. In the disk list, click the name of the to-be-detached disk.

 The disk details page is displayed.
 - b. Click the **Servers** tab to view the server where the target disk has been attached.
 - c. Click to select the server and click **Detach Disk**.

 The **Detach Disk** dialog box is displayed.
 - d. Click **Yes** to detach the disk.
 - If no, perform the following procedure:
 - a. In the disk list, locate the row that contains the target disk and choose **More** > **Detach** in the **Operation** column.

The **Detach Disk** dialog box is displayed.

b. Click Yes to detach the disk.

The disk list is displayed. The disk status is **Detaching**, indicating that the disk is being detached from the server.

When the status changes to **Available**, the disk is successfully detached.

----End

Detaching a Shared Disk

- **Step 1** Log in to the management console.
- Step 2 Under Storage, click Elastic Volume Service.

The disk list page is displayed.

- **Step 3** Choose a way to detach the disk by determining whether you want to check server information first.
 - If yes, perform the following procedure:
 - a. In the disk list, click the name of the to-be-detached disk.

 The disk details page is displayed.
 - b. Click the **Servers** tab to view the servers where the target disk has been attached.

C.	Click to select the server and click Detach Disk .
	Shared EVS disks support batch detachment so that you can select multiple servers at a time.
	The Detach Disk dialog box is displayed.
d.	Click Yes to detach the disk.

- If no, perform the following procedure:
 - In the disk list, locate the row that contains the target disk and choose More > Detach in the Operation column.

The **Detach Disk** dialog box is displayed.

- b. Click to select the server. Shared EVS disks support batch detachment so that you can select multiple servers at a time.
- Click **Yes** to detach the disk.

The disk list page is displayed. The disk status is **Detaching**, indicating that the disk is being detached from the server.

If the shared EVS disk has been attached to multiple servers and needs to be detached from only some of its servers, the disk status will go back to In-use after the disk has been detached from the target servers. The disk status changes to Available only when it has been detached from all the servers.

----End

2023-06-12 213

Attaching an Existing Disk

7.1 Attaching an Existing System Disk

Scenarios

This section describes how to attach an existing system disk.

System disks can only be attached offline, which means that the server must be in the **Stopped** state.

You can view the disk function in the disk list. A disk can be attached to a server and used as the system disk only when its function is **Bootable disk** and its status is **Available**.

 Bootable disks are the system disks detached from servers. A bootable disk can be reattached to a server and be used as a system disk or data disk depending on the disk function selected.

Constraints

- Cloud servers created from ISO images are only used for OS installation. They
 have limited functions and cannot have EVS disks attached.
- A non-shared disk can be attached to one server only.
- The disk and the server must be in the same region and AZ.

Procedure

- **Step 1** Log in to the management console.
- Step 2 Under Storage, click Elastic Volume Service.

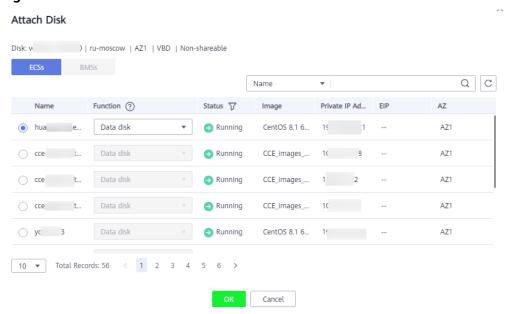
The disk list page is displayed.

Step 3 Locate the target disk in the list and click **Attach**.

Step 4 Select a server to attach the disk. Ensure that the disk and server are in the same AZ and that the server is in the **Stopped** state. After you select a the server, **System disk** is automatically entered as the disk function for the server.

One device name can be used for one disk only. For how to obtain the disk name in the OS, see section "How Do I Obtain My Disk Name in the ECS OS Using the Device Identifier Provided on the Console?" in the *Elastic Cloud Server User Guide*.

Figure 7-1 Attach Disk



Step 5 Click **OK** to return to the disk list page.

The status of the disk is **Attaching**, indicating that the disk is being attached to the server. When the function of the disk changes from **Bootable disk** to **System disk** and the disk status changes to **In-use**, the disk is successfully attached.

----End

Related Operations

For more attachment FAQs, see **Attachment**.

7.2 Attaching an Existing Non-Shared Disk

Scenarios

This section describes how to attach an existing non-shared disk to a server and use it as a data disk. A non-shared disk can be attached to one server only.

You can view the disk information in the disk list. A disk can be attached to a server and used as a data disk when all of the following conditions are met:

- Disk Sharing: Disabled
- Function: Bootable disk or Data disk

Status: Available

◯ NOTE

 Bootable disks are the system disks detached from servers. A bootable disk can be reattached to a server and be used as a system disk or data disk depending on the disk function selected.

Constraints

- Cloud servers created from ISO images are only used for OS installation. They have limited functions and cannot have EVS disks attached.
- A non-shared disk can be attached to one server only.
- The disk and the server must be in the same region and AZ.

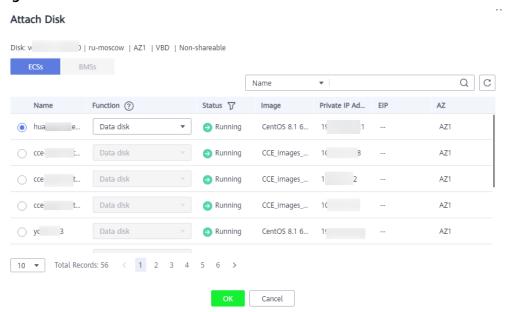
Procedure

- **Step 1** Log in to the management console.
- Step 2 Under Storage, click Elastic Volume Service.

The disk list page is displayed.

- **Step 3** Locate the target disk in the list and click **Attach**.
- **Step 4** Select a server to attach the disk. Ensure that the disk and server are in the same AZ. After you select a the server, **Data disk** is automatically entered as the disk function for the server.

Figure 7-2 Attach Disk



Step 5 Click **OK** to return to the disk list page.

The status of the disk is **Attaching**, indicating that the disk is being attached to the server. When the disk status changes to **In-use**, the disk is successfully attached.

Step 6 (Optional) Mount the existing disk partition to a mount point if you have attached the disk to a Linux server. The mount command is as follows:

mount Disk partition Mount point

----End

Related Operations

For more attachment FAQs, see Attachment.

7.3 Attaching an Existing Shared Disk

Scenarios

This section describes how to attach an existing shared disk to a server and use it as a data disk.

You can view the disk information in the disk list. A shared disk can be attached to servers and used as a data disk when all of the following conditions are met:

Disk Sharing: Enabled

• Function: Data disk

Status: In-use or Available

Constraints

NOTICE

If you simply attach a shared disk to multiple servers, files cannot be shared among them. Because there are no mutually agreed data read/write rules among servers, read and write operations from them may interfere with each other, or unpredictable errors may occur. To share files between servers, set up a shared file system or a clustered management system first.

- A shared disk can be attached to a maximum of 16 servers. These servers and the shared disk must be in the same AZ within a region.
- A shared, In-use disk can be attached to other servers only when the maximum number of servers that the disk can be attached to has not been reached.
- All the servers attached with a shared disk must run either Windows or Linux.
 For example, if you attach a shared disk to multiple Windows servers and then detach it from these servers, the shared disk cannot be attached to Linux servers later. This is because Windows and Linux support different file systems and cannot identify the original file system on the disk. Improper operations may damage the original file system.
- A shared disk can only be used as a data disk. It cannot be used as a system disk.
- Cloud servers created from ISO images are only used for OS installation. They have limited functions and cannot have EVS disks attached.

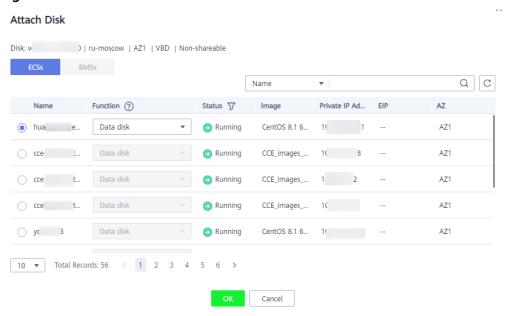
Procedure

- **Step 1** Log in to the management console.
- Step 2 Under Storage, click Elastic Volume Service.

The disk list page is displayed.

- **Step 3** Locate the target disk in the list and click **Attach**.
 - Shared disks support batch attachment so that you can attach a shared disk
 to multiple servers. The left area in the Attach Disk dialog box shows the
 server list. After you select the target servers, the selected servers will be
 displayed in the right area.
 - A shared disk can be attached to servers only when the disk status is **Available** or **In-use**.
- **Step 4** Select the target servers you want to attach the shared disk. Ensure that the disk and servers are in the same AZ. After you select the servers, **Data disk** is automatically entered as the disk function for each server.

Figure 7-3 Attach Disk



Step 5 Click **OK** to return to the disk list page.

The status of the disk is **Attaching**, indicating that the disk is being attached to the servers. When the disk status changes to **In-use**, the disk is successfully attached.

Step 6 (Optional) Mount the existing disk partition to a mount point if you have attached the disk to a Linux server. The mount command is as follows:

mount Disk partition Mount point

----End

Related Operations

For more attachment FAQs, see **Attachment**.

8 Deleting EVS Disks

Scenarios

If an EVS disk is no longer used, you can release the virtual resources by deleting it

- The disk status is Available, Error, Expansion failed, Restoration failed, or Rollback failed.
- The disk is not locked by any service.
- The shared disk has been detached from all its servers.

When a disk is deleted, EVS immediately destroys the metadata so that data can no longer be accessed. In addition, the physical storage space of the disk is reclaimed and cleared before being re-assigned. For any new disk created based on the re-assigned physical space, before data is written to the disk, EVS returns zero for all the read and write requests to the disk.

NOTICE

- When you delete a disk, all the disk data including the snapshots created for this disk will be deleted.
- A deleted disk cannot be recovered.

Procedure

- **Step 1** Log in to the management console.
- **Step 2** Under **Storage**, click **Elastic Volume Service**.

The disk list page is displayed.

- **Step 3** In the disk list, locate the row that contains the target disk, click **More** in the **Operation** column, and choose **Delete**.
- **Step 4** (Optional) If multiple disks are to be deleted, select and click **Delete** in the upper area of the list.

Step 5 In the displayed dialog box, confirm the information and click **Yes**.

----End

Related Operations

For more deletion FAQs, see **Deletion**.

9 Viewing EVS Disk Details

Scenarios

This section describes how to view disk details, including the disk status and specifications. You can view disk details:

- From the EVS Console
- From the Cloud Server Console

See EVS Disk Status to learn more about disk statuses.

From the EVS Console

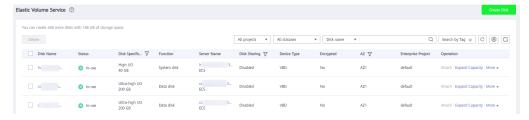
- **Step 1** Log in to the management console.
- Step 2 Under Storage, click Elastic Volume Service.

The disk list page is displayed.

Step 3 In the disk list, view disk information including the disk status, type, size, function, and device type.

You can filter disks by project, status, disk name, or tag.

Figure 9-1 Viewing the disk list



Step 4 In the disk list, locate the desired disk and click the disk name.

The disk details page is displayed for you to view the disk details.

Figure 9-2 Disk details page

Step 5 (Optional) Export disk information.

Click in the upper right corner of the list to export disk information.

----End

From the Cloud Server Console

- **Step 1** Log in to the management console.
- Step 2 Under Compute, click Elastic Cloud Server.

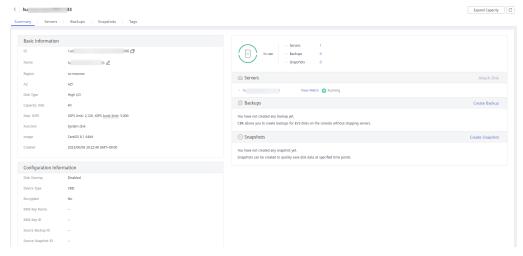
The **Elastic Cloud Server** page is displayed.

- **Step 3** In the server list, locate the desired server by server name and click the name.

 The server details page is displayed.
- **Step 4** On the **Disks** tab, click in front of the row containing the target disk. In the unfolded area, click the disk ID.

The disk details page is displayed for you to view the disk details.

Figure 9-3 Disk details page



----End

10 Managing Encrypted EVS Disks

Encryption Scenarios

System disk encryption

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

Table 10-1 Encryption relationship between images and system disks

Creating Server Using an Encrypted Image	Whether the System Disk Will Be Encrypted	Description
Yes	Yes	The encryption attribute of a systems disk is the same as that of the image. For details, see Managing Private Images > Encrypting Images in the Image Management Service User Guide.
No	No	-

• Data disk encryption

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

Table 10-2 Encryption relationship between backups, snapshots, images, and data disks

Where the Disk Is Created	Method of Creation	Whether the Data Disk Will Be Encrypted	Description
On the Cloud Server Console	Created together with the server	Yes/No	When a data disk is created together with a server, you can choose to encrypt the disk or not. For details, see Getting Started > Creating an ECS > Step 1: Configure Basic Settings in the Elastic Cloud Server User Guide.
Created Separately on the EVS console	No data source selected	Yes/No	When an empty disk is created, you can choose whether to encrypt the disk or not. The encryption attribute of the disk cannot be changed after the disk has been created.
	Creating from a backup	Yes/No	When a disk is created from a backup, you can choose whether to encrypt the disk or not. The encryption attributes of the disk and backup do not need to be the same.
			When you create a backup for a system or data disk, the encryption attribute of the backup will be the same as that of the disk.
	Creating from a snapshot (The snapshot's source disk is encrypted)	Yes	A snapshot created from an encrypted disk is also encrypted.
	Creating from a snapshot (The snapshot's source disk is not encrypted)	No	A snapshot created from a non-encrypted disk is not encrypted.

Where the Disk Is Created	Method of Creation	Whether the Data Disk Will Be Encrypted	Description
	Creating from an image (The image's source disk is encrypted)	Yes	-
	Creating from an image (The image's source disk is not encrypted)	No	-

Creating an Encrypted EVS Disk

Before you use the encryption function, KMS access rights need to be granted to EVS. If you have the Security Administrator permissions, grant the KMS access rights to EVS directly. If you do not have this permission, contact a user with the security administrator permissions to grant KMS access rights to EVS and then select the encryption option to create an encrypted disk.

For details about how to create an encrypted disk, see Create an EVS Disk.

Detaching an Encrypted EVS Disk

Before you detach a disk encrypted by a CMK, check whether the CMK is disabled or scheduled for deletion.

- If the CMK is available, the disk can be detached and re-attached, and data on the disk will not be lost.
- If the CMK is unavailable, the disk can still be used, but there is no guarantee for how long it will be usable. If the disk is detached, it will not be possible to re-attach it later. In this case, do not detach the disk without a working CMK.

The restoration method varies depending on the CMK status. For details, see **EVS Encryption**.

For details about how to detach an encrypted disk, see **Detaching a Data Disk**.

1 1 Managing Shared EVS Disks

How to Use Shared VBD and SCSI Disks

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

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

NOTICE

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

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

- The anti-affinity policy of an ECS group allows ECSs to be created on different physical servers to improve service reliability.
 For details about ECS groups, see Managing ECS Groups in the Elastic Cloud Server User Guide.
- The SCSI reservation mechanism uses a SCSI reservation command to perform SCSI reservation operations. If an ECS sends such a command to an EVS disk, the disk is displayed as locked to other ECSs, preventing the data damage that may be caused by simultaneous read/write operations to the disk from multiple ECSs.
- ECS groups and SCSI reservations have the following relationship: A SCSI reservation on a single EVS disk cannot differentiate multiple ECSs on the same physical host. For that reason, if multiple ECSs that use the same shared EVS disk are running on the same physical host, SCSI reservations

will not work properly. So you are advised to use SCSI reservations only on ECSs that are in the same ECS group, thus having a working antiaffinity policy.

Attaching a Shared EVS Disk

A non-shared EVS disk can only be attached to one server, whereas a shared EVS disk can be attached to up to 16 servers.

For details, see **Attaching a Shared Disk**.

Deleting a Shared EVS Disk

Because a shared EVS disk can be attached to multiple servers, ensure that the shared EVS disk is detached from all the servers before deletion.

For details, see **Deleting EVS Disks**.

Expanding a Shared EVS Disk

Shared EVS disks must be expanded when they are in the **Available** state. For details, see **Expanding Capacity for an In-use EVS Disk**.

Related Operations

For more disk sharing FAQs, see **Sharing**.

12 Managing EVS Backups

Scenarios

EVS backups are created using the CBR service. For details, see **Creating a Cloud Disk Backup** in the *Cloud Backup and Recovery User Guide*.

You can configure a backup policy for disks. With backup policies configured, data on EVS disks can be periodically backed up to improve data security.

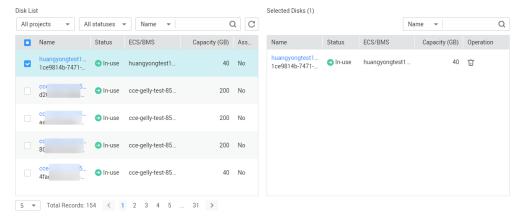
Constraints

Backups can be created only when the disks are in the Available or In-use state.

Creating a Disk Backup Vault and Applying a Backup Policy

- **Step 1** Log in to the CBR console.
 - 1. Log in to the management console.
 - 2. Choose Storage > Cloud Backup and Recovery > Cloud Disk Backups.
- Step 2 In the upper right corner, click Create Disk Backup Vault.
- **Step 3** (Optional) In the disk list, select the disks you want to back up. After disks are selected, they are added to the list of selected disks.

Figure 12-1 Selecting disks



□ NOTE

- Only Available and In-use disks can be selected.
- You can associate desired disks with the vault you created later if you skip this step.
- **Step 4** Specify a vault capacity ranging from the total sizes of disks to 10,485,760 GiB.
- **Step 5** Configure auto backup.
 - If you select **Configure**, you must then select an existing backup policy or create a new one. After the vault is created, the system applies this backup policy to the vault, and all disks associated with this vault will be automatically backed up based on this policy.
 - If you select **Skip**, disks associated with this vault are not automatically backed up.
- **Step 6** If you have subscribed to Enterprise Project, add the vault to an existing enterprise project.

An enterprise project facilitates project-level management and grouping of cloud resources and users. The default project is **default**.

Step 7 Specify a name for the vault.

The name can contain 1 to 64 characters including digits, letters, underscores (_), and hyphens (-), for example, **vault-612c**.

□ NOTE

You can use the default name, which is in the format of **vault_**xxxx.

- Step 8 Click Next.
- **Step 9** Complete the creation as prompted.
- **Step 10** Go back to the disk backup page. The vault you created is displayed in the list.

You can associate disks to the new vault or create backups for the disks. For details, see **Vault Management** in the *Cloud Backup and Recovery User Guide*.

----End

13 Managing EVS Snapshots

13.1 Snapshot Overview

What Is EVS Snapshot?

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

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

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

Application Scenarios

The snapshot function helps address your following needs:

Routine data backup

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

• Rapid data restoration

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

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

□ NOTE

When rolling back data from snapshots, data can only be rolled back to the original disk, and a rollback to a different disk is not possible.

Multi-service quick deployment

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

Operation Overview

You can create snapshots according to **Creating a Snapshot** to rapidly save the disk data at specified points in time.

If a data loss happened, you can roll back the disk data to the time when the snapshot was created based on **Rolling Back Data from a Snapshot**. In addition, you can create a new disk from the snapshot so that the disk will contain the snapshot data in the beginning. For details, see **Creating an EVS Disk from a Snapshot**.

When a snapshot is no longer needed, delete it according to **Deleting a Snapshot** to release the virtual resources.

Related Operations

For more snapshot FAQs, see **Snapshot**.

13.2 Creating a Snapshot

Scenarios

You can create an EVS snapshot on the management console to save the EVS disk data at a specific time point.

■ NOTE

Creating snapshots does not affect the performance of the disk.

Constraints

- A maximum of 7 snapshots can be created for one disk.
- Snapshots can be created for both system disks and data disks.
- Snapshots can be created only for available or in-use disks.
- Snapshots of encrypted disks are stored encrypted, and those of non-encrypted disks are stored non-encrypted.
- The enterprise project of a snapshot is the same as that of the snapshot's source disk.

Creating a Snapshot on the Disks Page

- **Step 1** Log in to the management console.
- Step 2 Under Storage, click Elastic Volume Service.

The disk list page is displayed.

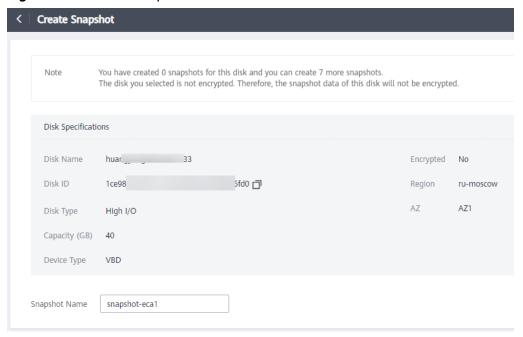
Step 3 In the disk list, locate the row that contains the target disk, click **Create Snapshot** in the **Operation** column.

Configure the basic settings for the snapshot according to Table 13-1.

Table 13-1 Snapshot parameter

Parameter	Description	Example Value
Snapshot Name	Mandatory	snapshot-01
	The name can contain a maximum of 64 characters.	

Figure 13-1 Create Snapshot



- Step 4 Click Create Now.
- **Step 5** Go back to the **Snapshots** page to view the snapshot creation information.

 After the snapshot status changes to **Available**, the snapshot has been created.

----End

Creating a Snapshot on the Snapshots Page

Step 1 Log in to the management console.

Step 2 Under **Storage**, click **Elastic Volume Service**.

The disk list page is displayed.

Step 3 In the navigation pane on the left, choose **Elastic Volume Service** > **Snapshots**.

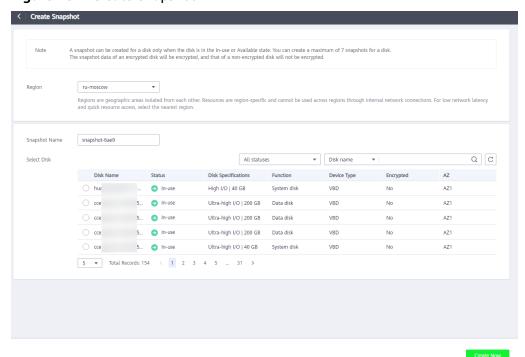
On the Snapshots page, click Create Snapshot.

Configure the basic settings for the snapshot according to Table 13-2.

Table 13-2 Snapshot parameters

Parameter	Description	Example Value
Region	Mandatory After you select a region, disks in the selected region will be displayed for you to choose from.	-
Snapshot Name	Mandatory The name can contain a maximum of 64 characters.	snapshot-01
Select Disk	Mandatory Select a disk based on which the snapshot will be created.	volume-01

Figure 13-2 Create Snapshot



Step 4 Click Create Now.

Step 5 Go back to the **Snapshots** page to view the snapshot creation information.

After the snapshot status changes to **Available**, the snapshot has been created.

----End

13.3 Deleting a Snapshot

Scenarios

If a snapshot is no longer needed, you can delete it to release the virtual resources.

Constraints

- The snapshot status must be Available or Error.
- If a disk is deleted, all the snapshots created for this disk will also be deleted.
- If a snapshot is deleted, disks rolled back and created from this snapshot are not affected.
- If you have reinstalled or changed the server OS, snapshots of the system disk are automatically deleted. Snapshots of the data disks can be used as usual.
- A snapshot whose name starts with autobk_snapshot_vbs_, manualbk_snapshot_vbs_, autobk_snapshot_csbs_, or manualbk_snapshot_csbs_ is automatically generated during backup. You can only view details of such snapshots and cannot delete them.

Procedure

Step 1	Log in to the management console.
Step 2	Under Storage , click Elastic Volume Service .
Step 3	In the navigation pane on the left, choose Elastic Volume Service > Snapshots .
	The snapshot list page is displayed.
Step 4	In the snapshot list, locate the row that contains the target snapshot and click Delete in the Operation column.
Step 5	(Optional) If multiple snapshots are to be deleted, select snapshot and click Delete in the upper area of the list.
Step 6	In the displayed dialog box, confirm the information and click Yes.
	If the snapshot disappears from the snapshot list, the snapshot is deleted successfully.
	End

13.4 Viewing Details of a Snapshot

Scenarios

This section describes how to view the details of a snapshot.

Procedure

- **Step 1** Log in to the management console.
- Step 2 Under Storage, click Elastic Volume Service.
- **Step 3** In the navigation pane on the left, choose **Elastic Volume Service** > **Snapshots**. The snapshot list page is displayed.
- **Step 4** Locate the row that contains the target snapshot and click in front of the snapshot name.

On the details page, view the snapshot information, such as the snapshot size.

----End

13.5 Rolling Back Data from a Snapshot

Scenarios

If data on an EVS disk is incorrect or damaged, you can roll back data from a snapshot to the source disk.

Constraints

- A snapshot can be rolled back only to its source disk. Rollback to another disk is not possible.
- A snapshot can be rolled back only when the snapshot status is Available and the source disk status is Available (not attached to any server) or Rollback failed
- A snapshot whose name starts with autobk_snapshot_vbs_, manualbk_snapshot_vbs_, autobk_snapshot_csbs_, or manualbk_snapshot_csbs_ is automatically generated during backup. Such a snapshot can only be viewed. It cannot be used to roll back the disk data.
- If you first roll back the snapshot to the original disk, you cannot use the snapshot to create a new disk then.

Procedure

- **Step 1** Log in to the management console.
- Step 2 Under Storage, click Elastic Volume Service.
- **Step 3** In the navigation pane on the left, choose **Elastic Volume Service** > **Snapshots**.

The snapshot list page is displayed.

- **Step 4** In the snapshot list, locate the row that contains the target snapshot and click **Roll Back Disk** in the **Operation** column.
- **Step 5** In the displayed dialog box, click **Yes**.

The snapshot list is displayed. After the snapshot status changes from **Rolling** back to **Available**, the data rollback is successful.

----End

13.6 Creating an EVS Disk from a Snapshot

Scenarios

This section describes how to create an EVS disk on the **Snapshots** page. Besides, you can also create an EVS disk from a snapshot by specifying the **Create from snapshot** parameter on the disk creation page. For details, see **Create an EVS Disk**.

Constraints

- The disk type, device type, and snapshot attributes of the new disk are the same as those of the snapshot's source disk.
- A maximum of 128 disks can be created from a snapshot.
- Batch disk creation is not possible, and the quantity parameter must be set to
 1.
- The AZ of the disk is the same as that of the snapshot's source disk and cannot be changed.
- A snapshot whose name starts with autobk_snapshot_vbs_, manualbk_snapshot_vbs_, autobk_snapshot_csbs_, or manualbk_snapshot_csbs_ is automatically generated during backup. Such a snapshot can only be viewed. It cannot be used to create new disks.
- If you first roll back the snapshot to the original disk, you cannot use the snapshot to create a new disk then.

Procedure

- **Step 1** Log in to the management console.
- Step 2 Under Storage, click Elastic Volume Service.
- **Step 3** In the navigation pane on the left, choose **Elastic Volume Service** > **Snapshots**.

 The snapshot list page is displayed.
- **Step 4** In the snapshot list, locate the row that contains the target snapshot and click **Create Disk** in the **Operation** column.
- **Step 5** Configure the disk parameters. For details, see parameter descriptions and operations provided in **Create an EVS Disk**.

MOTE

A maximum of 128 disks can be created from a snapshot.

If you create a disk from a snapshot, the disk capacity must be greater than or equal to the snapshot size. In the condition that you do not specify the disk capacity, if the snapshot size is smaller than 10 GiB, the default capacity 10 GiB will be used as the disk capacity; if the snapshot size is greater than 10 GiB, the snapshot size will be used as the disk capacity.

Step 6 Click Next.

Step 7 In the disk list, view the disk status.

When the disk status changes to **Available**, the disk is successfully created.

----End

14 Managing EVS Transfers

Scenarios

Through EVS transfer, EVS disks can be transferred from one account to another. After the transfer succeeds, the ownerships of the EVS disks belong to the target account only.

Users can use disk transfer via API only. For more information, see chapter "EVS Transfer" in the *Elastic Volume Service API Reference*.

Constraints

- EVS disks with backups and snapshots available cannot be transferred.
- EVS disks associated with backup policies cannot be transferred.
- EVS disks used as system disks cannot be transferred.

Procedure

The following example shows you how to transfer an EVS disk from account A to account B. User A belongs to account A, and user B belongs to account B. User A creates the transfer. User B accepts the transfer using the transfer ID (**transfer_id**) and authentication key (**auth_key**). After the transfer has been accepted, the transfer is complete. **Figure 14-1** shows the basic transfer process.

◯ NOTE

- **transfer_id** specifies the disk transfer ID. Each EVS disk transfer has a transfer ID, and user B uses this ID to accept the disk transfer. The transfer ID expires after user B accepts the transfer.
- **auth_key** specifies the identity authentication key of the disk transfer. Each EVS disk transfer has an authentication key, and user B uses this key for authentication when accepting the disk transfer.

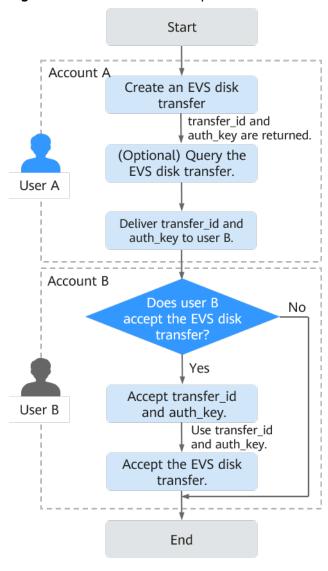


Figure 14-1 EVS disk transfer process

- **Step 1** User A creates the transfer. For details, see "Creating a Disk Transfer" in the *Elastic Volume Service API Reference*.
 - After the transfer is successfully created, transfer_id and auth_key are returned.
- **Step 2** (Optional) User A views the disk transfer. For details, see "Querying Details of a Disk Transfer" in the *Elastic Volume Service API Reference*. If multiple disk transfers have been created, user A can query all disk transfers. For details, see "Querying All Disk Transfers" or "Querying Details of All Disk Transfers" in the *Elastic Volume Service API Reference*.
- **Step 3** User A delivers the returned **transfer_id** and **auth_key** to user B.
- **Step 4** Check whether user B is going to accept the disk transfer.
 - If yes, go to **Step 5**.
 - If no, no further action is required.
 - User A can delete the unaccepted disk transfer. For details, see "Deleting a Disk Transfer" in the *Elastic Volume Service API Reference*.

- **Step 5** User B accepts **transfer_id** and **auth_key**.
- **Step 6** User B accepts the transfer through **transfer_id** and **auth_key**. For details, see "Accepting a Disk Transfer" in the *Elastic Volume Service API Reference*.

----End

15 Managing a Tag

15.1 Tag Overview

Tags identify EVS resources for purposes of easy categorization and quick search.

Table 15-1 Tag overview

Operation	Scenario
Adding a Tag	Add tags for existing disks or during disk creations.
Modifying a Tag	Change tag values for existing disks. Tag keys of existing disks cannot be changed.
Deleting a Tag	Delete tags that are no longer needed for existing disks.
Searching Disks by Tags	After tags are added, search for disks by tags.

15.2 Adding a Tag

Scenarios

You can add a tag for an existing EVS disk. You can also add tags during the disk creation. For details, see **Create an EVS Disk**.

- A tag is composed of a key-value pair.
 - A tag key can contain a maximum of 36 characters, including letters, digits, and underscores (_).
 - A tag value can contain a maximum of 43 characters, including letters, digits, underscores (_), periods (.), and hyphens (-).
- A maximum of 10 tags can be added for an EVS disk.
- Tag keys of the same EVS disk must be unique.

Procedure

- **Step 1** Log in to the management console.
- Step 2 Under Storage, click Elastic Volume Service.

The disk list page is displayed.

Step 3 In the disk list, locate the desired disk and click the disk name.

The disk details page is displayed.

- Step 4 Click the Tags tab.
- Step 5 Click Add Tag.

The Add Tag page is displayed.

- **Step 6** Enter a key and a value for a tag and click **OK**.
 - Key: This parameter is mandatory.
 - Value: This parameter is optional.

The **Tags** tab is displayed, and you can view the newly added tag.

----End

15.3 Modifying a Tag

Scenarios

You can change the value of a tag for an existing disk, but cannot change the key of a tag.

Procedure

- **Step 1** Log in to the management console.
- **Step 2** Under **Storage**, click **Elastic Volume Service**.

The disk list page is displayed.

Step 3 In the disk list, locate the desired disk and click the disk name.

The disk details page is displayed.

- **Step 4** Click the **Tags** tab.
- **Step 5** Locate the target tag and click **Edit** in the **Operation** column.

The **Edit Tag** page is displayed.

Step 6 Change the value of the tag and click **OK**.

Return to the tag list. If the tag value is changed, the modification is complete.

----End

15.4 Deleting a Tag

Scenarios

If an existing tag is no longer needed, you can delete it.

Procedure

- **Step 1** Log in to the management console.
- **Step 2** Under **Storage**, click **Elastic Volume Service**.

The disk list page is displayed.

Step 3 In the disk list, locate the desired disk and click the disk name.

The disk details page is displayed.

- Step 4 Click the Tags tab.
- **Step 5** Locate the target tag and click **Delete** in the **Operation** column.

The **Delete Tag** page is displayed.

Step 6 Confirm the information and click Yes.

The tag is deleted if it disappears from the tag list.

----End

15.5 Searching Disks by Tags

Scenarios

Tags can be used to categorize EVS disks, and users can quickly search for their desired EVS disks by tags. This section is used to guide users to search for EVS disk by existing tags.

Procedure

- **Step 1** Log in to the management console.
- Step 2 Under Storage, click Elastic Volume Service.

The disk list page is displayed.

Step 3 In the upper area of the disk list, click **Search by Tag**.

The **Search by Tag** page is displayed.

- **Step 4** Enter or select an existing tag in the text box under **Search by Tag**.
- **Step 5** (Optional) If disks containing multiple tags need to be queried, click to add tags.

A maximum of 10 tags can be added at a time.

For the added tags, you can delete them individually or click **Reset** to clear all of them.

Step 6 After the tags are added, click **Search**.

Disks owning the added tags are displayed in the list, and the search is complete.

----End

16 Changing EVS Disk Name

Scenarios

Disk names are used to identify disks. After a disk is created, you can perform operations in this section to change the disk name if needed.

Constraints

The name of a disk can be changed only when the disk status is **Available** or **Inuse**.

Procedure

- **Step 1** Log in to the management console.
- Step 2 Under Storage, click Elastic Volume Service.

The disk list page is displayed.

- **Step 3** Change the disk name in either of the following ways:
 - Perform the following steps to change the disk name in the disk list:
 - a. In the disk list, locate the target disk in the **Disk Name** column and click to the right of the disk name.

The Edit Disk Name dialog box is displayed.

- b. Enter a new name.
- c. Click **OK**.

After the change is successful, the new disk name is displayed in the disk list

- Perform the following steps to change the disk name on the disk details page:
 - a. In the disk list, locate the desired disk and click the disk name.

 The disk details page is displayed.
 - b. Click $\stackrel{\checkmark}{=}$ next to the disk name.
 - c. Enter a new name.
 - d. Click .

After the change is successful, the new disk name is displayed on the disk details page.

----End

1 Viewing EVS Monitoring Data

Description

This section describes monitored metrics reported by EVS to Cloud Eye as well as their namespaces and dimensions. You can use the console or APIs provided by Cloud Eye to query the metrics of the monitored objects and alarms generated for EVS.

Namespace

SYS.EVS

Metrics

Table 17-1 EVS metrics

Metric ID	Metric Name	Description	Value Range	Monitore d Object	Monitorin g Period
disk_de vice_re ad_byt es_rate	Disk Read Bandwidt h	Number of bytes read from the monitored disk per second Unit: Bytes/s	≥ 0 bytes/s	EVS disk	5 minutes in average
disk_de vice_wr ite_byt es_rate	Disk Write Bandwidt h	Number of bytes written to the monitored disk per second Unit: Bytes/s	≥ 0 bytes/s	EVS disk	5 minutes in average
disk_de vice_re ad_req uests_r ate	Disk Read IOPS	Number of read requests sent to the monitored disk per second Unit: Requests/s	≥ 0 Requests/ s	EVS disk	5 minutes in average

Metric	Metric	Description	Value	Monitore	Monitorin
ID	Name		Range	d Object	g Period
disk_de vice_wr ite_req uests_r ate	Disk Write IOPS	Number of write requests sent to the monitored disk per second Unit: Requests/s	≥ 0 Requests/ s	EVS disk	5 minutes in average

Dimension

Key	Value
disk_name	Server ID-drive letter, for example, 6f3c6f91-4b24-4e1b-b7d1-a94ac1cb011d-vda (vda is the drive letter)

Viewing Monitoring Data

- **Step 1** Log in to the management console.
- Step 2 Under Storage, click Elastic Volume Service.

The disk list page is displayed.

Step 3 In the EVS disk list, click the name of the disk you want to view the monitoring data.

The disk details page is displayed.

Step 4 On the **Servers** tab, locate the row that contains the server and click **View Metric** in the **Operation** column.

The Monitoring Metrics page is displayed.

Step 5 View the disk monitoring data by metric or monitored duration.

For more information about Cloud Eye, see the Cloud Eye User Guide.

----End

18 Viewing EVS Monitoring Data (Agent Installed and Simplified Monitoring Metrics Used)

Description

This section describes the EVS-related metrics included in the OS metrics supported by ECS. The agent of the latest version is used with simplified monitoring metrics.

After installing the agent on an ECS, you can view its EVS-related metrics included in the OS monitoring metrics.

For instructions about how to install and configure the agent, see section "Agent Installation and Configuration" in the *Cloud Eye User Guide*.

Monitoring Metrics

Table 18-1 EVS-related metrics

Metric	Name	Description	Valu e Ran ge	Monit ored Objec t	Monitori ng Period
mount PointPr efix_dis k_free	(Agent) Availabl e Disk Space	Available disk space on the monitored object Unit: GiB Linux: Run the df -h command to check the value in the Avail column. The path of the device name prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~). Windows: Obtain the metric value using the WMI API GetDiskFreeSpaceExW. The path of the device name prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).	≥0 GiB	ECS	5 minutes in average

Metric	Name	Description	Valu e Ran ge	Monit ored Objec t	Monitori ng Period
mount PointPr efix_dis k_used Percent	(Agent) Disk Usage	Percentage of total disk space that is used, which is calculated as follows: Disk Usage = Used Disk Space/ Disk Storage Capacity. Unit: Percent Linux: Obtain the metric value using the following formula: Disk Usage = Used Disk Space/Disk Storage Capacity. The path of the device name prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~). Windows: Obtain the metric value using the WMI API GetDiskFreeSpaceExW. The path of the device name prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).	0-10	ECS	5 minutes in average

Metric	Name	Description	Valu e Ran ge	Monit ored Objec t	Monitori ng Period
mount PointPr efix_dis k_ioUtil s and volume Prefix_ disk_io Utils	(Agent) Disk I/O Usage	Percentage of the time that the disk has had I/O requests queued to the total disk operation time Unit: Percent Linux: The disk I/O usage is calculated based on the data changes in the thirteenth column of the corresponding device in file /proc/diskstats in a collection period. The path of the device name prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~). Windows does not support this metric.	0-10 0%	ECS	5 minutes in average

Dimensions

Key	Value
instance_id	Specifies the ECS ID.

19 Auditing

Scenarios

EVS supports the recording of EVS operations through CTS. You can query EVS traces and use them for historical operation audits and backtracks.

Prerequisites

CTS has been enabled.

Key EVS Operations Recorded by CTS

Table 19-1 EVS operations that can be recorded by CTS

Operation	Resource	Trace
Create disk	evs	createVolume
Update disk	evs	updateVolume
Expand disk capacity	evs	extendVolume
Delete disk	evs	deleteVolume

Viewing Traces

To view audit logs, see **Querying Real-Time Traces** in the *Cloud Trace Service User Guide*.

20 Managing Quotas

20.1 Querying EVS Resource Quotas

Scenarios

Quotas are enforced for service resources on the platform to prevent unforeseen spikes in resource usage. Quotas can limit the number or amount of resources available to users, such as the number of EVS disks, the capacity of EVS disks, and the number of EVS snapshots.

Users can perform the following operations to view the resource quota details.

Procedure

- **Step 1** Log in to the management console.
- **Step 2** Click in the upper left corner and select the desired region and project.
- **Step 3** In the upper right corner of the page, click . .

The **Service Quota** page is displayed.

Step 4 View the used and total quota of each type of resources on the displayed page.

If a quota cannot meet service requirements, apply for a higher quota.

----End

20.2 Increasing EVS Resource Quotas

Scenarios

If your resource quotas fail to meet service requirements, perform the following operations to increase quotas.

How Do I Apply for a Higher Quota?

The system does not support online quota adjustment. If you need to adjust a quota, call the hotline or send an email to the customer service mailbox. Customer service personnel will timely process your request for quota adjustment and inform you of the real-time progress by making a call or sending an email.

Before dialing the hotline number or sending an email, make sure that the following information has been obtained:

- Account name, project name, and project ID, which can be obtained by performing the following operations:
 - Log in to the management console using the cloud account, click the username in the upper right corner, select **My Credentials** from the dropdown list, and obtain the account name, project name, and project ID on the **My Credentials** page.
- Quota information, which includes:
 - Service name
 - Quota type
 - Required quota

Learn how to obtain the service hotline and email address.

21 FAQ

21.1 Summary

General

- How Do I Start Using a Newly Created Disk?
- Can EVS Disks Be Used Directly for Storage?
- Can EVS Disks Be Used Alone?
- How Can I View My Disk Details?
- Can I Change the AZ of My Disk?
- Can I Change the Disk Type, Device Type, or Sharing Attribute of My Disk?
- What Should I Do If an Error Occurs on My EVS Disk?
- Why Do Some of My EVS Disks Not Have WWN Information?
- How Can I Migrate Data from an EVS Disk?
- What Are the Differences Between System Disks and Data Disks?
- How Can I Upload Files to My EVS Disk?

Capacity Expansion

- Can I Reduce or Temporarily Expand the Disk Capacity?
- What Are the Differences Between Expanding Capacity by Expanding an EVS Disk and Creating a New EVS Disk?
- Will My Disk Data Be Lost After I Expand the Disk Capacity?
- Can I Use Backups or Snapshots Created Before Capacity Expansion to Restore Data on Expanded Disks?
- Do I Need to Restart the Server After Expanding the Disk Capacity?
- Do I Need to Detach an EVS Disk Before Expanding Its Capacity?
- What Should I Do If My Disk Capacity Exceeds 2 TiB After Expansion?
- How Can I Allocate Newly Added Space to a New Partition?
- How Can I Allocate Newly Added Space to an Existing Partition?

- Why Did My Disk Capacity Remain Unchanged on the Server After Capacity Expansion?
- What Can't I Expand Capacity for My Disk?
- How Do I Extend the File System of an Unpartitioned Data Disk in Linux?
- How Do I Extend the Root Partition of a Quickly Provisioned BMS?

Attachment

- Why Can't I View the Attached Data Disk on the Server?
- Why Can't I Attach My Disk to a Server?
- Can I Attach a Disk to Multiple Servers?
- Can I Attach a Disk to a Server in Another AZ?
- How Can I Add a Data Disk to an Existing Server?
- Can I Attach Different Types of Disks to the Same Server?
- What Should I Do If a Linux EVS Disk Is Attached to a Windows Server?

Detachment

- If I Detach a Disk, Will I Lose the Data on My Disk?
- Why Can't I Detach My Disk?

Deletion

• How Can I Recover Data from a Disk That Was Accidentally Deleted?

Capacity

- What Is the Maximum Capacity Supported for the System and Data Disks?
- What Should I Do If My Disk Starts to Run Out of Space?
- What Can I Do If the Capacity of My Disk Reaches the Maximum But I Still Need More Space?
- What Should I Do If I Use fdisk to Initialize a Disk Larger Than 2 TiB and Then the Space in Excess of 2 TiB Cannot Be Displayed?
- How Can I View My Disk Usage?
- Can I Transfer the Data Disk Capacity to a System Disk?

Performance

- How Can I Test My Disk Performance?
- Why Does My Disk Performance Test Using Fio Have Incorrect Results?
- How Can I Handle a Slowdown in Disk Read/Write Speed and Increased I/O?

Sharing

- Do I Have to Deploy a Cluster to Use Shared Disks?
- How Many Servers Can I Attach a Shared Disk to?

- How Can I Attach a Shared Disk to Multiple Servers?
- Can I Attach a Shared Disk to Servers Running Different OSs?

Snapshot

- What Are the Typical Causes of a Snapshot Creation Failure?
- Does EVS Support Automatic Snapshot Creation?
- Can I Create Snapshots for Multiple Disks at a Time?
- How Is a Snapshot Created for My Disk?
- Why Can't I Roll Back My Disk Data from a Snapshot?
- Can I Roll Back the Data from a Snapshot After Reinstalling the OS or Formatting the Disk?
- How Is the Snapshot Size Calculated?
- Do Snapshots Take Space on the Disk?

Backup

- Do I Need to Stop the Server Before Performing a Backup?
- Can I Back Up and Restore My EVS Disk to a Different Region?

21.2 General

21.2.1 How Do I Start Using a Newly Created Disk?

A newly purchased disk must be attached to a server and then initialized in the server OS before you can use it.

For more information, see **Process Overview**.

21.2.2 Can EVS Disks Be Used Directly for Storage?

No.

EVS disks must be attached to cloud servers before use. You cannot use EVS disks alone to store data.

21.2.3 Can EVS Disks Be Used Alone?

No.

EVS disks must be attached to servers before you can use them.

21.2.4 How Can I View My Disk Details?

To do so, perform the following operations:

Step 1 Log in to the management console.

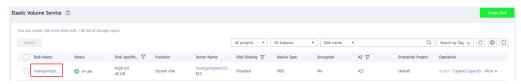
Step 2 Under Storage, click Elastic Volume Service.

The disk list page is displayed.

- **Step 3** Locate the row that contains the target disk and view the disk specifications, attributes, and billing mode.
- **Step 4** (Optional) Click the disk name to view more information, such as the disk backup and snapshot information.

View more information on the **Summary** tab.

Figure 21-1 Viewing disk details



----End

21.2.5 Can I Change the AZ of My Disk?

No.

The AZ of a disk cannot be changed after you have created the disk. If you want to change the AZ, delete the disk and create a new one.

21.2.6 Can I Change the Disk Type, Device Type, or Sharing Attribute of My Disk?

The following table describes whether the disk type, device type, sharing, and encryption attributes of a disk can be changed.

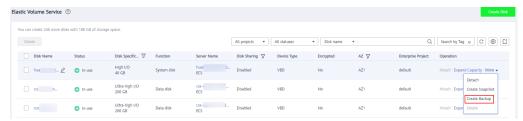
Table 21-1 EVS disk change description

Attribute	Allow Change	Change Direction
Disk type	No	 From high I/O to ultra-high I/O From ultra-high I/O to high I/O The preceding descriptions are examples only. EVS disk type cannot be changed.
Sharing	No	From shared to non-sharedFrom non-shared to shared
Device type	No	From SCSI to VBD From VBD to SCSI
Encryption	No	From non-encrypted to encryptedFrom encrypted to non-encrypted

However, you can:

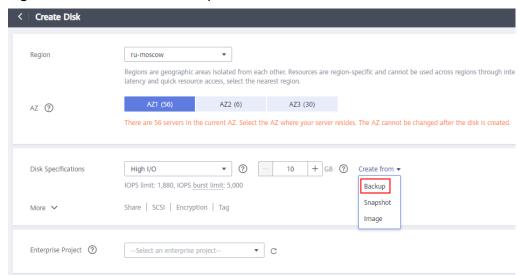
1. Create a backup for the disk.

Figure 21-2 Create Backup



2. Create a new disk from this backup. During the creation, select a new disk type and configure advanced settings (sharing, SCSI, and encryption) based on your service requirements.

Figure 21-3 Create from backup



21.2.7 What Should I Do If an Error Occurs on My EVS Disk?

If an error occurs, the disk may show one of the states listed in the following table. Take the measures described in the table to handle the exceptions.

Table 21-2 Solutions for disk errors

Error Status	Handling Suggestion
Error	Delete the disk in the Error state and create another one.
Deletion failed	Contact customer service.
Expansion failed	Our customer service personnel will contact you and help you handle this error. Do not perform any operations on the disk before the customer service personnel contact you. If the issue is urgent, you can contact our customer service personnel directly.

Error Status	Handling Suggestion
Restoratio n failed	Our customer service personnel will contact you and help you handle this error. Do not perform any operations on the disk before the customer service personnel contact you. If the issue is urgent, you can contact our customer service personnel directly.
Rollback failed	You can roll back the snapshot data to the disk again.

21.2.8 How Can I Obtain ECS NIC Information?

Scenarios

This topic shows how to obtain the ECS NIC information, such as the ECS IP address.

Procedure

- **Step 1** Log in to the management console.
- Step 2 Under Compute, click Elastic Cloud Server.

The ECS console is displayed.

Step 3 In the ECS list, click the target ECS name.

The ECS details page is displayed.

Step 4 Click the NICs tab.

The NIC details page is displayed.

- **Step 5** On the **NICs** tab, click to view ECS NIC information.
- **Step 6** To view the virtual IP address, click **Manage Virtual IP Address**.

The virtual IP address details page is displayed. You can find the bound virtual IP address based on the ECS NIC information.

----End

21.2.9 Why Do Some of My EVS Disks Not Have WWN Information?

EVS disks have two device types: VBD and SCSI. WWNs are used as the unique identifiers for SCSI EVS disks, and VBD EVS disks do not have WWNs.

You can view the WWN of a SCSI EVS disk on the management console. The details are as follows:

• If the SCSI EVS disk is brand new, you can view the disk WWN on the disk details page.

Figure 21-4 shows the query result.

Figure 21-4 Queried WWN information

WWN: ② 6888603000038430fa17a17502223655

 If the SCSI EVS disk was created before the WWN feature rollout, the disk WWN will fail to be obtained.

Figure 21-5 shows the query result.

Figure 21-5 No WWN information

WWN: (2) ---

21.2.10 How Can I Migrate Data from an EVS Disk?

Data migration involves the following scenarios:

- Cross-AZ data migration: Disk data can be migrated from one AZ to another through disk backups. You can create backups for your disks using the CBR service, and then use these backups to create new disks in the target AZ. For details, see sections "Creating a Cloud Disk Backup" and "Using a Backup to Create a Disk" in the Cloud Backup and Recovery User Guide.
- Cross-region data migration: You can create a data disk image from the data disk in the current region and replicate the image to the other region. Then you can use the data disk image to create data disks in that other region. For details, see section "Creating a Data Disk Image from an ECS" in the *Image Management Service User Guide*.
- Cross-account data migration: You can create a data disk image from the data disk under one account and then share the image with another account. Then you can use the data disk image under that second account to create data disks. For details, see section "Creating a Data Disk Image from an ECS" in the Image Management Service User Guide.

21.2.11 What Are the Differences Between System Disks and Data Disks?

- A system disk runs the server OS. It is like drive C in a PC.
 When a server is created, a system disk is automatically created and attached.
 You cannot create a system disk separately. The maximum size of a system disk is 1.024 GiB.
- Data disks store the server data. They are like drive D, drive E, and drive F in a PC

Data disks can be created during or after the server creation. If you create data disks during the server creation, the system will automatically attach the data disks to the server. If you create data disks after the server creation, you need to manually attach the data disks. The maximum size of a data disk is 32,768 GiB.

If one system disk already meets your business needs, you do not need to create data disks. As your business grows, you can create data disks when needed.

If the disk paths in your service systems cannot be changed or are difficult to change, you are advised to create data disks according to your system planning.

21.2.12 How Can I Download My EVS Disk Data to a Local PC?

EVS disk data cannot be directly saved to a local PC. It is recommended that you use a third-party tool, such as FTP, to download the data.

21.2.13 How Can I Export the Original Data After I Changed My Server OS from Windows to CentOS?

Solution:

- 1. Install the ntfsprogs software to enable Linux to access the NTFS file system. **yum install ntfsprogs**
- View the data disks previously attached to Windows.
 parted -l
- Mount the data disks.mount -t ntfs-3q Data disk path Mount point

21.2.14 What Are the Differences Between MBR and GPT Partition Styles?

Table 21-3 lists the common disk partition styles. In Linux, different partition styles require different partitioning tools.

Table 21-3 Disk partition styles

Disk Partition Style	Maximu m Disk Capacity Support ed	Maximum Number of Partitions Supported	Linux Partitioning Tool
Master Boot Record (MBR)	2 TiB	 4 primary partitions 3 primary partitions and 1 extended partition With MBR, you can create several primary partitions and one extended partition. The extended partition must be divided into logical partitions before use. For example, if 6 partitions need to be created, you can create them in the following two ways: 3 primary partitions and 1 extended partition, with the extended partition divided into 3 logical partitions 1 primary partition and 1 extended partition, with the extended partition divided into 5 logical partitions 	You can use either of the following tools: • fdisk • parted
Guid Partition Table (GPT)	18 EiB 1 EiB = 1048576 TiB	Unlimited Disk partitions created using GPT are not categorized.	parted

NOTICE

The maximum disk size supported by MBR is 2 TiB, and that supported by GPT is 18 EiB. Because an EVS data disk currently supports up to 32 TiB, use GPT if your disk size is larger than 2 TiB.

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

21.2.15 What Does the "reserveVolume" Trace Mean in CTS?

Before an EVS disk is attached, the system will call the reserveVolume EVS API to check whether the disk can be attached. If it can be attached, the system then changes the disk status to **attaching** to avoid conflicts with other operations.

21.2.16 How Can I Upload Files to My EVS Disk?

EVS disks must be attached to servers before you can use them. For how to upload files, see section "How Do I Upload Files to My ECS?" in the *Elastic Cloud Server FAQs*.

21.3 Capacity Expansion

21.3.1 Can I Reduce or Temporarily Expand the Disk Capacity?

No. The disk capacity can only be expanded, and temporary capacity expansion is not supported.

21.3.2 What Are the Differences Between Expanding Capacity by Expanding an EVS Disk and Creating a New EVS Disk?

The differences are as follows:

- Expanding an EVS disk is when you expand the capacity of an existing EVS disk. Some systems let you expand the capacity of EVS disks in use. In this case, services are not interrupted.
- If you create a new EVS disk and attach it to a server that already has an existing EVS disk, the new EVS disk and the original EVS disk are attached to the same server but independent from each other.

21.3.3 Will My Disk Data Be Lost After I Expand the Disk Capacity?

Data will not be deleted during a system disk or data disk capacity expansion. However, incorrect operations during an expansion may result in data loss or exceptions. Exercise caution when performing capacity expansions. You are advised to back up the disk data before expanding capacity.

- For details about disk backups, see the *Cloud Backup and Recovery User Guide*.
- For details about disk snapshots, see Creating a Snapshot.

21.3.4 Can I Use Backups or Snapshots Created Before Capacity Expansion to Restore Data on Expanded Disks?

Yes. If backups or snapshots have been created for disks before capacity was expanded, you can restore your disk data using these backups or snapshots after the capacity is expanded. Expansion operations do not affect backups and snapshots.

After the disk data is restored, the disk capacity is increased, but the additional space still needs to be partitioned and formatted before it can be used. You must log in to the server to extend the disk partition or file system.

To extend disk partitions and file systems, see the following sections:

- Extending Disk Partitions and File Systems (Windows)
- Extending Disk Partitions and File Systems (Linux)

21.3.5 Do I Need to Restart the Server After Expanding the Disk Capacity?

An EVS disk can be expanded either in the Available or In-use state. Expanding the disk capacity on the management console enlarges the disk capacity, but you still need to log in to the server and extend the disk partitions and file systems to make that additional space usable. You may need to restart the server during the partition and file system extension. The details are as follows:

- After expanding an In-use disk on the management console, log in to the server and view the disk capacity.
 - If the additional space can be viewed, you can extend the partition and file system and a restart is not required.
 - If the additional space cannot be viewed, the server OS may not be included in the compatibility list. In this case, you should stop and then start the server (do not restart the server). When the additional space can be viewed, extend the partition and file system.
- After expanding an Available disk on the management console, attach the
 disk to the server and extend the partition and file system on the server. In
 this case, a server restart is not required.

21.3.6 Do I Need to Detach an EVS Disk Before Expanding Its Capacity?

An expansion consists of two phases:

- 1. Expand the disk capacity on the management console.
 - A shared, in-use disk cannot be expanded. You must detach the shared disk from all its servers and then expand its capacity.
 - A non-shared, in-use disk can be expanded, and you can leave the disk attached during expansion as long as the following conditions are met:
 - The disk's server is in the Running or Stopped state.
 - The disk's server OS supports the expansion of In-use disks.

\sim		
	<i>(</i>)	-

Only some server OSs support capacity expansion of In-use disks. For details, see **Expanding Capacity for an In-use EVS Disk**.

- 2. Log in to the server and create a new partition or allocate the additional space to one that is already there.
 - In Windows, there are no partition extensions that require the partition to be unmounted first.
 - In Linux:
 - When allocating the additional space to an existing partition, that is, extending an existing partition, you must use the umount command to unmount the partition first.

When allocating the additional space to a new partition, that is, creating a new partition, you do not need to unmount the existing partition.

21.3.7 What Should I Do If My Disk Capacity Exceeds 2 TiB After Expansion?

An EVS system disk can be as large as 1 TiB (1,024 GiB). You can expand the capacity of a system disk to up to 1 TiB.

An EVS data disk can be as large as 32 TiB (32,768 GiB).

- With MBR, any disk space in excess of 2 TiB cannot be allocated and used, because the maximum disk capacity supported by MBR is 2 TiB (2,048 GiB). In this case, if you want to expand the disk capacity to over 2 TiB, change the partition style from MBR to GPT. Ensure that the disk data has been backed up before changing the partition style because services will be interrupted and data on the disk will be deleted during this change.
- With GPT, you can expand the capacity of a data disk to up to 32 TiB because the maximum disk capacity supported by GPT is 18 EiB (19,327,352,832 GiB). If the in-use partition style is GPT, see the following methods:
 - Windows:
 Extending Disk Partitions and File Systems (Windows Server 2016)
 - Linux:
 Extending Partitions and File Systems for Data Disks (Linux)

21.3.8 How Can I Allocate Newly Added Space to a New Partition?

Windows

In Windows, if your disk already has a D: drive and you want to create an E: drive, see System Disk: Add Additional Space to New Volume (F:) and Data Disk: Add Additional Space to New Volume (E:) in the following section:

Extending Disk Partitions and File Systems (Windows)

Linux

In Linux, if your disk already has, for example, partition /dev/vdb1 and you want to create a new partition, /dev/vdb2, and allocate the additional space to the new partition,

To extend disk partitions and file systems, see the following sections:

- Extending Partitions and File Systems for System Disks (Linux)
- Extending Partitions and File Systems for Data Disks (Linux)

21.3.9 How Can I Allocate Newly Added Space to an Existing Partition?

Windows

In Windows, if your disk already has, for example, a D: drive and you want to add space to this volume, see the following section:

See System Disk: Add Additional Capacity to New Volume (C:) and Data Disk: Add Additional Capacity to New Volume (D:) in the following section:

Extending Disk Partitions and File Systems (Windows Server 2016)

Linux

In Linux, if your disk already has, for example, partition /dev/vdb1 and you want to add space to this partition,

To extend disk partitions and file systems, see the following sections:

- Extending Partitions and File Systems for System Disks (Linux)
- Extending Partitions and File Systems for Data Disks (Linux)

21.3.10 Why Did My Disk Capacity Remain Unchanged on the Server After Capacity Expansion?

After expanding disk capacity on the management console, you must log in to the server and extend the disk partition and file system for the extra capacity to become available. Otherwise, you cannot view the additional space on the server.

To extend disk partitions and file systems, see the following sections:

- Extending Disk Partitions and File Systems (Windows Server 2008)
- Partition and File System Extension Preparations (Linux)

21.3.11 What Can't I Expand Capacity for My Disk?

Symptom

Capacity expansion is not allowed for the disk.

Troubleshooting

Possible causes are listed here in order of their probability.

If the fault persists after you have ruled out one cause, move on to the next one in the list.

Table 21-4 Troubleshooting

Possible Cause	Solution
A shared disk is still attached to a server.	See Shared Disk Is Still Attached to a Server.

Shared Disk Is Still Attached to a Server

Symptom: The **Expand Capacity** button is grayed out. When you attempt to click the capacity expansion button, the following hover tip is displayed: This operation can be performed only when the shared disk is in the Available state.

Solution: Detach the disk from all servers. If the **Expand Capacity** button becomes available, you can expand the disk capacity.

21.3.12 How Do I Extend the File System of an Unpartitioned Data Disk in Linux?

Scenarios

If no partition but only a file system is created on a data disk, extend the file system according to the following operations:

Run the **lsblk** command. Information similar to the following is displayed:

```
[root@ecs-test ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda 253:0 0 40G 0 disk

__vda1 253:1 0 40G 0 part /
vdb 253:16 0 60G 0 disk /mnt/sdc
```

In the command output, no partition but only a file system is created on data disk /dev/vdb.

In the following example, CentOS 7.4 64bit is used as the sample OS, data disk /dev/vdb has 10 GiB, no partition but only a file system is created on the disk, and additional 50 GiB has been added to this data disk on the management console. The following steps show how to extend this 50 GiB to the file system.

- Extending the EXT* File System
- Extending the XFS File System

The way you allocate additional space depends on the OS. This example is used for reference only. For the detailed operations and differences, see the corresponding OS documentations.

Extending the EXT* File System

Step 1 Run the following command to extend the file system:

resize2fs Disk name

In this example, run the following command:

resize2fs /dev/vdb

Information similar to the following is displayed:

```
[root@ecs-test ~]# resize2fs /dev/vdb
resize2fs 1.42.9 (28-Dec-2013)
Filesystem at /dev/vdb is mounted on /root/test; on-line resizing required
old_desc_blocs = 2, old_desc_blocs = 8
[17744.521535] EXT4-fs (vdb): resizing filesystem from 26214400 to 15728640 blocks
[17744.904470] EXT4-fs (vdb): resized filesystem to 15728640
The filesystem on /dev/vdb is now 15728640 blocks long.
```

Step 2 Run the following command to view the result:

df -TH

Information similar to the following is displayed:

```
[root@ecs-test ~]# df -TH
                Size Used Avail Use% Mounted on
Filesystem
         Type
/dev/vda1
          ext4
                43G 1.9G 39G 5% /
devtmpfs
          devtmpfs 2.0G 0 2.0G 0% /dev
         tmpfs 2.0G 0 2.0G 0% /dev/shm
tmpfs
tmpfs
         tmpfs 2.0G 9.1M 2.0G 1% /run
               tmpfs
         tmpfs
tmpfs
         tmpfs
         ext4 64G 55M 61G 1% /mnt/sdc
/dev/vdb
```

----End

Extending the XFS File System

Step 1 Run the following command to extend the file system:

```
xfs_growfs Disk name
```

In this example, run the following command:

xfs_growfs /dev/vdb

Information similar to the following is displayed:

```
[root@ecs-test ~]# xfs_growfs /dev/vdb
meta-data=/dev/vdb
                           isize=512
                                       agcount=4, agsize=655360 blks
                    sectsz=512 attr=2, projid32bit=1
                              finobt=0, spinodes=0
                    crc=1
data =
                     bsize=4096 blocks=2621440, imaxpct=25
                    sunit=0
                              swidth=0 blks
                          bsize=4096 ascii-ci=0 ftype=1
naming =version2
                        bsize=4096 blocks=2560, version=2
     =internal
log
                    sectsz=512 sunit=0 blks, lazy-count=1
realtime =none
                         extsz=4096 blocks=0, rtextents=0
data blocks changed from 2621440 to 15728640.
```

Step 2 Run the following command to view the result:

df -TH

Information similar to the following is displayed:

```
[root@ecs-test ~]# df -TH
Filesystem
           Type
                  Size Used Avail Use% Mounted on
                  40G 2.3G 35G 7%/
/dev/vda1
           ext4
           devtmpfs 1.9G 0 1.9G 0% /dev
devtmpfs
          tmpfs 1.9G 0 1.9G 0% /dev/shm
tmpfs
tmpfs
          tmpfs
                1.9G 8.6M 1.9G 1% /run
tmpfs
          tmpfs 1.9G 0 1.9G 0%/sys/fs/cgroup
```

```
tmpfs tmpfs 379M 0 379M 0% /run/user/0
/dev/vdb xfs 60G 34M 60G 1% /mnt/sdc
```

----End

21.3.13 How Do I Extend the Root Partition of a Quickly Provisioned BMS?

Scenarios

If the root partition of your quickly provisioned BMS is too small, extend the root partition by referring to the following procedure.

This example uses CentOS 7.3 and system disk, /dev/sdf. The way you allocate additional space depends on the OS. This example is used for reference only. For detailed operations and differences, see the corresponding OS documentations.

In this example, the initial size of the BMS system disk (**sdf**) is 40 GiB and needs to be expanded to 140 GiB. The initial disk partitions are as follows:

```
8:80
                      40G
                0
sdf 1
        8:81
                     500M
                              part /boot
        8:82
                0
                       5G
sdf2
                           0
                             part [SWAP]
        8:83
                0
                   34.5G
sdf3
                           0
                             part /
        8:84
                0
                     64M
```

Procedure

- **Step 1** Log in to the EVS console and expand the system disk capacity to 140 GiB.
- **Step 2** Log in to the BMS and run the following command to view the system disk capacity:

lsblk

Information similar to the following is displayed:

```
8:80
                     140G
                              part /boot
sdf 1
        8:81
                0
                     500M
                            0
                0
-sdf2
        8:82
                       5G
                            0
                              part [SWAP]
sdf3
        8:83
                0
                    34.5G
                            0 part /
sdf4
        8:84
                0
                      64M
```

The system disk (**sdf**) has been expanded from 40 GiB to 140 GiB. The **sdf4** partition (64 MiB) is the configurition that stores the BMS configuration information.

Step 3 Run the following command to back up the configdriver partition:

dd if=/dev/sdf4 of=/root/configdriver.img

Information similar to the following is displayed:

```
[root@bms-6acd ~]# dd if=/dev/sdf4 of=/root/configdriver.img
131072+0 records in
131072+0 records out
67108864 bytes (67 MB) copied, 0.291739 s, 230 MB/s
```

Step 4 Run the following command and delete the configdriver partition:

fdisk /dev/sdf

```
Iroot@bms-6acd ~]# fdisk /dev/sdf
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
Partition number (1-4, default 4): 4
Partition 4 is deleted

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 5 Run the **partprobe** command to update the partition information.

If partition configdriver has been deleted, information similar to the following is displayed:

```
140G
sdf
         8:80
                 0
                             0 disk
 sdf 1
         8:81
                 0
                      500M
                             0
                               part /boot
 sdf2
         8:82
                 0
                        5G
                             0 part [SWAP]
                     34.5G
  sdf3
         8:83
                 0
                             0 part /
                     97.7M
         8:84
                 0
                            0 part
```

Step 6 Recreate the configdriver partition with 100 MB.

If the available sectors range from 83755008 to 293601279, set 293401279 (293601279 – 200000) as the new partition's start sector and 293601279 (default value) as the end sector.

```
Command (m for help): n
Partition type:
    p primary (3 primary, 0 extended, 1 free)
    e extended
Select (default e): p
Selected partition 4
First sector (83755008-293601279, default 83755008): 293401279
Last sector, +sectors or +size{K,M,G} (293401279-293601279, default 293601279):
Using default value 293601279
Partition 4 of type Linux and of size 97.7 MiB is set

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.
```

Run the **partprobe** command to update the partition information.

sdf	8:80	0	140G	0	disk	
⊢sdf 1	8:81	0	500M	0	part	∕boot
-sdf2	8:82	0	5G	0	part	[SWAP]
-sdf3	8:83	0	34.5G	0	part	/
∟ <mark>sdf4</mark>	8:84	0	97.7M	0	part	

Step 7 Run the following command to extend the root partition:

growpart /dev/sdf 3

Information similar to the following is displayed:

```
[root@bms-6acd ~1# growpart /dev/sdf 3
CHANGED: partition=3 start=11511808 old: size=72243200 end=83755008 new:size=28
1889471,end=293401279
```

Run the **Isblk** command to view the new root partition size.

```
8:80
                     140G
sdf 1
        8:81
                0
                     500M
                            0 part /boot
                            0 part [SWAP]
sdf2
        8:82
                0
                       5G
sdf3
        8:83
                0
                  134.4G
                            0
                              part
sdf4
        8:84
                    97.7M
                0
                            0
                              part
```

Step 8 Run the following command to extend the file system of the root partition:

resize2fs /dev/sdf3

Information similar to the following is displayed:

```
[root@bms-6acd ~1# resize2fs /dev/sdf3
resize2fs 1.42.9 (28-Dec-2013)
Filesystem at /dev/sdf3 is mounted on /; on-line resizing required
old_desc_blocks = 5, new_desc_blocks = 17
The filesystem on /dev/sdf3 is now 35236183 blocks long.
You have new mail in /var/spool/mail/root
[root@bms-6acd ~]# df -h
                 Size Used Avail Use% Mounted on
Filesystem
/dev/sdf3
                133G
                        1.9G 125G
                                       2% /
devtmpfs
                  63G
                           0
                                63G
                                       0% /dev
tmpfs
                  63G
                           0
                                63G
                                       0% /dev/shm
                  63G
                                63G
tmpfs
                        9.0M
                                       1% /run
tmpfs
                  63G
                                63G
                           0
                                       0% /sys/fs/cgroup
/dev/sdf1
                 477M
                               344M
                                      24% /boot
                        104M
tmpfs
                   13G
                           0
                                13G
                                       0% /run/user/0
```

Step 9 Run the following command to restore the content of the configdriver partition:

dd if=/root/configdriver.img of=/dev/sdf4

Information similar to the following is displayed:

```
[root@bms-6acd ~]# dd if=/root/configdriver.img of=/dev/sdf4]
131072+0 records in
131072+0 records out
67108864 bytes (67 MB) copied, 0.372614 s, 180 MB/s
[root@bms-6acd ~]# __
/dev/sdf1: UUID="b9c472f9-6737-4200-910a-efa3af16629a" TYPE="ext4"
/dev/sdf2: UUID="b07ff4d0-8b0b-4c43-a40a-0b27290ea215" TYPE="swap"
/dev/sdf3: UUID="1e57f71e-6adc-4e98-9407-0f7d678d4525" TYPE="ext4"
/dev/sdf4: UUID="2018-09-27-19-13-01-00" LABEL="config-2" TYPE="iso9660"
[root@bms-6acd ~]#
```

The root partition of the quickly provisioned BMS has been extended.

----End

21.4 Attachment

21.4.1 Why Can't I View the Attached Data Disk on the Server?

Troubleshooting

Table 21-5 Possible causes

os	Possible Cause	Solution
Linux	New data disks are not formatted and partitioned by default, and an unformatted disk will not be listed in the command output. You must manually initialize the disk.	Linux Data Disk
	 If a data disk cannot be found after the server is restarted, automatic partition mounting at system start may not be configured. 	
Wind ows	New data disks are not formatted and partitioned by default. Only formatted and partitioned drives show up in the resource manager. You must manually initialize the disk.	Windows Data Disk

Linux Data Disk

Symptom: A data disk has been attached to a Linux server on the management console, but the disk cannot be viewed on the server.

Run **df -TH** to view the disk information. CentOS 7.4 is used in this example. The normal command output is as follows:

```
[root@ecs-test-0001 ~]# df -TH
          Type Size Used Avail Use% Mounted on
Filesystem
/dev/vda1
           ext4
                 43G 1.9G 39G 5%/
devtmpfs
           devtmpfs 2.0G 0 2.0G 0% /dev
         tmpfs 2.0G 0 2.0G 0% /dev/shm
tmpfs
                 2.0G 9.1M 2.0G 1% /run
tmpfs
          tmpfs
tmpfs
          tmpfs
                 2.0G 0 2.0G 0% /sys/fs/cgroup
                 398M 0 398M 0% /run/user/0
tmpfs
          tmpfs
/dev/vdb1 ext4
                106G 63M 101G 1% /mnt/sdc
```

Unlike the normal command output, only system disk /dev/vda1 is visible, but data disk /dev/vdb1 is missing from the command output.

Cause Analysis:

• Cause 1: New data disks are not formatted and partitioned by default, and an unformatted disk will not be listed in the command output. You must manually initialize the disk.

For details, see Introduction to Data Disk Initialization Scenarios and Partition Styles.

- Cause 2: If a data disk cannot be found after the server is restarted, automatic partition mounting at system start may not be configured. Perform the following steps:
 - a. Run the following command to mount the partition again:

mount Disk partition Mount point

In this example, run the following command:

mount /dev/vdb1 /mnt/sdc

Perform the following steps to enable automatic partition mounting at system start:

b. Run the following command to guery the partition UUID:

blkid Disk partition

In this example, run the following command to query the UUID of the /dev/vdb1 partition:

blkid /dev/vdb1

Information similar to the following is displayed:

[root@ecs-test-0001 ~]# blkid /dev/vdb1 /dev/vdb1: UUID="0b3040e2-1367-4abb-841d-ddb0b92693df" TYPE="ext4"

The UUID of the /dev/vdb1 partition is displayed.

c. Run the following command to open the **fstab** file using the vi editor:

vi /etc/fstab

Press i to enter editing mode.

d. Move the cursor to the end of the file and press Enter. Then, add the following information:

UUID=0b3040e2-1367-4abb-841d-ddb0b92693df /mnt/sdc

ext4 defaults 0.2

The preceding content is used for reference only. Add the information that is used in the environment. The parameters are described as follows:

- The first column indicates the partition UUID obtained in b.
- The second column indicates the directory on which the partition is mounted. You can query the mount point using the df -TH command.
- The third column indicates the file system format of the partition. You can query the file system format using the df -TH command.
- The fourth column indicates the partition mount option. Normally, this parameter is set to **defaults**.
- The fifth column indicates the Linux dump backup option.
 - **0**: Linux dump backup is not used. Normally, dump backup is not used, and you can set this parameter to **0**.
 - 1: Linux dump backup is used.

- The sixth column indicates the fsck option, that is, whether to use fsck to check the attached disk during startup.
 - o **0**: not use fsck.
 - If the mount point is the root partition (/), this parameter must be set to 1.

When this parameter is set to 1 for the root partition, this parameter for other partitions must start with 2 because the system checks the partitions in the ascending order of the values.

e. Press **Esc**, enter :wq, and press **Enter**.

The system saves the configurations and exits the vi editor.

Perform the following operations to verify the automatic mounting function:

i. Run the following command to unmount the partition:

umount Disk partition

In this example, run the following command:

umount /dev/vdb1

ii. Run the following command to reload all the content in the /etc/fstab file:

mount -a

iii. Run the following command to query the file system mounting information:

mount | grep Mount point

In this example, run the following command:

mount | grep /mnt/sdc

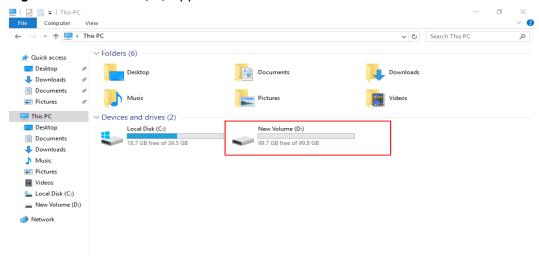
If information similar to the following is displayed, automatic mounting has been configured:

root@ecs-test-0001 ~]# mount | grep /mnt/sdc /dev/vdb1 on /mnt/sdc type ext4 (rw,relatime,data=ordered)

Windows Data Disk

Symptom: A data disk has been attached to a Windows server on the management console, but the disk cannot be viewed on the server. For example, Volume (D:) was not shown in **This PC** of a Windows server running Windows Server 2012. Normally, Volume (D:) appears, as shown in **Figure 21-6**.

Figure 21-6 Volume (D:) appears



Solution: New data disks are not formatted and partitioned by default. Only formatted and partitioned drives show up in **This PC**. You must manually initialize the disk before it can be viewed here.

For details, see Introduction to Data Disk Initialization Scenarios and Partition Styles.

21.4.2 Why Can't I Attach My Disk to a Server?

Symptom

My disk cannot be attached to a server.

Troubleshooting

Possible causes are listed here in order of their probability.

If the fault persists after you have ruled out one cause, move on to the next one in the list.

Table 21-6 Troubleshooting

Symptom	Solution
The target server on the Attach Disk page could not be found.	• Go to Check Whether the Disk and Server Are in the Same AZ.
	 Cloud servers created from ISO images are only used for OS installation. They have limited functions and cannot have EVS disks attached.
The Attach button is grayed out.	Go to Maximum Number of Disks That Can Be Attached to the Server Has Been Reached.

Symptom	Solution
An incorrect OS type warning is displayed when a shared disk is attached.	Go to Check Whether the Servers Attached with the Shared Disk Are Running the Same Type of OS.

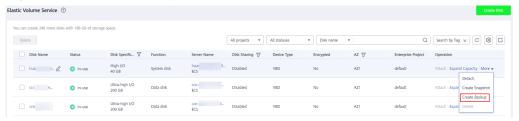
Check Whether the Disk and Server Are in the Same AZ

Symptom: After you click **Attach**, the target server cannot be found on the **Attach Disk** page.

Solution: A disk can only be attached to a server in the same AZ and region. The **Attach Disk** page filters and shows all the servers that the disk can be attached to. Determine whether your disk data is required.

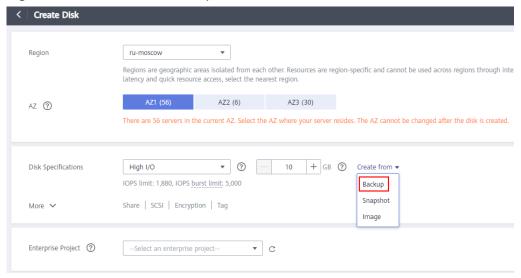
- If the disk data is no longer needed, delete the disk, and then create a new disk in the AZ where your target server is located.
- If the disk data is still required, create a new disk with the same data in the target AZ. The procedure is as follows:
 - a. Create a backup for the disk.

Figure 21-7 Create Backup



 Create a new disk from this backup. During the creation, select the target AZ. You can also change the settings of **Disk Type** and **Advanced Settings** if needed.

Figure 21-8 Create from backup



c. After the disk is created, click **Attach**. Your target server is displayed on the **Attach Disk** page.

Maximum Number of Disks That Can Be Attached to the Server Has Been Reached

Symptom: The Attach button is grayed out.

Solution:

- Non-shared disk: When you hover the mouse over the **Attach** button, message "This operation can be performed only when the disk is in the Available state" is displayed.
 - A non-shared disk can only be attached to one server. If the disk status is **Inuse**, the disk has been attached. You can detach the disk, wait until the disk status changes to **Available**, and then attach the disk to the target server.
- Shared disk: When you hover the mouse over the Attach button, message
 "This operation cannot be performed because the maximum number of servers that a shared disk can be attached to has been reached" is displayed.
 - A shared disk can be attached to a maximum of 16 servers, but you can detach the shared disk from one server and attach it to a new one if needed.

Check Whether the Servers Attached with the Shared Disk Are Running the Same Type of OS

Symptom: After you click **Attach**, the target server cannot be selected on the **Attach Disk** page, and message "A shared disk must be attached to servers with the same OS type" is displayed.

Solution: This message indicates that the OS type of the target server is inconsistent with that of the servers attached with the shared disk. You can change the OS type based your service requirements.

21.4.3 Can I Attach a Disk to Multiple Servers?

- A non-shared disk can only be attached to one server.
- A shared disk can be attached to up to 16 servers.

□ NOTE

- Shared disks are a type of EVS disks that can be attached to multiple servers.
- To use shared disks, you must set up a shared file system or cluster management system. If you directly attach a disk to multiple servers, the disk sharing attribute does not work and data may be overwritten.

21.4.4 Can I Attach a Disk to a Server in Another AZ?

No.

Disks and the servers you attach the disks to must be in the same AZ. The same is true for shared disks. Shared disks can only be attached to the servers in the same AZ.

21.4.5 How Can I Add a Data Disk to an Existing Server?

Data disks can be created during or after the server creation. If you create data disks during the server creation, the system will automatically attach the data disks to the server. If you create data disks after the server creation, you need to manually attach the data disks.

• On a Windows server:

- If a data disk is created along with the server, you need to log in to the server and initialize the disk. The data disk will be visible after the initialization succeeds.
- If no data disk is created along with the server, you need to create a data disk and attach it to the server. Then, you need to log in to the server and initialize the disk. The data disk will be visible after the initialization succeeds.

On a Linux server:

- If a data disk is created along with the server, you need to log in to the server and initialize the disk. The data disk will be visible after the initialization succeeds and the disk has been mounted via the mount command.
- If no data disk is created along with the server, you need to create a data disk and attach it to the server. Then, you need to log in to the server and initialize the disk. The data disk will be visible after the initialization succeeds and the disk has been mounted via the **mount** command.

For details, see Introduction to Data Disk Initialization Scenarios and Partition Styles.

21.4.6 Can I Attach Different Types of Disks to the Same Server?

Yes. Different types of EVS disks can be attached to the same server. You only need to make sure that these disks and the server are in the same AZ.

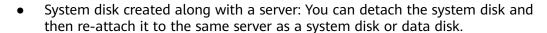
21.4.7 What Should I Do If a Linux EVS Disk Is Attached to a Windows Server?

You are not advised to attach a Linux EVS disk to a Windows server or attach a Windows EVS disk to a Linux server.

The disk information may fail to be displayed due to the inconsistent file systems. To solve this issue, you need to initialize and partition the disk again. Formatting the disk will destroy any data the disk has contained, so you should back up the data first.

21.4.8 Can I Change the Function of a System Disk or Data Disk Created Along with a Server?

This section describes how to change the function of a disk on the original server.



◯ NOTE

Each server can have only one system disk. If there is already a system disk attached to the server, then the detached system disk cannot be attached as a second system disk.

 Data disk created along with a server: You can detach the data disk and then re-attach it to the same server only as a data disk. It cannot be attached as a system disk.

21.4.9 How Do I Obtain My Disk Name in the ECS OS Using the Device Identifier Provided on the Console?

Scenarios

You find that the disk name displayed on the ECS OS is different from that displayed on the management console and you cannot determine which disk name is correct. This section describes how to obtain the disk name used on an ECS OS according to the device identifier on the console.

Obtaining the Disk ID of an ECS on the Console

- 1. Log in to the management console.
- 2. Under **Computing**, click **Elastic Cloud Server**.
- Click the target ECS name in the ECS list.
 The page providing details about the ECS is displayed.
- 4. Click the **Disks** tab and then to expand the disk information.
- 5. Check the device type and ID of the disk.

□ NOTE

If **Device Identifier** is not displayed on the web page, stop the ECS and restart it.

- KVM
 - If Device Type is VBD, use a serial number or BDF to obtain the disk device name.

If you use a serial number (recommended) to obtain the disk name, see Using a Serial Number to Obtain the Disk Name (Windows) and Using a Serial Number to Obtain a Disk Device Name (Linux).

If you use a BDF to obtain the disk device name, see **Using a VBD to Obtain a Disk Device Name (Linux)**. (BDF cannot be used to obtain the disk name of Windows ECSs.)

If Device Type is SCSI, use a WWN to obtain the disk name. For details, see Using a WWN to Obtain the Disk Name (Windows) and Using a WWN to Obtain a Disk Device Name (Linux).

Using a Serial Number to Obtain the Disk Name (Windows)

If a serial number is displayed on the console, use either of the following methods to obtain the disk name.

cmd

1. Start **cmd** in a Windows OS as an administrator and run either of the following commands:

wmic diskdrive get serialnumber wmic path win32_physicalmedia get SerialNumber wmic path Win32_DiskDrive get SerialNumber

MOTE

A serial number is the first 20 digits of a disk UUID.

For example, if the serial number of a VBD disk on the console is 97c876c0-54b3-460a-b, run either of the following commands to obtain the serial number of the disk on the ECS OS:

wmic diskdrive get serialnumber
wmic path win32_physicalmedia get SerialNumber
wmic path Win32_DiskDrive get SerialNumber
Information similar to the following is displayed:

Figure 21-9 Obtaining the disk serial number

```
C: Users Administrator>wmic diskdrive get serialnumber
SerialNumber
97c876c0-54b3-460a-b
C: Users Administrator>wmic path win32_physicalmedia get SerialNumber
SerialNumber
97c876c0-54b3-460a-b
C: Users Administrator>wmic path Win32_DiskDrive get SerialNumber
SerialNumber
97c876c0-54b3-460a-b
```

2. Run the following command to check the disk corresponding to the serial number:

wmic diskdrive get Name, SerialNumber

Figure 21-10 Checking the disk corresponding to the serial number

```
C:\Users\Administrator>wmic diskdrive get Name, SerialNumber
Name SerialNumber
\\.\PHYSICALDRIVEØ 97c876c0-54b3-460a-b
```

PowerShell

- Start PowerShell as an administrator in a Windows OS.
- 2. Run the following command to check the disk on which the logical disk is created:

- Windows Server 2012 or later
 - Run the following command to check the disk on which the logical disk is created:

Get-CimInstance -ClassName Win32_LogicalDiskToPartition | select Antecedent, Dependent |fl

As shown in Figure 21-11, the disk is Disk 0.

ii. Run the following command to view the mapping between the serial number and the disk:

Get-Disk |select Number, SerialNumber

As shown in Figure 21-11, the disk is Disk 0.

Figure 21-11 Viewing the disk on which the logical disk is created

```
PS C:\Users\Administrator> Get-CimInstance -ClassName Win32_LogicalDiskToPartition |select Anteceder

Antecedent : Win32_DiskPartition (DeviceID = "Disk #0, Partition #1")

PS C:\Users\Administrator> Get-Disk |select Number, SerialNumber

Number SerialNumber

0 97c876c0-54b3-460a-b
1 dswfal6520d39517815206127
```

- Versions earlier than Windows 2012
 - i. Run the following command to check the disk on which the logical disk is created:

Get-WmiObject -Class Win32_PhysicalMedia |select Tag, Serialnumber

ii. Run the following command to view the mapping between the serial number and the disk:

Get-WmiObject -Class Win32_LogicalDiskToPartition |select Antecedent, Dependent |fl

Using a Serial Number to Obtain a Disk Device Name (Linux)

If a serial number is displayed on the console, run either of the following commands to obtain the device name.

udevadm info --query=all --name=/dev/xxx | grep ID SERIAL

ll /dev/disk/by-id/*

Ⅲ NOTE

A serial number is the first 20 digits of a disk UUID.

For example, if the serial number of the VBD disk is 62f0d06b-808d-480d-8, run either of the following commands:

udevadm info --query=all --name=/dev/vdb | grep ID_SERIAL

ll /dev/disk/by-id/*

The following information is displayed:

[root@ecs-ab63 ~]# udevadm info --query=all --name=/dev/vdb | grep ID_SERIAL E: ID SERIAL=62f0d06b-808d-480d-8

[root@ecs-ab63 ~]# ll /dev/disk/by-id/* lrwxrwxrwx 1 root root 9 Dec 30 15:56 /dev/disk/by-id/virtio-128d5bfd-f215-487f-9 -> ../../vda lrwxrwxrwx 1 root root 10 Dec 30 15:56 /dev/disk/by-id/virtio-128d5bfd-f215-487f-9-part1 -> ../../vda1 lrwxrwxrwx 1 root root 9 Dec 30 15:56 /dev/disk/by-id/virtio-62f0d06b-808d-480d-8 -> ../../vdb

/dev/vdb is the disk device name.

Using a VBD to Obtain a Disk Device Name (Linux)

1. Run the following command to use a BDF to obtain the device name:

ll /sys/bus/pci/devices/BDF disk /D/virtio*/block

For example, if the BDF disk ID of the VBD disk is 0000:02:02.0, run the following command to obtain the device name:

ll /sys/bus/pci/devices/0000:02:02.0/virtio*/block

The following information is displayed:

[root@ecs-ab63 ~]# ll /sys/bus/pci/devices/0000:02:02.0/virtio*/block total 0 drwxr-xr-x 8 root root 0 Dec 30 15:56 **vdb**

/dev/vdb is the disk device name.

Using a WWN to Obtain the Disk Name (Windows)

- Obtain the device identifier on the console by referring to Obtaining the Disk ID of an ECS on the Console.
- 2. Manually convert the WWN.

For example, the obtained WWN (device identifier) is 688860300003252ffa16520d39517815.

- a. Obtain the 21st to 17th digits that are counted backwards (3252f).
- b. Convert a hexadecimal (3252f) to a decimal (206127).
- Start PowerShell as an administrator in a Windows OS.
- 4. Run the following command:

Get-CimInstance Win32_DiskDrive | Select-Object DeviceID, SerialNumber

5. In the command output, the disk whose serial number ends with **206127** is the disk corresponding to the WWN.

Figure 21-12 Disk with the serial number ending with 206127

```
PS C:\Users\Administrator> Get-CimInstance Win32_DiskDrive | Select-Object DeviceID, SerialNumber

DeviceID SerialNumber

\\.\PHYSICALDRIVEO 97c876c0-54b3-460a-b
dswfai6520d39517815206127
```

Using a WWN to Obtain a Disk Device Name (Linux)

- 1. Log in to the ECS as user **root**.
- 2. Run the following command to view the disk device name:

ll /dev/disk/by-id |grep WW/Mgrep scsi-3

For example, if the WWN obtained on the console is 6888603000008b32fa16688d09368506, run the following command:

ll /dev/disk/by-id |grep 6888603000008b32fa16688d09368506|grep scsi-3

The following information is displayed:

[root@host-192-168-133-148 block]# ll /dev/disk/by-id/ |grep 6888603000008b32fa16688d09368506 | grep scsi-3

lrwxrwxrwx 1 root root 9 May 21 20:22 scsi-3688860300008b32fa16688d09368506 -> ../../sda

21.5 Detachment

21.5.1 If I Detach a Disk, Will I Lose the Data on My Disk?

Data on a disk will not be lost after the disk is detached, and the disk can be reattached later if needed.

To ensure your data safety, you are advised to follow the instructions described in **Disk Detachment Process**.

Disk Detachment Process

- For disks not supporting online detachment:
 - a. Stop the server where the disk was attached.
 - After the server has stopped, detach the disk.
- For disks supporting online detachment:

Detach the disk from a running ECS. For details, see **Management** > **Detaching an EVS Disk from a Running ECS** in the *Elastic Cloud Server User Guide*.

21.5.2 Why Can't I Detach My Disk?

EVS disks can be used as system disks or data disks, but the way you detach each one is different.

• System disks: A system disk can only be detached offline. You must first stop the server that uses this system disk and then detach the disk.

In Linux, a system disk is typically mounted on /dev/vda. In Windows, a system disk is normally Volume (C:).

 Data disks: A data disk can be detached regardless of whether it is offline or online.

™ NOTE

In Linux, a data disk is typically mounted on a mount point other than /dev/vda. In Windows, a data disk is normally a volume other than Volume (C:).

- Offline detachment: The server must be in the **Stopped** state. If it is not, stop the server and then detach the data disk.
- Online detachment: Some OSs support online detachment. In this case, you do not need to stop the server before detaching the data disk. For more information, see Storage > Detaching an EVS Disk from a Running ECS in the Elastic Cloud Server User Guide.

21.6 Deletion

21.6.1 How Can I Recover Data from a Disk That Was Accidentally Deleted?

Check whether the disk has any snapshots or backups created.

 If there are, use the snapshot or backup to restore the disk data to the state when the snapshot or backup was created. For details, see Rolling Back Data from a Snapshot or .

NOTICE

If the disk was deleted after the last snapshot or backup was created, that incremental data cannot be restored.

• If there are not, the disk data cannot be restored.

21.7 Capacity

21.7.1 What Is the Maximum Capacity Supported for the System and Data Disks?

The maximum capacity supported for a system disk is 1024 GiB.

The maximum capacity supported for a data disk is 32768 GiB.

21.7.2 What Should I Do If My Disk Starts to Run Out of Space?

If your disk space starts to fill up, you can:

- Create a new disk and attach it to the server. For details, see Create an EVS
 Disk
- Expand the capacity of the existing disk. Both system disks and data disks can be expanded. For details, see **Disk Capacity Expansion**

Differences Between Expanding an EVS Disk and Creating an EVS Disk

The differences are as follows:

- Expanding an EVS disk is when you expand the capacity of an existing EVS disk. Some systems let you expand the capacity of EVS disks in use. In this case, services are not interrupted.
- If you create a new EVS disk and attach it to a server that already has an existing EVS disk, the new EVS disk and the original EVS disk are attached to the same server but independent from each other.

21.7.3 What Can I Do If the Capacity of My Disk Reaches the Maximum But I Still Need More Space?

The capacity of a single disk can be expanded to up to 32 TiB. If this still fails to meet your needs, it is recommended that you create RAID arrays with EVS disks or manage EVS disks using LVM.

21.7.4 What Should I Do If I Use fdisk to Initialize a Disk Larger Than 2 TiB and Then the Space in Excess of 2 TiB Cannot Be Displayed?

If your disk capacity is greater than 2 TiB, do not use fdisk to partition the disk. Or any space in excess of 2 TiB will be unable to show up after the disk is partitioned.

In this case, use parted to repartition the disk and choose the GPT partition style because MBR does not support disks over 2 TiB.

For details, see Introduction to Data Disk Initialization Scenarios and Partition Styles.

21.7.5 How Can I View My Disk Usage?

You can view your disk usages in either of the following ways:

View disk usages manually.

The details depend on the OS. This FAQ uses Windows Server 2008, Windows Server 2016, and Linux as samples to describe how to view the disk usage.

- Viewing Disk Usage in Linux
- Viewing Disk Usage in Windows Server 2008
- Viewing Disk Usage in Windows Server 2016
- Installing Agent to View Disk Usage

Viewing Disk Usage in Linux

In this section, CentOS 7.4 64bit is used as an example. The details depend on if you need to view the available space or not.

• To guery the total capacity only, run **lsblk**.

Information similar to the following is displayed: [root@ecs-test-0001 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda 253:0 0 40G 0 disk

__vda1 253:1 0 40G 0 part /
vdb 253:16 0 40G 0 disk
__vdb1 253:17 0 40G 0 part

In the command output, the server has two disks, /dev/vda and /dev/vdb. System disk /dev/vda has 40 GiB of capacity, as does data disk /dev/vdb.

• To query the total capacity and display the space available as well, run **df -TH**. Ensure that the disk has been attached and initialized before running this command.

Information similar to the following is displayed:

```
[root@ecs-0001 ~]# df -TH
           Type Size Used Avail Use% Mounted on
Filesystem
                   43G 2.0G 39G 5% /
/dev/vda1
            ext4
           devtmpfs 509M 0 509M 0% /dev
devtmpfs
          tmpfs 520M 0 520M 0% /dev/shm
tmpfs 520M 7.2M 513M 2% /run
tmpfs
tmpfs
                  520M 0 520M 0% /sys/fs/cgroup
tmpfs
           tmpfs
tmpfs
           tmpfs
                  104M
                          0 104M 0% /run/user/0
                  43G 51M 40G 1% /mnt/sdc
/dev/vdb1
           ext4
```

In the command output, the server has two partitions, /dev/vda1 and /dev/vdb1. Partition /dev/vda1 is used to deploy the OS, and its total capacity, used capacity, and available capacity are 43 GiB, 2 GiB, and 39 GiB, respectively. Partition /dev/vdb1's total capacity, used capacity, and available capacity are 43 GiB, 51 MiB, and 40 GiB, respectively.

Viewing Disk Usage in Windows Server 2008

In this section, Windows Server 2008 R2 Enterprise 64bit is used as an example.

Step 1 On the desktop of the server, right-click **Computer** and choose **Manage** from the shortcut menu.

The **Server Manager** window is displayed.

Step 2 In the navigation tree on the left, choose **Storage** > **Disk Management**.

The sizes and available spaces of the volumes on the current disk are displayed in the middle pane.

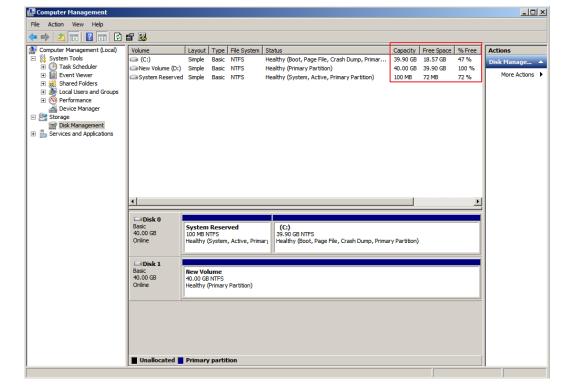


Figure 21-13 Disk Management page

----End

Viewing Disk Usage in Windows Server 2016

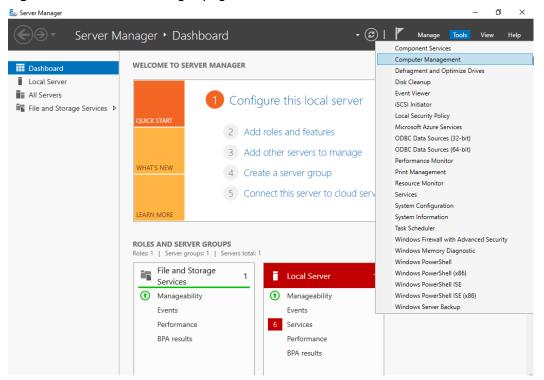
In this section, Windows Server 2016 Standard 64bit is used as an example.

Step 1 On the desktop of the server, click the start icon in the lower left corner. The **Windows Server** window is displayed.

Step 2 Click Server Manager.

The **Server Manager** window is displayed.

Figure 21-14 Server Manager page



- **Step 3** In the upper right corner, choose **Tools** > **Computer Management**.
- **Step 4** Choose **Storage** > **Disk Management**.

In the middle pane, you can view the sizes and available spaces of the volumes on the disk.

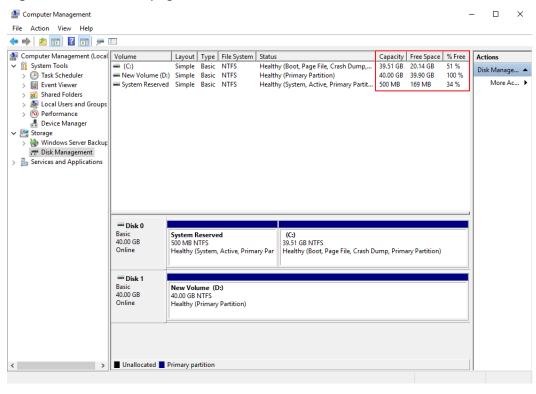


Figure 21-15 Disk list page

----End

Installing Agent to View Disk Usage

Some disk monitoring metrics require that the agent to be installed.

- For instructions about how to install the agent on a Windows ECS, see "Installing and Configuring the Agent (Windows)" in the *Cloud Eye User Guide*
- For instructions about how to install the agent on a Linux ECS, see "Installing and Configuring the Agent (Linux)" in the *Cloud Eye User Guide*.

Table 21-7 Disk metrics

Metric	Paramet er	Description	Value Range	Monit ored Objec t & Dime nsion	Monitori ng Period (Raw Data)
disk_fre e	(Agent) Available Disk Space	Free space on the disks Unit: GB Linux: Run the df -h command to check the value in the Avail column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). Windows: Use the WMI interface to call GetDiskFreeSpaceExW API to obtain disk space data. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	≥ 0	ECS - Moun t point	1 minute

Metric	Paramet er	Description	Value Range	Monit ored Objec t & Dime nsion	Monitori ng Period (Raw Data)
disk_tot al	(Agent) Disk Storage Capacity	Total space on the disks, including used and free Unit: GB Linux: Run the df -h command to check the value in the Size column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). Windows: Use the WMI interface to call GetDiskFreeSpaceExW API to obtain disk space data. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	≥ 0	ECS - Moun t point	1 minute

Metric	Paramet er	Description	Value Range	Monit ored Objec t & Dime nsion	Monitori ng Period (Raw Data)
disk_use d	(Agent) Used Disk Space	Used space on the disks Unit: GB Linux: Run the df -h command to check the value in the Used column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). Windows: Use the WMI interface to call GetDiskFreeSpaceExW API to obtain disk space data. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	≥ 0	ECS - Moun t point	1 minute

Metric	Paramet er	Description	Value Range	Monit ored Objec t & Dime nsion	Monitori ng Period (Raw Data)
disk_use dPercen t	(Agent) Disk Usage	Percentage of total disk space that is used, which is calculated as follows: Disk Usage = Used Disk Space/Disk Storage Capacity Unit: percent Linux: It is calculated as follows: Used/Size. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). Windows: Use the WMI interface to call GetDiskFreeSpaceExW API to obtain disk space data. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	0-100	ECS - Moun t point	1 minute

21.7.6 How Can I Monitor My Disk Usages?

Some disk monitoring metrics require that the agent to be installed.

- For instructions about how to install the agent on a Windows ECS, see "Installing and Configuring the Agent (Windows)" in the *Cloud Eye User Guide*.
- For instructions about how to install the agent on a Linux ECS, see "Installing and Configuring the Agent (Linux)" in the *Cloud Eye User Guide*.

Table 21-8 Disk metrics

Metric	Paramet er	Description	Value Range	Monit ored Objec t & Dime nsion	Monitori ng Period (Raw Data)
disk_fre e	(Agent) Available Disk Space	Free space on the disks Unit: GB Linux: Run the df -h command to check the value in the Avail column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). Windows: Use the WMI interface to call GetDiskFreeSpaceExW API to obtain disk space data. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	≥ 0	ECS - Moun t point	1 minute

Metric	Paramet er	Description	Value Range	Monit ored Objec t & Dime nsion	Monitori ng Period (Raw Data)
disk_tot al	(Agent) Disk Storage Capacity	Total space on the disks, including used and free Unit: GB Linux: Run the df -h command to check the value in the Size column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). Windows: Use the WMI interface to call GetDiskFreeSpaceExW API to obtain disk space data. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	≥ 0	ECS - Moun t point	1 minute

Metric	Paramet er	Description	Value Range	Monit ored Objec t & Dime nsion	Monitori ng Period (Raw Data)
disk_use d	(Agent) Used Disk Space	Used space on the disks Unit: GB Linux: Run the df -h command to check the value in the Used column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). Windows: Use the WMI interface to call GetDiskFreeSpaceExW API to obtain disk space data. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	≥ 0	ECS - Moun t point	1 minute

Metric	Paramet er	Description	Value Range	Monit ored Objec t & Dime nsion	Monitori ng Period (Raw Data)
disk_use dPercen t	(Agent) Disk Usage	Percentage of total disk space that is used, which is calculated as follows: Disk Usage = Used Disk Space/Disk Storage Capacity	0-100	ECS - Moun t point	1 minute
		 Unit: percent Linux: It is calculated as follows: Used/Size. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). Windows: Use the WMI interface to call GetDiskFreeSpaceExW API to obtain disk space data. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). 			

21.7.7 Can I Transfer the Data Disk Capacity to a System Disk?

Sorry, you cannot.

Currently, the capacity of an EVS disk cannot be transferred to another disk. Multiple EVS disks cannot be combined into a single, larger disk, either.

Common Scenarios

- 1. You want to expand the system disk capacity, but created a new data disk.
- 2. You want to expand the system disk capacity, but expanded the data disk capacity.

Recommended Solution

- If you do not need the data on the data disk, you can just delete the data disk, and then **expand the system disk**.
- If you need the data on the data disk, you can create a small-capacity data disk, copy data from the original data disk to the new data disk, and then expand the system disk capacity.
 - a. Back up the data disk using the CBR service or by creating a snapshot. For details about backups, see section "Managing EVS Backups" in the *Elastic Volume Service User Guide*. For details about snapshots, see **Creating a Snapshot**.
 - b. Create a new data disk with the desired capacity and attach it to the server. After initializing the disk, copy the data from the original data disk to the new data disk.
 - c. Confirm that the services on the new data disk are available. Then, delete the original data disk, and delete the created backup.
 - d. Expand the system disk capacity by referring to Disk Capacity Expansion

21.8 Performance

21.8.1 How Can I Test My Disk Performance?

Precautions

In the disk performance test, if the start sector number is not 4-KiB aligned, the disk performance will be greatly affected. Ensure that the start sector number is 4-KiB aligned before you start the test.

□ NOTE

To test the performance of a shared disk, the following requirements must be met:

- The shared disk must be attached to multiple servers (ECSs or BMSs).
- If the shared disk is attached to multiple ECSs, these ECSs must belong to the same anti-affinity ECS group.

If these ECSs fail to meet the anti-affinity requirement, the shared disk cannot reach the optimal performance.

The testing process for Windows and Linux is different.

- Windows
- Linux

Windows

The way you test disk performance depends on the server OS. This section uses Windows 7 Professional 64-bit as an example. For other Windows OSs, see the corresponding OS documentations.

Install the performance measurement tool Iometer before the test. You can obtain the tool at http://www.iometer.org/.

- **Step 1** Log in to the server.
- **Step 2** Press win+R to open the Run window. Enter msinfo32 and click OK. The system information window is displayed.
- **Step 3** Choose **components** > **storage** > **disks**. In the right pane, view the partition offset.
 - If 4096 can be divided by the parameter value, the partition is 4-KiB aligned. Go to **Step 4**.
 - If 4096 cannot be divided by the parameter value, the partition is not 4-KiB aligned. Ensure 4-KiB alignment for the partition before continuing the test.

NOTICE

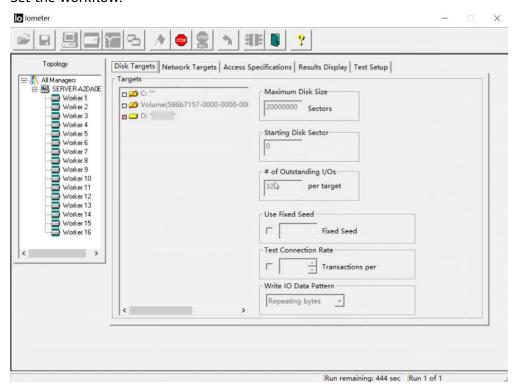
If you delete the partition and select another start sector number for 4-KiB alignment, you will lose all the data on that partition. Exercise caution when performing this operation.

Step 4 Use lometer to test the disk performance. For details, see the lometer product documentation.

When the disk IOPS and throughput are tested, the parameter settings for lometer and fio are the same. For details, see **Table 21-9**.

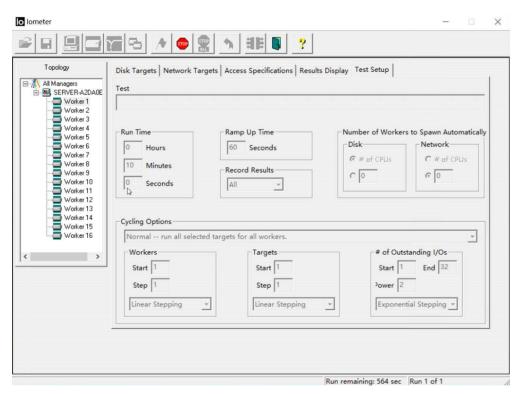
The following example uses Iometer to test the disk performance.

1. Set the workflow.

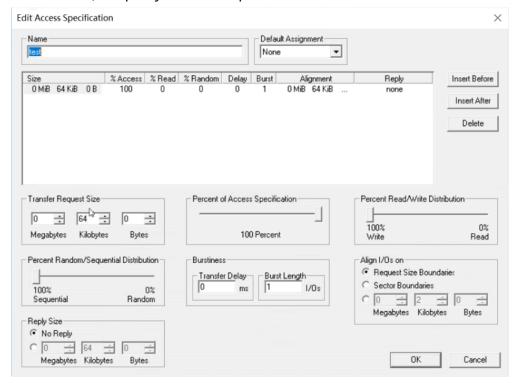


2. Set the test run time.

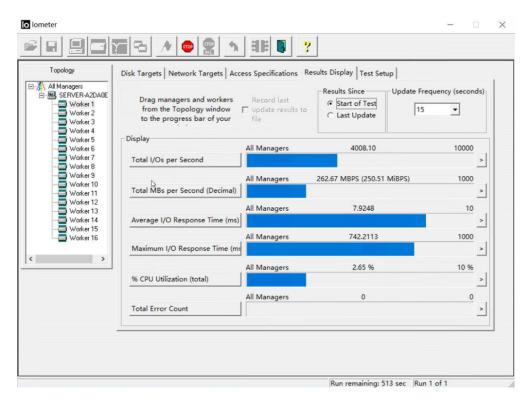
In this example, the test run time is set to 10 minutes, with 60 seconds ramp up time. Disk performance is tested after the writes are stable.



3. Set the data block size and read/write policy. In this example, the disk size is set to 64 KiB, the policy is 100% sequential write.



4. View the test results.



----End

Linux

If you use an old version Linux OS, for example CentOS 6.5, and run **fdisk** to create partitions, the default start sector number will not be 4-KiB aligned, which will greatly affect the test performance. For this reason, if such an OS is used, you are advised to select a new start sector number, one that is 4-KiB aligned, when creating partitions.

The way you test disk performance depends on the server OS. This section uses CentOS 7.2 64-bit as an example. For other Linux OSs, see the corresponding OS documentations.

Install the performance measurement tool, fio, before the test.

- **Step 1** Log in to the server and switch to user **root**.
- **Step 2** Before you start the test, run the following command to check whether the start sector number is 4-KiB aligned:

fdisk -lu

Information similar to the following is displayed: [root@ecs-centos sdc]# fdisk -lu

Disk /dev/xvda: 10.7 GiB, 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 GiB, 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 GiB, 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 sector number, the number is 4-KiB aligned. Go to **Step 3**.

• If 8 cannot be divided by the start sector number, the number is not 4-KiB aligned. Delete the partition and select a 4-KiB aligned start sector number for the new partition before continuing the test.

NOTICE

If you delete the partition and select another start sector number for 4-KiB alignment, you will lose all the data on that partition. Exercise caution when performing this operation.

Step 3 Run the following commands and use fio to test the disk performance:

- To test random write IOPS, run the following command: fio -direct=1 iodepth=128 -rw=randwrite -ioengine=libaio -bs=4k -size=10G -numjobs=1 runtime=600 -group_reporting -filename=/opt/fiotest/fiotest.txt name=Rand_Write_IOPS_Test
- To test random read IOPS, run the following command: **fio** -**direct**= 1 **iodepth**= 128 -**rw**= randread -**ioengine**= libaio -**bs**= 4k -**size**= 10G -**numjobs**= 1 **runtime**= 600 -**group**_reporting -**filename**= /opt/fiotest/fiotest.txt **name**= Rand_Read_IOPS_Test
- To test write throughput, run the following command: fio -direct=1 iodepth=32 -rw=write -ioengine=libaio -bs=1024k -size=10G -numjobs=1 runtime=600 -group_reporting -filename=/opt/fiotest/fiotest.txt name=Write BandWidth Test
- To test read throughput, run the following command: fio -direct=1 iodepth=32 -rw=read -ioengine=libaio -bs=1024k -size=10G -numjobs=1 runtime=600 -group_reporting -filename=/opt/fiotest/fiotest.txt name=Read BandWidth Test
- To test random read latency, run the following command: fio -direct=1 iodepth=1-rw=randread -ioengine=libaio -bs=4k -size=10G -numjobs=1 runtime=60 -group_reporting -filename=/opt/fiotest/fiotest.txt name=Rand_Read_LATE_Test

NOTICE

- When using fio to perform a raw disk performance test, ensure that no partitions and file systems have been created on the disk and there is no data stored on the disk. Or, the raw disk test will damage the file system, and data on the disk will become read-only. In this case, your only option will be to delete the disk and buy a new one to continue the test.
- Do not perform the test on a disk with service data on it. If such test is a must, you are advised to perform the test as follows:
 - Back up the disk data before the test as you may damage the data on the disk.
 - Specify a file, for example -filename=/opt/fiotest/fiotest.txt, to test the performance of the file system.

Table 21-9 lists the fio test parameters.

Table 21-9 Parameter description

Parameter	Description
direct	Defines whether direct I/O is used. - Set to 0 : buffered I/O is used. - Set to 1 : direct I/O is used.
iodepth	Defines the I/O queue depth. This queue depth refers to the queue depth of each thread regardless of whether a single or multiple threads are used in the test. Total concurrent I/Os of fio = iodepth x numjobs Examples: - If there is a single thread and -iodepth=32, the I/O queue depth of this thread is 32 and the total concurrent I/Os of fio is 32 (32 x 1). - If there are three threads and -iodepth=32, the I/O queue depth of each thread is 32 and the total concurrent I/Os of fio is 96 (32 x 3).
rw	Defines the test read/write policy. - randread: random read - randwrite: random write - read: sequential read - write: sequential write - randrw: mixed random read/write

Parameter	Description
ioengine	Defines how fio delivers the I/O request (synchronously or asynchronously).
	 Synchronous I/O: Only one I/O request is delivered at a time, and the response is returned after the kernel has processed the request. That said, the single-thread I/O queue depth is always less than 1, and multi-thread concurrent processing can be used to handle such issues. Normally, 16 to 32 concurrent working threads fully occupy the I/O queue depth.
	 Asynchronous I/O: Multiple I/O requests are delivered using libaio at a time. Wait for the process to complete and reduce the interaction times to improve efficiency.
bs	Defines the I/O block size. The unit can be KiB, Kb, MiB, and Mb, and the default value is 4 KiB.
size	Defines the amount of data processed by the test I/Os. If parameters, such as runtime , are not specified, the test ends when 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-GiB read/write data amount. A percentage indicates the ratio of read/write data amount to the total size of 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_rep orting	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.
	 If a file is specified, the performance of the file system is tested. Example: -filename=/opt/fiotest/fiotest.txt
	 If a device name is specified, the performance of the raw disk is tested. Example: -filename=/dev/vdb
	NOTICE If the test is performed on a disk already has partitions and file systems created as well as data on it, user parameter filename to specify a file so that the original file system is not damaged and the data is not overwritten.
name	Defines the test task name.

----End

21.8.2 Why Does My Disk Performance Test Using Fio Have Incorrect Results?

Symptom

You have followed the test performance method, but the test results do not meet expectations.

Troubleshooting

During a disk performance test, the disk and stress test conditions play an important role.

Possible causes are listed here in order of their probability.

If the fault persists after you have ruled out one cause, move on to the next one in the list.

NOTICE

Some operations may result in data loss. It is recommended that you use raw disks for performance test.

Figure 21-16 Troubleshooting

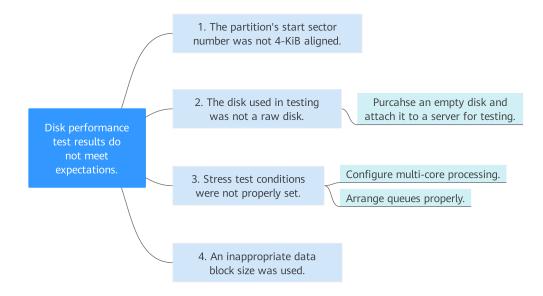


Table 21-10 Troubleshooting

Possible Cause	Solution	
The partition's start sector number is not 4-	Go Check Whether Partition's Start Sector Number Is 4-KiB Aligned.	
KiB aligned.	Delete the partition and select a 4-KiB aligned start sector number for the new partition.	
The disk used in testing was not a raw disk.	Purchase an empty disk and attach it to a server for testing.	
Stress test conditions were not properly set.	Configure multi-core processing and arrange queues properly to maximize the concurrent performance.	
An inappropriate data block size was used.	 Set a suitable data block size. When testing the disk IOPS, set the data block size to a small value, for example, 4 KiB. 	
	When testing the disk throughput, set the data block size to a large value, for example, 1024 KiB.	

Check Whether Partition's Start Sector Number Is 4-KiB Aligned

- **Step 1** Log in to the server and switch to user **root**.
- **Step 2** Before you start the test, run the following command to check whether the start sector number is 4-KiB aligned:

fdisk -lu

Information similar to the following is displayed:

[root@ecs-centos sdc]# fdisk -lu

Disk /dev/xvda: 10.7 GiB, 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 GiB, 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 GiB, 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 sector number, the number is 4-KiB aligned.
- If 8 cannot be divided by the start sector number, the number is not 4-KiB aligned. Delete the partition and select a 4-KiB aligned start sector number for the new partition before continuing the test.

NOTICE

If you delete the partition and select another start sector number for 4-KiB alignment, you will lose all the data on that partition. Exercise caution when performing this operation.

----End

21.8.3 How Can I Handle a Slowdown in Disk Read/Write Speed and Increased I/O?

Symptom

If you are aware of a service slowdown, depending on if you are examining a Windows or Linux server, you can take the following actions:

- Windows: Open **Task Manager** and view the average response time.
- Linux: Run **iostat -dx** to view the I/O performance.

If the disk read/write speed is slowed down, disk I/O increases, or the await time increases, the disk is likely encounters a performance bottleneck.

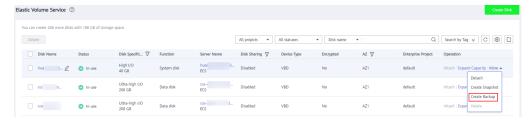
Solution

It is recommended that you choose a disk type with higher performance.

If your disk contains important data, create a new disk from the backup, so you do not lose any of the disk data. The procedure is as follows:

1. Create a backup for the disk.

Figure 21-17 Create Backup



2. Create a new disk from this backup. During the creation, select a new disk type and configure advanced settings (sharing, SCSI, and encryption) based on your service requirements.

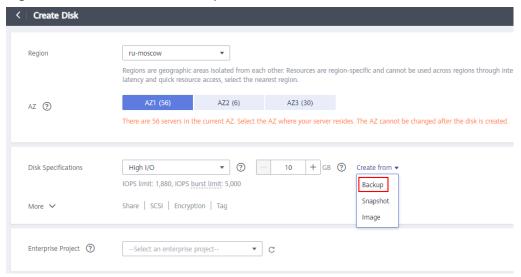


Figure 21-18 Create from backup

21.8.4 Why My Disk's Read IOPS Can't Reach the Theoretical Maximum IOPS When the Disk I/O Usage Is Almost 100%?

Symptom

A 500 GiB ultra-high I/O disk had an I/O usage of 99.94%, but it only had 12,000 IOPS.

Description

 100% disk I/O usage does not mean that the disk IOPS reaches the maximum.

Disk I/O usage calculates the read or write operations performed by a disk in a measurement period. It describes how busy a disk is, not the disk I/O performance.

EVS disks can process I/O requests concurrently, so 100% disk I/O usage does not mean that the disk encounters the performance bottleneck. For example, an EVS disk takes 0.1 second to process an I/O request and can process 10 I/O requests concurrently. If 10 I/O requests are submitted serially, the disk takes 1 second to process all I/O requests. In this 1-second measurement period, the disk I/O usage reaches 100%. However, if 10 I/O requests are submitted concurrently, the disk takes just 0.1 second to process all the requests. This way, the disk I/O usage in a 1-second measurement period is only 10%. This means that a disk can still process I/O requests even if its I/O usage reaches 100%.

• Why does the disk not reach the theoretical maximum IOPS?

The actual maximum IOPS that a disk can reach is calculated as follows: Disk IOPS = Min. (Max. IOPS, Min. IOPS + IOPS per GiB x Disk size). For a 500 GiB disk, its IOPS is calculated as follows: Disk IOPS = Min. (50,000, 1,800 + 50 \times 500) = 26,800

The disk read IOPS is the number of read operations performed by the disk per second. IOPS is also affected by latency. In a single-queue access scenario with 4 KiB data blocks, the access latency of an ultra-high I/O disk is 1 ms,

which means the disk can process 1,000 requests (IOPS) in a second. 12,000 IOPS indicates that the queue depth is 12. To reach the theoretical maximum IOPS (26,800), the queue depth should reach 26.

21.9 Sharing

21.9.1 Do I Have to Deploy a Cluster to Use Shared Disks?

Yes.

If you simply attach a shared disk to multiple servers, files cannot be shared among them. Because there are no mutually agreed data read/write rules among servers, read and write operations from them may interfere with each other, or unpredictable errors may occur.

Shared EVS disks do not have cluster management capabilities. You need to build a clustered system for data sharing, such as Windows MSCS, Veritas VCS, and Veritas CFS clusters.

21.9.2 How Many Servers Can I Attach a Shared Disk to?

A shared disk can be attached to up to 16 servers.

21.9.3 How Can I Attach a Shared Disk to Multiple Servers?

A shared disk can be attached to multiple servers on the management console. You can choose to attach it to servers one by one or in a batch.

For details, see **Attaching a Shared Disk**.

21.9.4 Can a Shared Disk Be Attached to Servers That Belong to Different Accounts?

No. A Shared disk can only be attached to servers that belong to the same account and are in the same AZ.

21.9.5 Can I Attach a Shared Disk to Servers Running Different OSs?

- It is recommended that you do not simultaneously attach a shared disk to servers running Linux and Windows.
- If a shared disk is attached to servers running different versions of the same OS type, it can be used normally. For example, a shared disk attached to one server running CentOS 6 and another server running CentOS 7 can work fine.

21.10 Snapshot

21.10.1 What Are the Typical Causes of a Snapshot Creation Failure?

A snapshot creation will fail if its source disk is in an intermediate state, such as **Attaching** and **Expanding**, or an abnormal state, such as **Error** and **Restoration failed**.

Ensure that a disk is in the **In-use** or **Available** state before creating a snapshot.

21.10.2 Does EVS Support Automatic Snapshot Creation?

No, but you can develop a program to periodically create snapshots through APIs. Alternatively, you may use the CBR service and enable auto backup to create backups periodically. For details, see section "Creating a Backup Policy" in the *Cloud Backup and Recovery User Guide*.

21.10.3 Can I Create Snapshots for Multiple Disks at a Time?

No.

Snapshots cannot be created for multiple disks at the same time. You must create snapshots for disks one by one.

21.10.4 How Is a Snapshot Created for My Disk?

There are two types of snapshots: manual snapshots you create on-demand and automatic snapshots created by the system.

- Manual snapshots: You may manually create snapshots to rapidly save the disk data at specific points of time. For details about how to create snapshots, see Creating a Snapshot.
- Automatic snapshots: During the creation of a cloud server backup, or a disk backup with the CBR service, the system automatically creates a snapshot and saves the latest snapshot for each disk. If the disk already has a backup and a new backup is created, the system will automatically delete the old snapshot and save the latest one generated. This snapshot is free of charge. You can view the snapshot details only but cannot perform any operations on it.

21.10.5 Why Can't I Roll Back My Disk Data from a Snapshot?

Possible causes are as follows:

- Snapshot data can be rolled back only when the status of the snapshot's source disk is Available or Rollback failed. If the snapshot's source disk is in the In-use state, detach the disk first and then roll back the snapshot data. After the rollback succeeds, re-attach the disk.
- A snapshot whose name starts with autobk_snapshot_vbs_, manualbk_snapshot_vbs_, autobk_snapshot_csbs_, or manualbk_snapshot_csbs_ is automatically generated during backup. Such a snapshot can only be viewed. It cannot be used to roll back the disk data.

21.10.6 Can I Roll Back the Data from a Snapshot After Reinstalling the OS or Formatting the Disk?

- If you have reinstalled or changed the server OS, snapshots of the system disk are automatically deleted. Snapshots of the data disks can be used as usual.
- If you have formatted a disk, snapshots can still be used to roll back the disk data.

21.10.7 How Is the Snapshot Size Calculated?

When you create a snapshot for a disk for the first time, the snapshot is a full snapshot, and its size is the same as the disk size. Subsequent snapshots created for the same disk are incremental snapshots, so they are only as large as the amount of data that has been added.

21.10.8 Do Snapshots Take Space on the Disk?

No.

Snapshots are stored on the physical disks that provide storage resources for EVS disks. Therefore, snapshots do not use the EVS disk space.

21.10.9 Can I Perform Multiple Rollback Operations for a Snapshot?

Yes. You can roll back a snapshot to the source disk multiple times as needed.

21.10.10 Can I Replicate Snapshots to Other Regions or Accounts?

No. Snapshots cannot be replicated to other regions or accounts.

21.10.11 Why Can't I Find My Snapshot?

Possible causes are:

- You have manually deleted the snapshot.
- You have deleted the disk based on which the snapshot was created. If a disk
 is deleted, all the snapshots created for this disk will be automatically deleted.
- You have reinstalled or changed the server OS. System disk snapshots will be automatically deleted after a server OS has been reinstalled or changed.

21.10.12 Can I Use a Snapshot to Create a Disk and Roll Back the Snapshot at Almost the Same Time?

There are two conditions:

- If you first use the snapshot to create a new disk, you can then roll back the snapshot to the original disk.
- If you first roll back the snapshot to the original disk, you cannot use the snapshot to create a new disk then.

21.11 Backup

21.11.1 Do I Need to Stop the Server Before Performing a Backup?

No. You can back up servers that are in use. When a server is running, data is written into disks on the server, and some newly generated data is cached in the server memory. During a backup task, data in the memory will not be automatically written into disks, so the disk data and their backups may be inconsistent.

To ensure data integrity, you are advised to perform the backup during off-peak hours when no data is written to the disks.

21.11.2 Can I Back Up and Restore My EVS Disk to a Different Region?

EVS disks that have been attached to servers can be backed up along with the servers using CBR. In addition, CBR cloud server backup allows you to replicate backups to a different region and create images in that region using replicated backups.

EVS disks that have not been attached do not support cross-region backup and restoration.

21.11.3 How Do I View My Backup Data?

You can check your backup data in the following ways:

◯ NOTE

Backup data cannot be viewed on the CBR console.

Server Backups

- 1. Create an image from a server backup. For details, see section "Using a Backup to Create an Image" in the *Cloud Backup and Recovery User Guide*.
- 2. Use the image to create a server. For details, see section "Creating an ECS from an Image" in the *Image Management Service User Guide*.
- 3. Log in to the server to view the data.

Disk Backups

- 1. Create a new disk from a disk backup. For details, see section "Using a Backup to Create a Disk" in the *Cloud Backup and Recovery User Guide*.
- 2. Attach the disk to a server. For details, see section "Attaching a Non-Shared Disk" or "Attaching a Shared Disk" in the *Elastic Volume Service Getting Started*.
- 3. Log in to the server to view the data.

SFS Turbo Backups

- 1. Create a new SFS Turbo file system from an SFS turbo backup. For details, see section "Using a Backup to Create a File System" in the *Cloud Backup and Recovery User Guide*.
- 2. Mount the file system to a server.
 - To mount the file system to a Linux server, see section "Mounting an NFS File System to ECSs (Linux)" in the *Scalable File Service Getting Started*.
 - To mount the file system to a Windows server, see section "Mounting an NFS File System to ECSs (Windows)" in the Scalable File Service Getting Started.
- 3. Log in to the server to view the data.



A.1 EVS Disk Status

An EVS disk has several statuses. **Table A-1** lists EVS disk statuses, the meaning of each status, and the operations a disk in each status allows.

Table A-1 Disk status details

EVS Disk Status	Status Description	Allowed Operation	
In-use	The EVS disk is attached to a server and is in use.	DetachingExpanding capacity	
Available	The EVS disk has not been attached to any server and can be attached.	 Attaching Expanding capacity Deleting Rolling back snapshots to EVS disks 	
Creating	The EVS disk is being created.	None	
Attaching	The EVS disk is being attached to a server.	None	
Detaching	The EVS disk is being detached from a server.	None	
Deleting	The EVS disk is being deleted.	None	
Expandin g	The capacity of the EVS disk is being expanded.	None	
Uploading	Data on the EVS disk is being uploaded to an image. This status occurs when you create an image from a server.	None	

EVS Disk Status	Status Description	Allowed Operation
Download ing	Data is being downloaded from an image to the EVS disk. This status occurs when you create a server.	None
Error	An error occurs when you try to create an EVS disk.	Deleting
Deletion failed	An error occurs when you try to delete the EVS disk.	None
Expansion failed	An error occurs when you try to expand the capacity of the EVS disk.	Deleting
Rolling back	Data on the EVS disk is being restored from a snapshot. NOTE • When you roll back a snapshot, you can only restore the data to the original disk. Data restoration to a specific disk is not possible. • A snapshot can only be rolled back when the original disk is in the available or error_rollbacking state.	None
Rollback failed	An error occurs when the EVS disk is rolled back from a snapshot.	DeletingRolling back snapshots to EVS disks

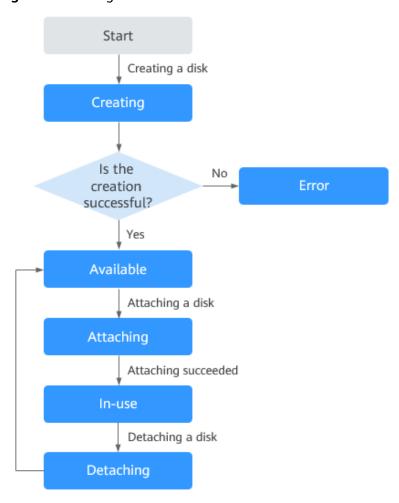


Figure A-1 Change of some EVS disk statuses

◯ NOTE

If an EVS disk status is Error, Deletion failed, Expansion failed, Restoration failed, or Rollback failed, you can rectify the error by referring to What Should I Do If an Error Occurs on My EVS Disk?.

A.2 EVS Snapshot Status

An EVS snapshot has several statuses. **Table A-2** lists the EVS snapshot statuses, the meaning of each status, and the operations a snapshot in each status allows.

Table A-2 Snapshot status details

Snapshot Status	Description	Allowed Operation
Creating	The snapshot is being created.	No operations are allowed.

Snapshot Status	Description	Allowed Operation
Available	The snapshot is successfully created.	Creating EVS disks using snapshots
		Deleting snapshots
		 Rolling back snapshots to source EVS disks
Deleting	The snapshot is being deleted.	No operations are allowed.
Error	An error occurs when you try to create a snapshot.	Deleting
Deletion failed	An error occurs when you try to delete a snapshot.	No operations are allowed.
Rolling back	The snapshot is rolling back data.	No operations are allowed.
	 When you roll back a snapshot to an EVS disk, you can only roll back the snapshot to the source EVS disk. Rollback to a specified disk is not supported. 	
	 You can roll back an EVS disk from a snapshot only when the disk is in the Available or Rollback failed state. 	

B Change History

Released On	Description
2023-06-12	This issue is the fourth official release, which incorporates the following change:
	Added screenshots in sections Create an EVS Disk, Attaching a Non-Shared Disk, Attaching a Shared Disk, Attaching an Existing System Disk, Attaching an Existing Non-Shared Disk, Viewing EVS Disk Details, Managing EVS Backups, Creating a Snapshot, How Can I View My Disk Details?, Can I Change the Disk Type, Device Type, or Sharing Attribute of My Disk?, Why Can't I Attach My Disk to a Server?, and How Can I Handle a Slowdown in Disk Read/Write Speed and Increased I/O?.

Released On	Description
2023-03-13	This issue is the third official release, which incorporates the following changes: Modified the following content:
	Document content has been supplemented in sections What Is EVS?, Disk Types and Performance, EVS Encryption, EVS Backup, Permissions, EVS and Other Services, Process Overview, Create an EVS Disk, Attaching a Non-Shared Disk, Attaching a Shared Disk, Initializing a Windows Data Disk (Windows Server 2008), Initializing a Windows Data Disk (Windows Server 2019), Initializing a Linux Data Disk (fdisk), Initializing a Linux Data Disk (parted), Initializing a Windows Data Disk Larger Than 2 TiB (Windows Server 2008), Initializing a Windows Data Disk Larger Than 2 TiB (Windows Server 2012), Initializing a Linux Data Disk Larger Than 2 TiB (parted), Expansion Overview, Expanding Capacity for an In-use EVS Disk, Expanding Capacity for an Available EVS Disk, Extending Disk Partitions and File Systems (Windows Server 2008), Partition and File System Extension Preparations (Linux), Extending Partitions and File Systems for System Disks (Linux), Detaching a Data Disk, Deleting EVS Disks, Managing EVS Backups, Snapshot Overview, Creating a Snapshot, Deleting a Snapshot, Rolling Back Data from a Snapshot, Creating an EVS Disk from a Snapshot, How Do I Start Using a Newly Created Disk?, Will My Disk Data Be Lost After I Expand the Disk Capacity?, Do I Need to Detach an EVS Disk Before Expanding Its Capacity?, Can I Attach a Disk to a Server in Another AZ?, How Can I Attach a Shared Disk to Multiple Servers?, and EVS Disk Status.
	Added the following content: Added sections EVS Three-Copy Redundancy, Constraints, EVS Concepts, Extending Disk Partitions and File Systems (Windows Server 2016), Attaching an Existing Disk, Viewing EVS Disk Details, Viewing Details of a Snapshot, Managing a Tag, Changing EVS Disk Name, Viewing EVS Monitoring Data (Agent Installed and Simplified Monitoring Metrics Used), Managing Quotas, Summary, Can I Change the Disk Type, Device Type, or Sharing Attribute of My Disk?, What Should I Do If an Error Occurs on My EVS Disk?, How Can I Obtain ECS NIC Information?, Why Do Some of My EVS Disks Not Have WWN Information?, How Can I Migrate Data from an EVS Disk?, What Does the "reserveVolume" Trace Mean in CTS?, How Can I Upload Files to My EVS Disk?, Can I Use Backups or Snapshots Created Before Capacity Expansion to Restore Data on Expanded Disks?, How Can I Allocate Newly Added Space to a New Partition?, How Can I Allocate Newly Added Space to an Existing Partition?, What Can't I Expand Capacity for My Disk?, How Do I Extend the Root Partition of a Quickly Provisioned

Released On	Description
	BMS?, Why Can't I Attach My Disk to a Server?, Can I Change the Function of a System Disk or Data Disk Created Along with a Server?, How Do I Obtain My Disk Name in the ECS OS Using the Device Identifier Provided on the Console?, Deletion, What Can I Do If the Capacity of My Disk Reaches the Maximum But I Still Need More Space?, Can I Transfer the Data Disk Capacity to a System Disk?, Performance, Snapshot, and Backup.
2022-08-08	This issue is the second official release, which incorporates the following change: • Optimized some descriptions.
2020-02-26	This issue is the first official release.