



## Elastic Volume Service

# User Guide

Date 2025-03-07

---

# Contents

---

<b>1 Overview.....</b>	<b>1</b>
1.1 What Is EVS?.....	1
1.2 Disk Types and Performance.....	4
1.3 Device Types.....	7
1.4 Permissions.....	9
1.5 Notes and Constraints.....	11
1.6 EVS and Other Services.....	16
1.7 Basic Concepts.....	18
1.7.1 EVS Concepts.....	18
1.7.2 Region and AZ.....	18
1.7.3 Three-Copy Redundancy.....	19
1.7.4 Disk Sharing.....	22
1.7.5 Disk Encryption.....	23
1.7.6 Disk Backup.....	23
1.7.7 Disk Snapshot.....	24
<b>2 Getting Started.....</b>	<b>25</b>
2.1 Quickly Creating an EVS Disk and Using It on a Linux Server.....	25
2.2 Quickly Creating an EVS Disk and Using It on a Windows Server.....	33
<b>3 Permissions Management.....</b>	<b>43</b>
3.1 Creating a User and Granting EVS Permissions.....	43
3.2 EVS Custom Policies.....	44
<b>4 Creating and Using an EVS Disk.....</b>	<b>46</b>
4.1 Overview.....	46
4.2 Creating an EVS Disk.....	48
4.3 Attaching an EVS Disk.....	56
4.3.1 Attaching a Non-Shared Disk.....	56
4.3.2 Attaching a Shared Disk.....	59
4.4 Initializing EVS Data Disks.....	62
4.4.1 Initialization Overview.....	62
4.4.2 Initializing a Linux Data Disk (Less Than or Equal to 2 TiB).....	64
4.4.3 Initializing a Linux Data Disk (Greater Than 2 TiB).....	70
4.4.4 Initializing a Windows Data Disk.....	75

<b>5 Viewing EVS Disk Details.....</b>	<b>84</b>
<b>6 Expanding the EVS Disk Capacity.....</b>	<b>89</b>
6.1 Expansion Overview.....	89
6.2 Step 1: Expand Disk Capacity.....	90
6.3 Step 2: Extend Disk Partitions and File Systems.....	93
6.3.1 Extending Disk Partitions and File Systems (Linux).....	93
6.3.2 Extending Disk Partitions and File Systems (Windows).....	111
<b>7 Detaching and Deleting an EVS Disk.....</b>	<b>129</b>
7.1 Detaching an EVS Disk.....	129
7.2 Deleting an EVS Disk.....	132
<b>8 Managing EVS Snapshots.....</b>	<b>134</b>
8.1 EVS Snapshot Overview.....	134
8.2 Using EVS Snapshots.....	137
8.2.1 Creating an EVS Snapshot.....	137
8.2.2 Rolling Back Disk Data from a Snapshot.....	140
8.2.3 Creating a Disk from a Snapshot.....	141
8.2.4 Checking EVS Snapshot Details.....	142
8.2.5 Deleting an EVS Snapshot.....	143
<b>9 Managing Encrypted EVS Disks.....</b>	<b>145</b>
<b>10 Managing Shared EVS Disks.....</b>	<b>150</b>
<b>11 Managing EVS Disk Backups.....</b>	<b>155</b>
11.1 CBR Overview.....	155
11.2 Backing Up EVS Disks.....	157
<b>12 Managing EVS Transfers.....</b>	<b>159</b>
<b>13 Managing EVS Tags.....</b>	<b>162</b>
13.1 Tag Overview.....	162
13.2 Adding a Tag.....	162
13.3 Modifying a Tag.....	164
13.4 Deleting a Tag.....	165
13.5 Searching for Disks by Tag.....	165
<b>14 Managing EVS Quotas.....</b>	<b>167</b>
14.1 Querying EVS Resource Quotas.....	167
14.2 Increasing EVS Resource Quotas.....	167
<b>15 Cloud Eye Monitoring.....</b>	<b>169</b>
15.1 Viewing Basic EVS Monitoring Data.....	169
15.2 Viewing EVS Monitoring Data Included in OS Metrics (with Agent Installed).....	171
<b>16 Recording EVS Operations Using CTS.....</b>	<b>174</b>
<b>17 FAQ.....</b>	<b>175</b>

17.1 Summary.....	175
17.2 General.....	177
17.2.1 How Do I Start Using a Newly Created Disk?.....	177
17.2.2 Can EVS Disks Be Used Directly for Storage?.....	177
17.2.3 Can EVS Disks Be Used Alone?.....	177
17.2.4 How Can I View My Disk Details?.....	177
17.2.5 How Do I Change the Name of My Disk?.....	178
17.2.6 Can I Change the AZ of My Disk?.....	179
17.2.7 Can I Change the Disk Type, Device Type, or Sharing Attribute of My Disk?.....	179
17.2.8 What Should I Do If an Error Occurs on My EVS Disk?.....	180
17.2.9 How Can I Obtain ECS NIC Information?.....	180
17.2.10 Why Do Some of My EVS Disks Not Have WWN Information?.....	181
17.2.11 How Can I Migrate Data from an EVS Disk?.....	181
17.2.12 What Are the Differences Between System Disks and Data Disks?.....	182
17.2.13 How Can I Export the Original Data After I Changed My Server OS from Windows to CentOS?.....	182
17.2.14 What Are the Differences Between MBR and GPT Partition Styles?.....	182
17.2.15 What Does the "reserveVolume" Trace Mean in CTS?.....	183
17.2.16 How Can I Download My EVS Disk Data to a Local PC?.....	184
17.2.17 How Can I Upload Files to My EVS Disk?.....	184
17.2.18 What EVS Disk Types Are Available?.....	184
17.3 Attachment.....	185
17.3.1 Why Can't I View the Attached Data Disk on the Server?.....	185
17.3.2 Why Can't I Attach My Disk to a Server?.....	188
17.3.3 Can I Attach a Disk to Multiple Servers?.....	190
17.3.4 Can I Attach a Disk to a Server in a Different AZ?.....	190
17.3.5 How Can I Add a Data Disk to an Existing Server?.....	191
17.3.6 Can I Attach Different Types of Disks to the Same Server?.....	191
17.3.7 What Should I Do If a Linux EVS Disk Is Attached to a Windows Server?.....	191
17.3.8 Can I Change the Function of a System Disk or Data Disk Created Along with a Server?.....	191
17.3.9 How Do I Obtain My Disk Device Name in the ECS OS Using the Device Identifier Provided on the Console?.....	192
17.4 Capacity Expansion.....	196
17.4.1 Can I Reduce or Temporarily Expand the Disk Capacity?.....	196
17.4.2 What Are the Differences Between Expanding Capacity by Expanding an EVS Disk and Creating a New EVS Disk?.....	196
17.4.3 Will My Disk Data Be Lost After I Expand the Disk Capacity?.....	196
17.4.4 Can I Use Backups or Snapshots Created Before Capacity Expansion to Restore Data on Expanded Disks?.....	196
17.4.5 Do I Need to Restart the Server After Expanding the Disk Capacity?.....	197
17.4.6 Do I Need to Detach an EVS Disk Before Expanding Its Capacity?.....	197
17.4.7 What Should I Do If My Disk Capacity Exceeds 2 TiB After Expansion?.....	197
17.4.8 How Can I Allocate Newly Added Space to a New Partition?.....	198
17.4.9 How Can I Allocate Newly Added Space to an Existing Partition?.....	198

17.4.10 Why Did My Disk Capacity Remain Unchanged on the Server After a Capacity Expansion?.....	199
17.4.11 Why Can't I Expand Capacity for My Disk?.....	199
17.4.12 How Do I Extend the File System of an Unpartitioned Data Disk in Linux?.....	199
17.4.13 How Do I Extend the Root Partition of a Quickly Provisioned BMS?.....	201
17.4.14 How Do I Check the Disk Partition Style in Linux?.....	204
17.4.15 How Do I Extend the Disk Partition and File System in Windows Server 2008?.....	207
17.4.16 How Do I Extend the Disk Partition and File System in Linux?.....	233
17.5 Detachment.....	237
17.5.1 If I Detach a Disk, Will I Lose the Data on My Disk?.....	238
17.5.2 Why Can't I Detach My Disk?.....	238
17.6 Deletion.....	239
17.6.1 How Can I Recover Data from a Disk That Was Accidentally Deleted?.....	239
17.7 Capacity.....	239
17.7.1 What Is the Maximum Capacity Supported for the System and Data Disks?.....	239
17.7.2 What Should I Do If My Disk Starts to Run Out of Space?.....	239
17.7.3 How Do I Clean Up My Disk Space on a Windows Server?.....	240
17.7.4 What Can I Do If the Capacity of My Disk Reaches the Maximum But I Still Need More Space?..	243
17.7.5 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?.....	243
17.7.6 How Can I View the Available Capacity of My Disk?.....	243
17.7.7 How Can I Monitor My Disk Usages?.....	251
17.7.8 Can I Transfer the Data Disk Capacity to a System Disk?.....	255
17.7.9 Why the Space of My New Disk Is Full After I Uploaded Only 500 MB of Files to the Disk?.....	256
17.8 Snapshot.....	257
17.8.1 What Are the Typical Causes of a Snapshot Creation Failure?.....	258
17.8.2 Does EVS Support Automatic Snapshot Creation?.....	258
17.8.3 Can I Create Snapshots for Multiple Disks at a Time?.....	258
17.8.4 How Is a Snapshot Created for My Disk?.....	258
17.8.5 Why Can't I Roll Back My Disk Data from a Snapshot?.....	258
17.8.6 Can I Roll Back Data from a Snapshot After Reinstalling the OS or Formatting a Disk?.....	259
17.8.7 How Is the Snapshot Size Calculated?.....	259
17.8.8 Do Snapshots Take Space on the Disk?.....	259
17.8.9 Can I Perform Multiple Rollback Operations for a Snapshot?.....	259
17.8.10 Can I Replicate Snapshots to Other Regions or Accounts?.....	259
17.8.11 Why Can't I Find My Snapshot?.....	259
17.8.12 Can I Use a Snapshot to Create a Disk and Roll Back Disk Data at Almost the Same Time?.....	259
17.8.13 Can I Modify Data in My EVS Snapshot?.....	260
17.9 Performance.....	260
17.9.1 How Do I Test My Disk Performance?.....	260
17.9.2 Why Does My Disk Performance Test Using Fio Have Incorrect Results?.....	267
17.9.3 How Can I Handle a Slowdown in Disk Read/Write Speed or Increased I/Os?.....	269
17.9.4 Why My Disk's Read IOPS Can't Reach the Theoretical Maximum IOPS When the Disk I/O Usage Is Almost 100%?.....	270

17.10 Sharing.....	271
17.10.1 Do I Have to Deploy a Cluster to Use Shared Disks?.....	271
17.10.2 How Many Servers Can I Attach a Shared Disk to?.....	271
17.10.3 How Can I Attach a Shared Disk to Multiple Servers?.....	271
17.10.4 Can a Shared Disk Be Attached to Servers That Belong to Different Accounts?.....	271
17.10.5 Can I Attach a Shared Disk to Servers Running Different OSs?.....	271
17.11 Backup.....	271
17.11.1 Do I Need to Stop the Server Before Performing a Disk Backup?.....	271
17.11.2 Can I Back Up and Restore My EVS Disk to a Different Region?.....	272
17.11.3 How Do I Check the Backup Data of My Disk?.....	272
<b>A Change History.....</b>	<b>273</b>

# 1 Overview

## 1.1 What Is EVS?


### Overview

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

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

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

**Figure 1-1** EVS architecture



## EVS Advantages

EVS has the following advantages.

**Table 1-1** EVS advantages

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

## Differences Among EVS, SFS, and OBS

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



**Table 1-2** Differences among EVS, SFS, and OBS

Service	Overall Introduction	Typical Application Scenarios	Storage Capacity
EVS	EVS provides scalable block storage that features high reliability, high performance, and a variety of specifications for servers.	<ul style="list-style-type: none"> <li>• Enterprise office applications</li> <li>• Development and testing</li> <li>• Enterprise applications, including SAP, Microsoft Exchange, and Microsoft SharePoint</li> <li>• Distributed file systems</li> <li>• Various databases, including MongoDB, Oracle, SQL Server, MySQL, and PostgreSQL</li> </ul>	EVS disks start at 10 GiB and can be expanded as required in 1 GiB increments to up to 32 TiB.
SFS	SFS provides completely hosted shared file storage for cloud servers. Compatible with the Network File System (NFS) protocol, SFS is expandable to pebibyte and seamlessly handles data-intensive and bandwidth-intensive applications.	<ul style="list-style-type: none"> <li>• HPC scenarios, such as gene sequencing, animation rendering, and CAD/CAE</li> <li>• File sharing</li> <li>• Media processing</li> <li>• Content management and web services</li> <li>• Offline file backup</li> </ul>	SFS storage capacity is available on demand and can be expanded to a maximum of 2 PiB.

Service	Overall Introduction	Typical Application Scenarios	Storage Capacity
OBS	OBS provides cloud storage for unstructured data, such as files, pictures, and videos. With multiple options for migration to the cloud, OBS provides low-cost, reliable storage access for massive data and supports online multimedia processing.	<ul style="list-style-type: none"><li>• Enterprise backup and archive</li><li>• Big data analysis</li><li>• Enterprise cloud box</li><li>• Static website hosting</li><li>• Cloud-native applications</li></ul>	OBS has limitless storage capacity, and storage resources are available for linear and nearly infinite expansion.

## Access Methods

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

- APIs  
Use APIs if you need to integrate EVS into a third-party system for secondary development. For details, see *Elastic Volume Service API Reference*.
- Console  
Use the console if you do not need to integrate EVS with a third-party system. Log in to the 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. You can choose whichever disk type that is the best fit 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	6	50
Max. IOPS/EVS disk	5,000	33,000
Baseline IOPS/EVS disk	1,200	1,500
Disk IOPS	Min. [5,000, 1,200 + 6 x Capacity (GiB)]	Min. [33,000, 1,500 + 50 x Capacity (GiB)]
IOPS burst limit/EVS disk	5,000	16,000
Max. throughput	150 MiB/s	350 MiB/s
API name <b>NOTE</b> This API name is the value of the <b>volume_type</b> parameter in the EVS API. It does not represent the type of the underlying hardware device.	SAS	SSD
Typical scenarios	Mainstream applications requiring high performance and high reliability, such as large-scale development and test environments, web server logs, and enterprise applications. Typical enterprise applications include SAP applications, Microsoft Exchange, and Microsoft SharePoint.	Read/write-intensive workloads that demand 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.

## Calculating Disk IOPS Limit

To calculate the IOPS limit of a disk, obtain the smaller value of the following two values:

- Max. IOPS/disk
- Baseline IOPS/disk + IOPS per GiB x Disk capacity

Take an ultra-high I/O EVS disk with a maximum IOPS of 33,000 for example.

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

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

The disk IOPS limit is 6,500, the smaller of the two values (33,000 and 6,500).

- If the disk capacity is 1,000 GiB, the disk IOPS limit is calculated as follows:

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

The disk IOPS limit is 33,000, the smaller of the two values (33,000 and 51,500).

## Disk Burst Capability and Principles

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

Disks with burst capability are well-suited for speeding up server startup. In most cases, system disks have small capacities. For example, the IOPS of a 50-GiB ultra-high I/O disk without burst capability can reach only 4,000 (1,500 + 50 x 50). If the disk has burst capability, its IOPS can burst up to 16,000. 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 ultra-high I/O disk without burst capability can reach only 4,300 (1,800 + 50 x 50). 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 6,500, but it can temporarily burst to 16,000 IOPS.
- If the disk capacity is 1,000 GiB, the disk has a maximum IOPS of 33,000. 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 1800 seconds. 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 6,500 tokens/s.
- Token consumption rate: This rate is based on the I/O usage. Each I/O request consumes a token. The maximum consumption rate is 16,000 tokens/s, which is the larger value of the disk burst IOPS and the maximum IOPS.

Consumption principles

When tokens are consumed faster than they are produced, the number of tokens decreases accordingly, and eventually the disk IOPS will be consistent with the token production rate (the maximum IOPS). In this example, the disk can burst for approximately 3,032 seconds [28,800,000/(16,000 - 6,500)].

Reservation principles

When tokens are consumed more slowly than they are produced, the number of tokens increases accordingly, and the disk regains burst capability. In this example, if the disk is suspended for approximately 4,431 seconds ( $28,800,000/6,500$ ), the token bucket will be filled up with tokens.


**NOTE**

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

- As long as there are tokens, the disk IOPS can exceed 6,500 and can burst to up to 16,000, the IOPS burst limit.
- When there are no more tokens, the disk loses the burst capability, and the disk IOPS can reach up to 6,500.
- Anytime the disk IOPS is less than 6,500, the number of tokens starts to increase, and the disk regains the burst capability.

**Figure 1-2** Burst capability diagram



## Performance Testing

For details about how to test the EVS disk performance, see [How Do I Test My Disk Performance?](#)

## 1.3 Device Types

### What Device Types Are Available?

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

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

Device type is configured during 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.

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

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

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

- Bare Metal Server (BMS)

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

- KVM ECS

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

### NOTE

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

- Xen ECS

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

However, a few Windows and Linux images support SCSI EVS disks on Xen ECSs. For the supported images, see [Table 1-4](#).

 **NOTE**

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

- Public Windows images are preinstalled with the Paravirtual SCSI (PVSCSI) driver. Therefore, no driver needs to be installed.
- Private Windows images are not preinstalled with the PVSCSI driver. You need to download and install it explicitly.

For details, see **(Optional) Optimizing Windows Private Images** in the *Image Management Service User Guide*.

- Linux images are not preinstalled with the PVSCSI driver. You need to obtain the source code of the open-source Linux driver at <https://github.com/UVP-Tools/SAP-HANA-Tools>, compile the code, and then install the driver.

**Table 1-4** OSs supporting SCSI EVS disks

Virtualization Type	OS	
Xen	Windows	See the Windows images listed on the <b>Public Images</b> page.  Log in to the console, choose <b>Image Management Service</b> , click the <b>Public Images</b> tab, and select <b>ECS image</b> and <b>Windows</b> from the drop-down lists, respectively.
	Linux	<ul style="list-style-type: none"><li>• SUSE Linux Enterprise Server 11 SP4 64bit (The kernel version is 3.0.101-68-default or 3.0.101-80-default.)</li><li>• SUSE Linux Enterprise Server 12 64bit (The kernel version is 3.12.51-52.31-default.)</li><li>• SUSE Linux Enterprise Server 12 SP1 64bit (The kernel version is 3.12.67-60.64.24-default.)</li><li>• SUSE Linux Enterprise Server 12 SP2 64bit (The kernel version is 4.4.74-92.35.1-default.)</li></ul>

## 1.4 Permissions

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

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

If your 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 .

## EVS Permissions

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

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

You can grant users permissions by using roles and policies.

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

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

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

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



Role/Policy Name	Description	Type	Dependencies
Server Administrator	Full permissions for EVS	System-defined role	None

**Table 1-6** lists the common operations supported by each system-defined policy of EVS. Select the policies as required.

**Table 1-6** Common operations supported by each system-defined policy of EVS

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

## Helpful Links

- Section "Service Overview" in the *Identity 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*

## 1.5 Notes and Constraints

This section describes the constraints on using EVS.

## Specifications

**Table 1-7** Specifications

Resource Type	Item	Description
Disk performance	Major disk performance metrics include IOPS, throughput, and latency.	Different types of EVS disks have different performance. For details, see <a href="#">Disk Types and Performance</a> .

## Security

**Table 1-8** Security

Item	Description
Disk encryption	<ul style="list-style-type: none"> <li>• The encryption attribute of a disk cannot be changed after the disk is created.</li> <li>• If you use an encrypted disk to create a backup, the backup generated will be an encrypted backup. You cannot modify the encryption attribute of the backup.</li> <li>• If you use an encrypted disk to create an image, the image generated will be an encrypted image. You cannot modify the encryption attribute of the image.</li> <li>• If you use an encrypted disk to create a snapshot, the snapshot generated will be an encrypted snapshot. You cannot modify the encryption attribute of the snapshot.</li> <li>• If you use an image that does not support lazyloading to create a disk, the disk created will be an encrypted disk. You cannot modify the encryption of the disk.</li> </ul>

## Quotas

You can log in to the console to view default quotas. You can to apply for a larger quota if needed.

**Table 1-9** Quotas

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

## Operations

**Table 1-10** Operations

Scenario	Item	Description
Disk creation	Maximum number of disks that can be created at a time	100
Disk creation	Disk creation from snapshot	<ul style="list-style-type: none"> <li>Batch disk creation from a snapshot is not supported.</li> <li>A disk created from a snapshot has the same device type (SCSI or VBD), encryption attribute, AZ, region, and disk type as the snapshot's source disk.</li> <li>A snapshot whose name starts with <b>autobk_snapshot_vbs_</b>, <b>manualbk_snapshot_vbs_</b>, <b>autobk_snapshot_csbs_</b>, or <b>manualbk_snapshot_csbs_</b> is automatically generated during backup. Such a snapshot can only be viewed. It cannot be used to create new disks.</li> </ul>

Scenario	Item	Description
Disk creation	Disk creation from backup	<ul style="list-style-type: none"> <li>• Batch creation is not supported. One can create only one disk from a backup at a time.</li> <li>• 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.</li> <li>• If a disk is created from a backup of a system disk, the new disk can be used as a data disk only.</li> </ul>
Disk creation	Disk creation from image	<ul style="list-style-type: none"> <li>• The device type of the new disk is the same as that of the image's source disk.</li> <li>• The encryption attribute of the new disk is the same as that of the image's source disk.</li> </ul>
Disk creation	Device type	The device type of a disk cannot be changed after the disk is created.
Disk creation	Disk sharing	The sharing attribute of a disk cannot be changed after the disk is created.
Disk creation	Disk encryption	The encryption attribute of a disk cannot be changed after the disk is created.
Disk attachment	Constraints on region and AZ	The disk and the server must be in the same region and AZ.
	Maximum number of servers that a non-shared disk can be attached to	1
	Maximum number of servers that a shared disk can be attached to	16
	Maximum number of disks that can be attached to an ECS	<p>This number varies depending on the ECS type.</p> <p>For details, see section "Can I Attach Multiple Disks to an ECS?" in the <i>Elastic Cloud Server User Guide</i>.</p> <p>For a BMS, a maximum of 60 EVS disks (1 system disk and 59 data disks) are supported.</p>

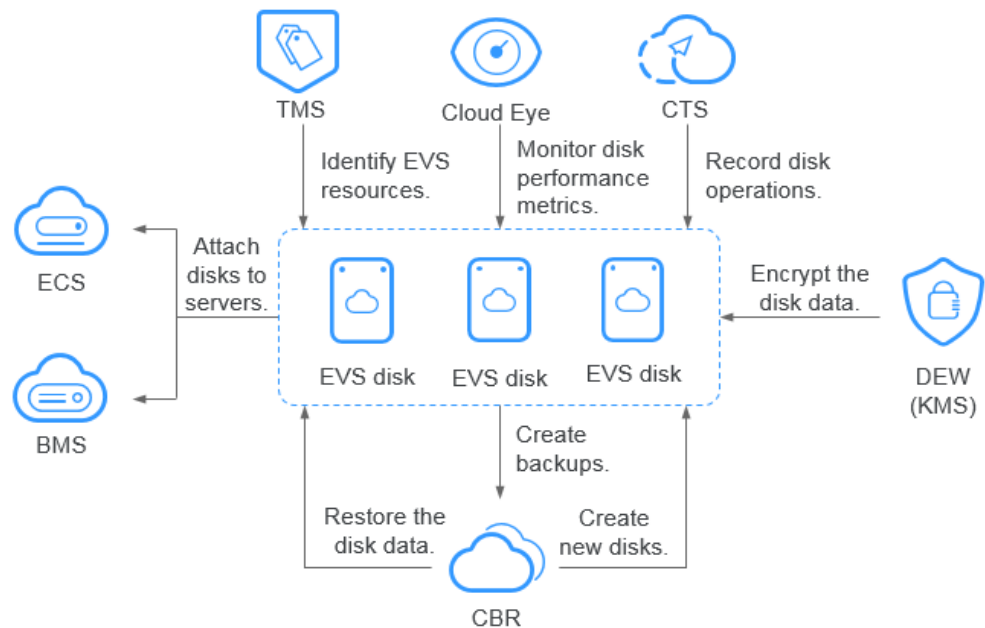
Scenario	Item	Description
	Device name	<ul style="list-style-type: none"> <li>System disk: /dev/vda, /dev/sda, and /dev/xvda</li> <li>Data disk: /dev/vd[b-z], /dev/sd[b-z], and /dev/xvd[b-z]</li> </ul>
Disk capacity expansion	Capacity expansion	Disk capacity can be expanded, but cannot be reduced.
	Capacity expansion of non-shared disks	Some server OSs support the capacity expansion of non-shared, In-use disks. For details, see .
	Capacity expansion of shared disks	A shared disk must be detached from all its servers before expansion. This means that the shared disk status must be <b>Available</b> .
	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 <b>Stopped</b> state.
	Data disk detachment	A data disk can be detached online or offline, meaning that its server can either be in the <b>Running</b> or <b>Stopped</b> state.
Disk deletion	Deletion of disks	<ul style="list-style-type: none"> <li>The disk status is <b>Available, Error, Expansion failed, Restoration failed, or Rollback failed</b>.</li> <li>The disk is not locked by any service.</li> <li>The shared disk has been detached from all its servers.</li> </ul>
	/	<ul style="list-style-type: none"> <li>Snapshots can be created for both system disks and data disks.</li> <li>Snapshots of encrypted disks are stored encrypted, and those of non-encrypted disks are stored non-encrypted.</li> <li>You can manually create a maximum of seven snapshots for a disk.</li> <li>The enterprise project of a snapshot is the same as that of the snapshot's source disk.</li> </ul>

Scenario	Item	Description
Snapshot data rollback to disks	/	<ul style="list-style-type: none"> <li>• Snapshot data can only be rolled back to source EVS disks. Rollback to a different disk is not possible.</li> <li>• You can only roll back disk data from a snapshot when the source disk status is <b>Available</b> (not attached to any server) or <b>Rollback failed</b>. If the source disk is attached, detach the disk first.</li> <li>• If a snapshot is being created, it cannot be used to roll back disk data.</li> <li>• A snapshot whose name starts with <b>autobk_snapshot_vbs_</b>, <b>manualbk_snapshot_vbs_</b>, <b>autobk_snapshot_csbs_</b>, or <b>manualbk_snapshot_csbs_</b> is automatically generated during backup. Such a snapshot can only be viewed. It cannot be used to roll back the disk data.</li> </ul>
	/	<ul style="list-style-type: none"> <li>• If a snapshot is deleted, disks rolled back or created from this snapshot are not affected.</li> <li>• If a snapshot's source disk is deleted, all snapshots of this disk will also be deleted.</li> <li>• If you reinstall or change the server OS, snapshots of the system disk are automatically deleted. Those of the data disks can be used as usual.</li> <li>• A snapshot whose name starts with <b>autobk_snapshot_vbs_</b>, <b>manualbk_snapshot_vbs_</b>, <b>autobk_snapshot_csbs_</b>, or <b>manualbk_snapshot_csbs_</b> is automatically generated during backup. You can only check details of such snapshots and cannot delete them.</li> </ul>

## 1.6 EVS and Other Services

Figure 1-3 shows the relationships between EVS and other services.

**Figure 1-3 Relationships between EVS and other services**



**Table 1-11 EVS and other services**

Interactive Function	Related Service	Reference
EVS disks can be attached to ECSs and used as scalable block storage devices.	ECS	<a href="#">Attaching an EVS Disk</a>
SCSI EVS disks can be attached to BMSs and used as scalable block storage devices.	BMS	<a href="#">Attaching an EVS Disk</a>
Backups can be created for EVS disks to guarantee the reliability and security of the server data.	CBR	<a href="#">Managing EVS Disk Backups</a>
EVS disk encryption depends on the KMS service in DEW. You can use keys provided by KMS to encrypt EVS disks (both system and data disks), thus improving EVS disk data security.	DEW	
After EVS is enabled, the performance metrics of monitored disks can be viewed through Cloud Eye without any additional plug-in installed. The monitored metrics include Disk Read Rate, Disk Write Rate, Disk Read Requests, and Disk Write Requests.	Cloud Eye	<a href="#">Viewing Basic EVS Monitoring Data</a>

Interactive Function	Related Service	Reference
Cloud Trace Service (CTS) records operations of EVS resources, facilitating user query, audit, and backtracking.	CTS	<a href="#">Recording EVS Operations Using CTS</a>
Tag Management Service (TMS) tags are used to identify EVS resources for purposes of easy categorization and quick search.	TMS	<a href="#">Adding a Tag</a>

## 1.7 Basic Concepts

### 1.7.1 EVS Concepts

Table 1-12 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	A small-capacity disk can surpass its official maximum IOPS for a short period of time.
VBD	A device type of EVS disks. VBD EVS disks only support basic SCSI read/write commands.
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.

### 1.7.2 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.

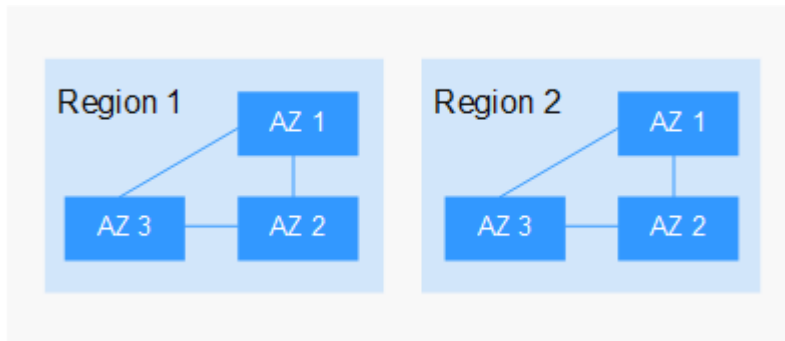
- A region is a physical data center, which is completely isolated to improve fault tolerance and stability. The region that is selected during resource creation cannot be changed after the resource is created.



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

**Figure 1-4** shows the relationship between regions and AZs.

**Figure 1-4** 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](#).

## 1.7.3 Three-Copy Redundancy

### What Is the Three-Copy Redundancy?

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


Three-copy redundancy has the following characteristics:

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

- The storage system guarantees strong consistency between the data copies.

For example, for data block P1 on physical disk A of server A, the storage system backs up its data to P1'' on physical disk B of server B and to P1' on physical disk C of server C. Data blocks P1, P1', and P1'' are the three copies of the same data block. If physical disk A where P1 resides is faulty, P1' and P1'' can continue providing storage services, ensuring service continuity.

**Figure 1-5** Three-copy redundancy



## How Does the Three-Copy Redundancy Keep Data Consistency?

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

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

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

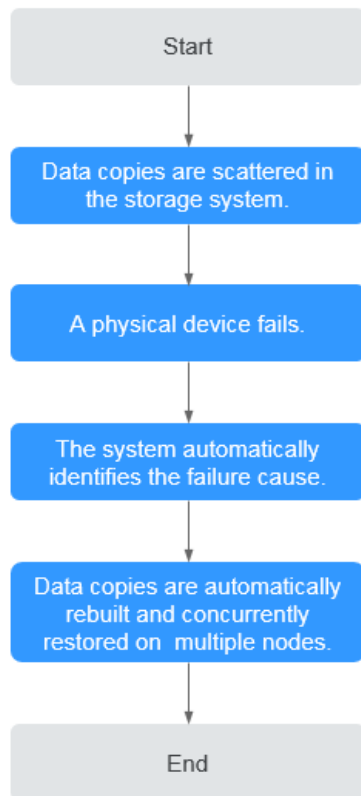
When an application fails to read data, the system automatically identifies the failure cause. If the data cannot be read from a physical disk sector, the system reads the data from another copy of the data on another node and writes it back to the original disk sector. This ensures the correct number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

## How Does Three-Copy Redundancy Rapidly Rebuild Data?

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

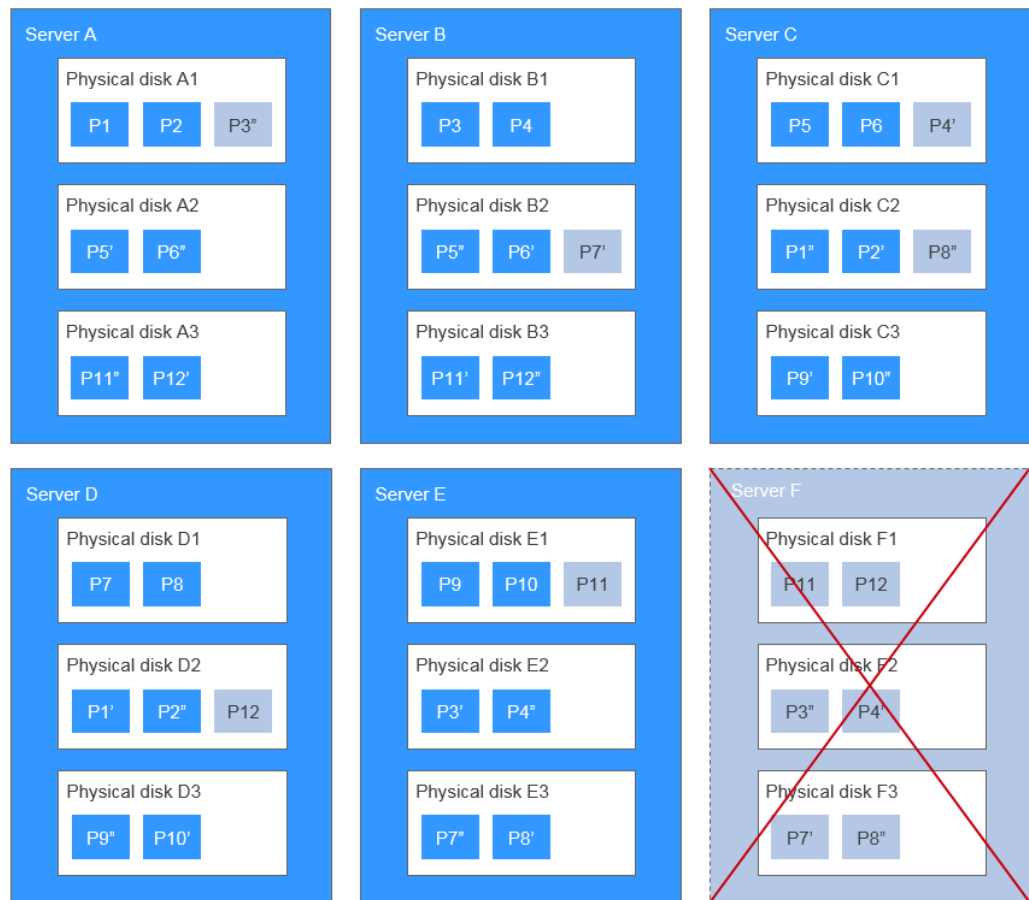
**Figure 1-6** shows the data rebuild process.

**Figure 1-6** Data rebuild process



**Figure 1-7** 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 1-7** Data rebuild principle



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

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

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

### 1.7.4 Disk Sharing


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

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

**NOTICE**

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

**Figure 1-8** Application scenario of shared EVS disks



## 1.7.5 Disk Encryption

EVS enables you to encrypt data on newly created disks as required. Keys used to encrypt EVS disks are provided by the Key Management Service (KMS) of Data Encryption Workshop (DEW), which is secure and convenient. You do not need to establish and maintain the key management infrastructure.

**NOTICE**

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

## 1.7.6 Disk Backup

Cloud Backup and Recovery (CBR) provides the cloud disk backup function, which allows you to back up EVS disks while servers are running. In case of a virus, accidental deletion, or software or hardware fault, you can restore data to any backup point to guarantee data integrity and security.

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

## 1.7.7 Disk Snapshot

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

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

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

### Snapshot Usage 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 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 data from snapshots. However, you can create disk B using an existing snapshot of system disk A and attach disk B to a properly running server, for example server B. In this case, server B 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 source 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. 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 Restrictions

For details about the limitations on using snapshots, see [Notes and Constraints](#) or the "Notes and Constraints" part in the corresponding snapshot section.

### Usage Instructions

For details about how to use snapshots, see [Managing EVS Snapshots](#).

# 2 Getting Started

## 2.1 Quickly Creating an EVS Disk and Using It on a Linux Server

### Scenarios

You can use EVS disks as system disks or data disks. System disks are created together with servers, while data disks can be created together with servers or separately. If you have created data disks separately, you must attach and initialize them before they can be used.


This section describes how a non-shared data disk can be created on the EVS console, attached to a Linux server, and initialized on the server.

### Operation Process

Procedure	Description
<a href="#">Step 1: Create an EVS Disk</a>	Create a data disk on the EVS console.
<a href="#">Step 2: Attach the EVS Disk</a>	Attach the data disk to a Linux server.
<a href="#">Step 3: Initialize the EVS Disk</a>	Initialize the data disk on the server.

### Step 1: Create an EVS Disk

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

1. Log in to the console.
2. Click  in the upper left corner of the page and choose **Storage > Elastic Volume Service**.

3. Click **Create Disk** in the upper right corner.

**Step 2** Configure mandatory parameters based on the following table and retain the default settings for other parameters.

Parameter	Example Value	Description
Region		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	-	You can only attach EVS disks to servers in the same AZ. After a disk is created, its AZ cannot be changed.
Data Source	Do not configure it.	If you want to create an empty data disk, do not configure a data source.
Disk Type	Ultra-high I/O	To learn more about disk types, see <a href="#">Disk Types and Performance</a> .
Capacity	100 GiB	Enter a disk capacity.
Automatic Backup	Do not use	Automatic backup allows you to back up the disk data to ensure your data security and integrity.
More > Share	Do not selection this option.	A non-shared disk can only be attached to one server. The sharing attribute of a disk cannot be changed after the disk has been created.
More > SCSI	Select this option.	A SCSI disk allows the server OS to directly access the underlying storage media and send SCSI commands to the disk. The device type of a disk cannot be changed after the disk has been created.
More > Encryption	Select this option and use the default key.	EVS uses the industry-standard XTS-AES-256 cryptographic algorithm and keys to encrypt EVS disks. The encryption attribute of a disk cannot be changed after the disk has been created.
Enterprise Project	default	This parameter shows up only when you use an enterprise account to create disks. It enables unified management of cloud resources by project.
Disk Name	volume-0001	Enter a disk name.



Parameter	Example Value	Description
Quantity	1	The preset disk quantity is <b>1</b> , which means only one disk is created.

**Step 3** Click **Next**.

**Step 4** Go back to the disk list page. When the status of the **volume-0001** disk changes to **In-use**, the disk is successfully created.

----End

## Step 2: Attach the EVS Disk

EVS disks cannot be used alone. You need to attach them to cloud servers first.

**Step 1** In the disk list, find the **volume-0001** disk and click **Attach** in the **Operation** column.

**Step 2** Attach the **volume-0001** disk to your desired server. Ensure that the server and disk are in the same AZ.

**Step 3** Click **OK** to go back to the disk list page. When the status of the **volume-0001** disk changes to **In-use**, the disk is successfully attached.

----End

## Step 3: Initialize the EVS Disk

After attaching the **volume-0001** disk, you need to initialize it before it can be used. The following example uses `fdisk` to format the disk into two primary MBR partitions, with one 40 GiB and the other 60 GiB.

**Step 1** Log 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*.

**Step 2** Create two primary partitions, **/dev/vdb1** and **/dev/vdb2** for data disk **/dev/vdb**.

1. Check that the capacity of the **/dev/vdb** data disk is 100 GiB.

### lsblk

```
[root@ecs-centos76 ~]# 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 100G 0 disk
```

2. Create the first primary partition **/dev/vdb1**.

### fdisk /dev/vdb

n

p

1

 **NOTE**

- Entering **p** for **Partition type** creates a primary partition, and entering **e** creates an extended partition.
- Value **1** is the primary partition number.

```
[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): n
Partition type:
  p primary (0 primary, 0 extended, 4 free)
  e extended
Select (default p): p
Partition number (1-4, default 1): 1
```

Set **First sector** to **2048** and **Last sector** to **83886079** for partition **/dev/vdb1** (40 GiB).

```
First sector (2048-209715199, default 2048): 2048
Last sector, +sectors or +size{K,M,G} (2048-209715199, default 209715199):83886079
Partition 1 of type Linux and of size 40 GB is set
```

3. Create the second primary partition **/dev/vdb2**.

**n**

**p**

**2**

```
Command (m for help): n
Partition type:
  p primary (0 primary, 0 extended, 4 free)
  e extended
Select (default p): p
Partition number (1-4, default 2): 2
```

Set the **First sector** to **83886080** and **Last sector** to **209715199** for partition **/dev/vdb2**.

```
First sector (83886080-209715199, default 83886080): 83886080
Last sector, +sectors or +size{K,M,G} (83886080-209715199, default 209715199):209715199
Partition 2 of type Linux and of size 60 GB is set
```

 NOTE

First and last sectors of the partitions in this example are calculated as follows:

**Sector value = Capacity/512 bytes, 1 GiB = 1073741824 bytes**

- **First sector (2048-209715199, default 2048)** shows the sector value range of the `/dev/vdb` data disk (100 GiB).

First sector = 2048

Last sector = Sector value - 1 =  $(100 \times 1073741824/512) - 1 = 209715200 - 1 = 209715199$

- For the first partition `/dev/vdb1` (40 GiB) of the `/dev/vdb` data disk:

First sector = 2048 (The start sector of the `/dev/vdb` data disk is used.)

Last sector = Sector value - 1 =  $(40 \times 1073741824/512) - 1 = 83886079$

- For the second partition `/dev/vdb2` (60 GiB) of the `/dev/vdb` data disk:

First sector = Last sector of `/dev/vdb1` + 1 =  $83886079 + 1 = 83886080$

Last sector = First sector + Sector value - 1 =  $83886080 + (60 \times 1073741824/512) - 1 = 209715199$

### Step 3 Check the sizes and partition styles of the new partitions.

1. Check whether the partitioning is successful.

**p**

Command (m for help): p

```
Disk /dev/vdb: 107.4 GB, 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: 0x994727e5
```

Device	Boot	Start	End	Blocks	Id	System
/dev/vdb1		2048	83886079	41942016	83	Linux
/dev/vdb2		83886080	209715199	62914560	83	Linux

Command (m for help):

 NOTE

In case that you want to discard the changes made before, you can exit `fdisk` by entering `q` and press **Enter**. Then, re-create the partitions by referring to step 1.

2. Write the changes to the partition table and synchronize the new partition table to the OS.

**w**

**partprobe**

 NOTE

If error message **-bash: partprobe: command not found** is returned, the system cannot identify the command. In this case, run **yum install -y parted** to install the command. Then run the command again.

3. Confirm that the partition style is MBR.

**parted /dev/vdb**

**p**

 NOTE

If **Partition Table: msdos** is returned, the partition style is MBR.

```
[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
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 107GB
Sector size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:

Number  Start  End    Size  Type  File system  Flags
  1      1049kB 42.9GB 42.9GB primary
  2      42.9GB 107GB  64.4GB primary

(parted) q
[root@ecs-test-0001 ~]#
```

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

**Step 4** Create ext4 file systems for partitions **/dev/vdb1** (40 GiB) and **/dev/vdb2** (60 GiB).

```
mkfs -t ext4 /dev/vdb1
```

```
mkfs -t ext4 /dev/vdb2
```

#### NOTE

It takes some time to create file systems. Do not exit before the system returns the following information:

```
[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
2621440 inodes, 10485504 blocks
524275 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
```

Check whether the file system format is ext4.

```
parted /dev/vdb
```

```
p
```

```
[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
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 107GB
Sector size (logical/physical): 512B/512B
```

```
Partition Table: msdos
Disk Flags:

Number Start  End    Size  Type  File system  Flags
 1    1049kB 42.9GB 42.9GB primary ext4
 2    42.9GB 107GB  64.4GB primary ext4

(parted) q
[root@ecs-test-0001 ~]#
```

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

- Step 5** Create directories (mount points) and mount the new partitions on the created mount points.

```
mkdir -p /mnt/sdc
```

```
mkdir -p /mnt/sdd
```

```
mount /dev/vdb1 /mnt/sdc
```

```
mount /dev/vdb2 /mnt/sdd
```

```
lsblk
```

View the mount results.

```
[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 40G  0 part /mnt/sdc
└─vdb2 253:18  0 60G  0 part /mnt/sdd
```

You should now see that partitions **/dev/vdb1** and **/dev/vdb2** are mounted on **/mnt/sdc** and **/mnt/sdd**.

- Step 6** Use the partition UUIDs to configure auto mount at startup.

 **NOTE**

- Mounts become invalid after a system reboot. You can configure auto mount at startup by adding information of the new partition into the **/etc/fstab** file.
- You are advised not to use device names to identify disks in the **/etc/fstab** file because device names are assigned dynamically and may change (for example, from **/dev/vdb1** to **/dev/vdb2**) after a stop or start. This can even prevent your server from booting up.
- UUIDs are the unique character strings for identifying partitions in Linux.
- This operation will not affect the existing data on the ECS.

1. Query the partition UUIDs.

```
blkid /dev/vdb1
```

```
blkid /dev/vdb2
```

```
[root@ecs-test-0001 ~]# blkid /dev/vdb1
/dev/vdb1: UUID="0b3040e2-1367-4abb-841d-ddb0b92693df" TYPE="ext4"
/dev/vdb2: UUID="0d6769k2-1745-9dsf-453d-hgd0b34267dj" TYPE="ext4"
```

The UUIDs of partitions **/dev/vdb1** and **/dev/vdb2** are **0b3040e2-1367-4abb-841d-ddb0b92693df** and **0d6769k2-1745-9dsf-453d-hgd0b34267dj**.

2. Configure auto mount at startup.

```
vi /etc/fstab
```

Press **i** to enter the editing mode, move the cursor to the end of the file, press **Enter**, and add the following content:

```
UUID=0b3040e2-1367-4abb-841d-ddb0b92693df /mnt/sdc ext4 defaults 0 2
UUID=0d6769k2-1745-9dsf-453d-hgd0b34267dj /mnt/sdd ext4 defaults 0 2
```

Press **Esc**, enter **:wq**, and press **Enter** to save the settings and exit the vi editor.

**Table 2-1** Parameter description

Example Value	Description
UUID=0b3040e2-1367-4abb-841d-ddb0b92693df	The UUID of the partition.
/mnt/sdc	The mount point of the partition.
ext4	The file system format of the partition.
defaults	The partition mount option. Normally, this parameter is set to <b>defaults</b> .
0	<ul style="list-style-type: none"> <li>- The Linux dump backup option. <ul style="list-style-type: none"> <li>▪ <b>0</b>: Linux dump backup is not used. Usually, dump backup is not used, and you can set this parameter to <b>0</b>.</li> <li>▪ <b>1</b>: Linux dump backup is used.</li> </ul> </li> </ul>
2	<ul style="list-style-type: none"> <li>- The fsck option, which means whether to use fsck to check the disk during startup. <ul style="list-style-type: none"> <li>▪ <b>2</b>: The check starts from the partitions whose mount points are non-root directories. / is the root directory.</li> <li>▪ <b>1</b>: The check starts from the partitions whose mount points are root directories.</li> <li>▪ <b>0</b>: The fsck option is not used.</li> </ul> </li> </ul>

**Step 7** Verify that auto mount takes effect.

```
umount /dev/vdb1
```

```
umount /dev/vdb2
```

```
mount -a
```

The system reloads all the content in the **/etc/fstab** file.

Query file system mounting information.

```
mount | grep /mnt/sdc
```

```
mount | grep /mnt/sdd
```

If information similar to the following is displayed, auto mount has taken effect:

```
root@ecs-test-0001 ~]# mount | grep /mnt/sdc
/dev/vdb1 on /mnt/sdc type ext4 (rw,relatime,data=ordered)
root@ecs-test-0001 ~]# mount | grep /mnt/sdd
/dev/vdb2 on /mnt/sdd type ext4 (rw,relatime,data=ordered)
```

----End

You can use the disk after it is initialized.

## 2.2 Quickly Creating an EVS Disk and Using It on a Windows Server

### Scenarios

You can use EVS disks as system disks or data disks. System disks are created together with servers, while data disks can be created together with servers or separately. If you have created data disks separately, you must attach and initialize them before they can be used.


This section describes how a non-shared data disk can be created on the EVS console, attached to a Windows server, and initialized on the server.

### Operation Process

Procedure	Description
<a href="#">Step 1: Create an EVS Disk</a>	Create a data disk on the EVS console.
<a href="#">Step 2: Attach the EVS Disk</a>	Attach the data disk to a Windows server.
<a href="#">Step 3: Initialize the EVS Disk</a>	Initialize the data disk on the server.

### Step 1: Create an EVS Disk

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

1. Log in to the console.
2. Click  in the upper left corner of the page and choose **Storage > Elastic Volume Service**.
3. Click **Create Disk** in the upper right corner.

**Step 2** Configure mandatory parameters based on the following table and retain the default settings for other parameters.

Parameter	Example Value	Description
Region		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	-	You can only attach EVS disks to servers in the same AZ. After a disk is created, its AZ cannot be changed.
Data Source	Do not configure it.	If you want to create an empty data disk, do not configure a data source.
Disk Type	Ultra-high I/O	To learn more about disk types, see <a href="#">Disk Types and Performance</a> .
Capacity	100 GiB	Enter a disk capacity.
Automatic Backup	Do not use	Automatic backup allows you to back up the disk data to ensure your data security and integrity.
More > Share	Do not selection this option.	A non-shared disk can only be attached to one server. The sharing attribute of a disk cannot be changed after the disk has been created.
More > SCSI	Select this option.	A SCSI disk allows the server OS to directly access the underlying storage media and send SCSI commands to the disk. The device type of a disk cannot be changed after the disk has been created.
More > Encryption	Select this option and use the default key.	EVS uses the industry-standard XTS-AES-256 cryptographic algorithm and keys to encrypt EVS disks. The encryption attribute of a disk cannot be changed after the disk has been created.
Enterprise Project	default	This parameter shows up only when you use an enterprise account to create disks. It enables unified management of cloud resources by project.
Disk Name	volume-0001	Enter a disk name.
Quantity	1	The preset disk quantity is <b>1</b> , which means only one disk is created.



**Step 3** Click **Next**.

**Step 4** Go back to the disk list page. When the status of the **volume-0001** disk changes to **In-use**, the disk is successfully created.

----End

## Step 2: Attach the EVS Disk

EVS disks cannot be used alone. You need to attach them to cloud servers first. In the following example, the **volume-0001** disk is attached.

**Step 1** In the disk list, find the **volume-0001** disk and click **Attach** in the **Operation** column.

**Step 2** Attach the **volume-0001** disk to your desired server. Ensure that the server and disk are in the same AZ.

**Step 3** Click **OK** to go back to the disk list page. When the status of the **volume-0001** disk changes to **In-use**, the disk is successfully attached.

----End

## Step 3: Initialize the EVS Disk

After attaching the **volume-0001** disk, you need to initialize it before it can be used. In the following example, the disk is formatted into a 100 GiB GPT partition with the NTFS file system.

**Step 1** Log 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*.


**Step 2** On the desktop of the server, click the start icon in the lower left corner.

The **Windows Server** window is displayed.

**Step 3** Click **Server Manager**.


The **Server Manager** window is displayed.

Figure 2-1 Server Manager



**Step 4** In the upper right corner, choose **Tools > Computer Management**.

Figure 2-2 Computer Management



**Step 5** 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.

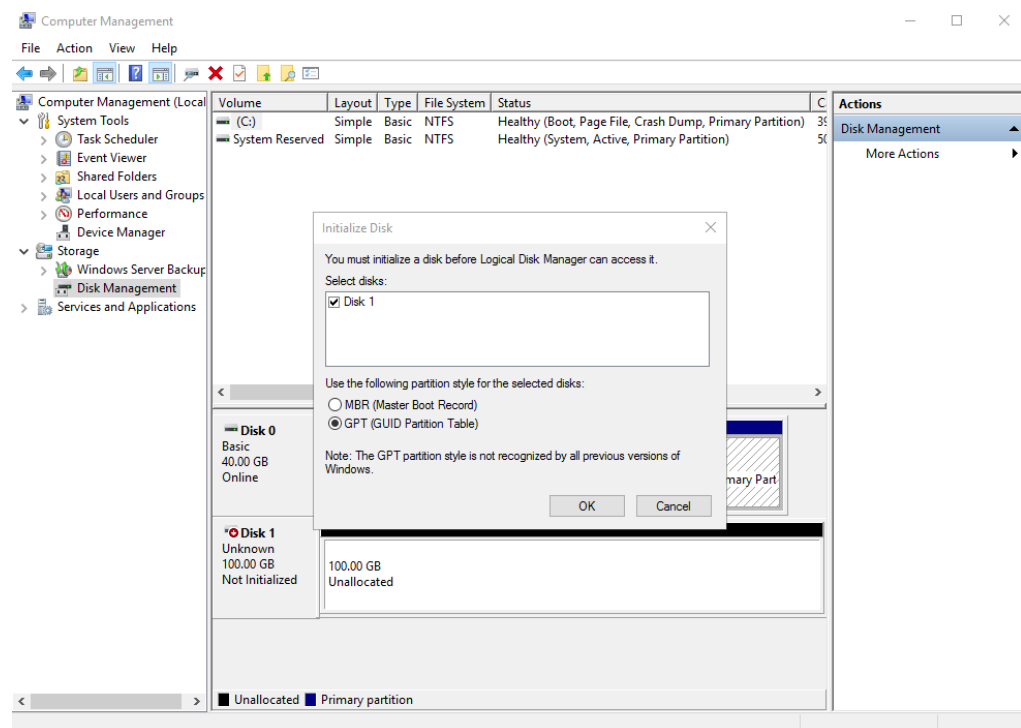
In the **Initialize Disk** dialog box, the to-be-initialized disk is selected. Select a partition style and click **OK**. In this example, **GPT (GUID Partition Table)** is selected.

**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 greater than 2 TiB.


If the partition style of an in-use disk is changed, 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, it is recommended that you back up the disk data before the change.

**Figure 2-3** Disk list




**Step 6** In the **Unallocated** area of **Disk 1**, right-click the blank area and choose **New Simple Volume**.

Figure 2-4 Computer Management




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

Figure 2-5 New Simple Volume Wizard




**Step 7** Click **Next** to go to the **Specify Volume Size** page.

Figure 2-6 Specify Volume Size



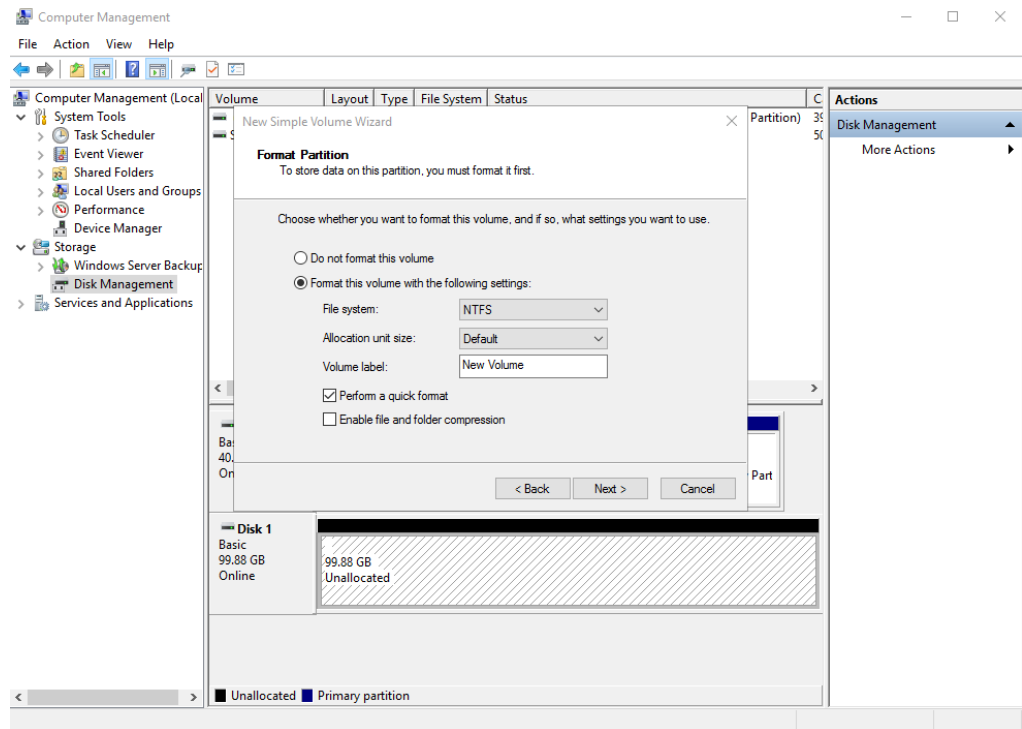
**Step 8** Specify the volume size and click **Next**. The system selects the maximum volume size by default. You can specify the volume size as required. In this example, the default setting is used.

Figure 2-7 Assign Drive 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.

Figure 2-8 Format Partition



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

Figure 2-9 Completing the New Simple Volume Wizard



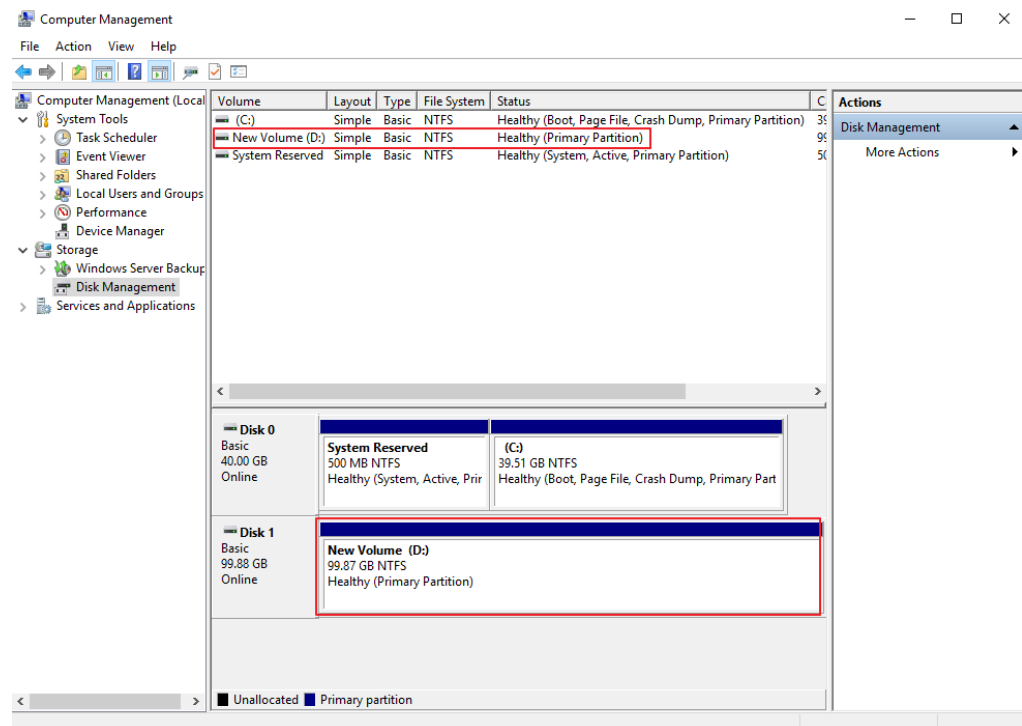
**NOTICE**


The partition sizes supported by file systems vary. Choose an appropriate file system format 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 succeeded.


**Figure 2-10** Disk initialized



**Step 12** After the volume is created, click  on the task bar and check whether a new volume appears in the File Explorer. 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 2-11** File Explorer



----End

You can use the disk after it is initialized.



# 3 Permissions Management

---

## 3.1 Creating a User and Granting EVS Permissions

You can use IAM for fine-grained permissions control for your EVS resources. With IAM, you can:

- Create IAM users for personnel based on your enterprise's organizational structure. Each IAM user has their own identity 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 efficient O&M on your EVS resources.

If your account does not require individual IAM users, you may skip over this section.

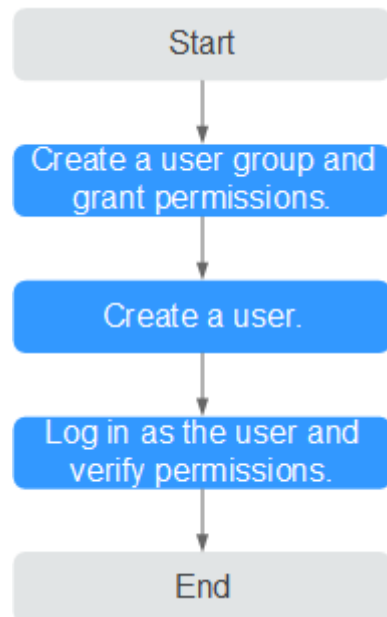
This section describes the procedure for granting permissions (see [Figure 3-1](#)).

### Prerequisites

Before granting permissions to user groups, learn about system-defined permissions in "Service Overview" > "Permissions" of the Elastic Volume Service.

## Process Flow

**Figure 3-1** Process for granting EVS permissions



1. On the IAM console, create a user group and grant it permissions (**EVS ReadOnlyAccess** as an example).
2. Create an IAM user and add it to the created user group.  
Create a user on the IAM console and add the user to the group created in **1**.
3. Log in as the IAM user and verify permissions.  
In the authorized region, perform the following operations:
  - Choose **Service List** > **Elastic Volume Service**. Then click **Create Disk** on the EVS console. If a message appears indicating that you have insufficient permissions to perform the operation, the **EVS ReadOnlyAccess** policy is in effect.
  - Choose another service from **Service List**. If a message appears indicating that you have insufficient permissions to access the service, the **EVS ReadOnlyAccess** policy is in effect.

## 3.2 EVS Custom Policies

You can create custom policies 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:


```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": [
        "evs:volumes:delete"
      ]
    }
  ]
}
```

# 4 Creating and Using an EVS Disk

## 4.1 Overview

Figure 4-1 shows the process for using EVS.

Figure 4-1 Process overview



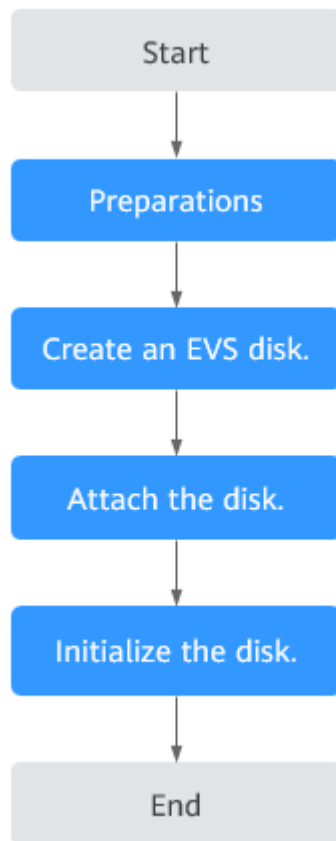
EVS disks can be attached to servers to be used as system disks or data disks. For details, see [Table 4-1](#).

Table 4-1 Method for creating disks

Function	Description	Method
System disk	System disks are created together with servers. You cannot create them separately.	<ul style="list-style-type: none"> <li>See section "Creating an ECS" in the <i>Elastic Cloud Server User Guide</i>.</li> <li>See section "Creating a BMS" in the <i>Bare Metal Server User Guide</i>.</li> </ul>
Data disk	You can create data disks together with servers or separately.	<ul style="list-style-type: none"> <li>See section "Creating an ECS" in the <i>Elastic Cloud Server User Guide</i>.</li> <li>See section "Creating a BMS" in the <i>Bare Metal Server User Guide</i>.</li> <li><a href="#">Creating an EVS Disk</a></li> </ul>

**Figure 4-2** shows the process of creating and using a data disk.

**Figure 4-2** Process overview



1. **Make preparations:** Register an account on the console and obtain permissions required for creating ECSs and EVS disks.
2. **Create an EVS disk:** Configure the disk parameters, including the disk type, capacity, name, and other information by referring to [Creating an EVS Disk](#).
3. **Attach the data disk:** Attach the separately created disk to an ECS by referring to [Attaching an EVS Disk](#).
4. **Initialize the data disk:** After the data disk is attached, log in to the ECS and initialize the disk before using it. For details about how to initialize the disk, see the following sections:
  - [Initialization Overview](#)
  - [Initializing a Linux Data Disk \(Less Than or Equal to 2 TiB\)](#)
  - [Initializing a Linux Data Disk \(Greater Than 2 TiB\)](#)
  - [Initializing a Linux Data Disk \(Less Than or Equal to 2 TiB\)](#)

## 4.2 Creating an EVS Disk

### Scenarios

You can use EVS disks as system disks or data disks for servers. You can create data disks on the EVS console, or create them together with system disks on the cloud server console.

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


### Notes and Constraints


**Table 4-2** Constraints on creating disks

Create On	Description
EVS console	<ul style="list-style-type: none"><li>• Disks created on the EVS console are data disks. You need to manually attach them to servers.</li><li>• Disks can only be attached to servers in the same region and AZ. Once created, the region and AZ cannot be changed.</li><li>• There are quantity and capacity quotas on EVS disks, so properly plan the number of disks and total disk capacity your workloads require. For details, see <a href="#">Managing EVS Quotas</a>.</li></ul>
Cloud server console	<ul style="list-style-type: none"><li>• System disks can only be created together with servers and are automatically attached.</li><li>• Data disks created together with servers or added after the server creation are automatically attached.</li><li>• Disks will have the same billing mode as their server if the disks are created together with the server.</li><li>• Disks created together with BMSs are SCSI disks.</li></ul>
-	Capacities of multiple disks cannot be combined, and the capacity of a single disk cannot be split.

### Procedure

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

**Step 2** Click  in the upper left corner and select the desired region and project.


**Step 3** Click  in the upper left corner and choose **Storage > Elastic Volume Service**.

The **Elastic Volume Service** page is displayed.

**Step 4** In the upper right corner, click **Create Disk**.

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

**Figure 4-3** Non-encrypted disks



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

**Figure 4-4** Encrypted disks



**Step 5** Configure disk parameters according to **Table 4-3**.

**Table 4-3** Disk parameters

Parameter	Sub-Parameter	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. <b>NOTE</b> <ul style="list-style-type: none"> <li>Disks can only be attached to the servers in the same AZ.</li> <li>The AZ of a disk cannot be changed after the disk has been created.</li> </ul>	-

Parameter	Sub-Parameter	Description	Example Value
Data Source (Optional)	Create from <ul style="list-style-type: none"> <li>• Backup</li> <li>• Snapshot</li> <li>• Image</li> </ul>	Optional <ul style="list-style-type: none"> <li>• <b>Create from Backup:</b> The backup data is used to create the disk. Click <b>Create from</b> and choose <b>Backup</b>. On the displayed page, select the target backup and click <b>OK</b>.               <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>- 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.</li> <li>- If a disk is created from a backup of a system disk, the new disk can be used as a data disk only.</li> </ul> </li> <li>• <b>Create from Snapshot:</b> The snapshot data is used to create the disk. Click <b>Create from</b> and choose <b>Snapshot</b>. On the displayed page, select the target snapshot and click <b>OK</b>.               <p><b>NOTE</b></p> <p>For details about how to create disks from snapshots, see <a href="#">Creating a Disk from a Snapshot</a>.</p> </li> <li>• <b>Create from Image:</b> The image data is used to create the disk. Click <b>Create from</b> and choose <b>Image</b>. On the displayed page, select the target image and click <b>OK</b>.               <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>- The device type of the new disk is the same as that of the image's source disk.</li> <li>- The encryption attribute of the new disk is the same as that of the image's source disk.</li> </ul> </li> </ul>	<ul style="list-style-type: none"> <li>• Create from Backup: autobackup-001</li> </ul>



Parameter	Sub-Parameter	Description	Example Value
Disk Specifications	Disk Type	<p>Mandatory</p> <p>EVS disk types vary depending on regions. See the EVS types displayed on the console.</p> <p>For more information about disk types, see <a href="#">Disk Types and Performance</a>.</p>	Ultra-high I/O
Disk Specifications	Capacity (GiB)	<p>Mandatory</p> <p>The disk capacity. Only data disks can be created on the current page, and the disk size ranges from 10 GiB to 32,768 GiB.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>When you use a backup to create a disk, the disk capacity must be greater than or equal to the backup size. In the condition that you do not specify a disk capacity, if the backup size is smaller than 10 GiB, the default capacity 10 GiB will be used as the disk capacity; if the backup size is greater than 10 GiB, the disk capacity will be consistent with the backup size.</li> <li>When you use a snapshot to create a disk, the disk capacity must be greater than or equal to the snapshot size. In the condition that you do not specify a 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.</li> </ul>	100 GiB

Parameter	Sub-Parameter	Description	Example Value
More	Advanced Settings <ul style="list-style-type: none"> <li>• Share</li> <li>• SCSI</li> <li>• Encryption</li> </ul>	<p>Optional</p> <ul style="list-style-type: none"> <li>• <b>Share</b> If you select <b>Share</b>, a shared disk is created. A shared disk can be attached to up to 16 servers. If you do not select <b>Share</b>, a non-shared disk is created, and the disk can be attached to one a server only.  If you select both <b>SCSI</b> and <b>Share</b>, a shared SCSI disk is created.</li> </ul> <p><b>NOTE</b> The sharing attribute of a disk cannot be changed after the disk is created.  For details about shared EVS disks, see <a href="#">Managing Shared EVS Disks</a>.</p> <ul style="list-style-type: none"> <li>• <b>SCSI</b> If you select <b>SCSI</b>, 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 <b>SCSI</b>, a VBD disk is created. That said, the disk device type is VBD, the default device type.</li> </ul> <p><b>NOTE</b> 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 <a href="#">Device Types</a>.</p> <ul style="list-style-type: none"> <li>• <b>Encryption</b> This option is only used to encrypt data disks, and you need to create an agency to grant KMS access rights to EVS. After the access rights are granted, configure the KMS key name on the <b>Encryption Settings</b> page displayed: A key name is the identifier of the key, and you can use <b>KMS Key Name</b> to specify a KMS key and use it for encryption. You can select one of the following keys:               <ul style="list-style-type: none"> <li>– Default Master Key: After the KMS access rights have been granted to EVS, the system</li> </ul> </li> </ul>	-

Parameter	Sub-Parameter	Description	Example Value
		<p>automatically creates a Default Master Key <b>evs/default</b>.</p> <ul style="list-style-type: none"> <li>- An existing or newly created CMK. For details, see "Key Management Service" &gt; "Creating a Key" in the <i>Data Encryption Workshop User Guide</i>.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>- System disk encryption relies on the image. For details, see the <i>Image Management Service User Guide</i>.</li> <li>- Before using the encryption function, you need to create an agency to grant KMS access rights to EVS. If you have the right to grant the permission, grant the KMS access rights to EVS directly. After the KMS access rights have been granted, follow-up operations do not require the rights to be granted again. If you do not have this permission, contact a user with the security administrator permissions to grant KMS access rights to EVS, then repeat the preceding operations.</li> <li>- The encryption attribute of a disk cannot be changed after the disk has been created.</li> </ul> <p>For details, see <a href="#">Managing Encrypted EVS Disks</a>.</p>	
More	Tag	<p>Optional</p> <p>You can add tags when creating disks. Tags can help you identify, classify, and search for your disks. For details about tag rules, see <a href="#">Adding a Tag</a>.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• Except for tagging the disk during disk creation, you can also add, modify, or delete tags for existing disks. For more information about tags, see <a href="#">Managing EVS Tags</a>.</li> </ul>	-

Parameter	Sub-Parameter	Description	Example Value
Disk Name	-	<p>Mandatory</p> <ul style="list-style-type: none"> <li>If you create a single disk, the name you entered will be used as the disk name. The name can contain a maximum of 64 characters.</li> <li>If you create multiple disks in a batch, the name you entered will be used as the prefix of disk names. An actual disk name will be composed of the name you entered and a four-digit number. The name can contain a maximum of 59 characters.</li> </ul>	<p>For example, if you create two disks and set <b>volume</b> for <b>Disk Name</b>, the EVS disk names will be <b>volume-0001</b> and <b>volume-0002</b>.</p>

Parameter	Sub-Parameter	Description	Example Value
Automatic Backup	-	<p>CBR lets you back up EVS disks and ECSs and use the backups to restore data. After you configure automatic backup, the system will associate the EVS disk with the backup vault and apply the selected policy to the vault to periodically back up the disk.</p> <ul style="list-style-type: none"> <li>• <b>Do not use:</b> Skip this configuration if backup is not required. If you need backup protection after a disk has been , log in to the CBR console, locate the desired vault, and associate the disk with the vault.</li> <li>• <b>Use existing:</b> <ol style="list-style-type: none"> <li>1. <b>Vault:</b> Select an existing vault from the drop-down list.</li> <li>2. <b>Backup Policy:</b> Select a backup policy from the drop-down list, or go to the CBR console and configure a desired one.</li> </ol> </li> <li>• <b>Auto assign:</b> <ol style="list-style-type: none"> <li>1. Enter a vault name, which can contain a maximum of 64 characters, including letters, digits, underscores (_), and hyphens (-), for example, <b>vault-f61e</b>. The default naming rule is <b>vault_XXXX</b>.</li> <li>2. Enter the vault capacity, which is required for backing up the disk. The vault capacity cannot be less than the size of the disk to be backed up. The value ranges from the disk size to 10,485,760 in the unit of GiB.</li> <li>3. Select a backup policy from the drop-down list, or go to the CBR console and configure a desired one.</li> </ol> </li> </ul>	-

Parameter	Sub-Parameter	Description	Example Value
Enterprise Project	-	<p>Mandatory</p> <p>When creating EVS disks, you can add the disks to an existing enterprise project.</p> <p>An enterprise project facilitates project-level management and grouping of cloud resources and users. The default project is <b>default</b>.</p>	default
Quantity	-	<p>Optional</p> <p>The preset disk quantity is <b>1</b>, which means only one disk is created. You can create a maximum of 100 disks at a time.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• If the disk is created from a backup, batch creation is not possible, and this parameter must be set to <b>1</b>.</li> <li>• If the disk is created from a snapshot, batch creation is not possible, and this parameter must be set to <b>1</b>.</li> </ul>	1

**Step 6** Click **Next**.

**Step 7** On the **Details** page, check the disk configuration.

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

**Step 8** In the disk list, view the disk status.

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

----End

## 4.3 Attaching an EVS Disk

### 4.3.1 Attaching a Non-Shared Disk

#### Scenarios

This section describes how to attach a non-shared EVS disk to a cloud server. Disks supporting this operation include:

- Separately created data disks
- Detached data disks

- Detached system disks

**NOTE**

After a system disk is detached from an ECS, the disk function changes to **Bootable disk**, and the status changes to **Available**. You can attach a bootable disk to an ECS to be used as a system disk or data disk depending on the device name selected.

### Prerequisites


- The status of the non-shared disk is **Available**.
- To attach a data disk, the status of the server must be **Running** or **Stopped**.
- To attach a system disk, the status of the server must be **Stopped**.

### Notes and 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 only be attached to one server.
- The disk and the server must be in the same region and AZ.
- A shared disk can be attached only when the servers' statuses are **Running** or **Stopped**.
- A frozen disk cannot be attached.

### Attaching the Disk on the EVS Console

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

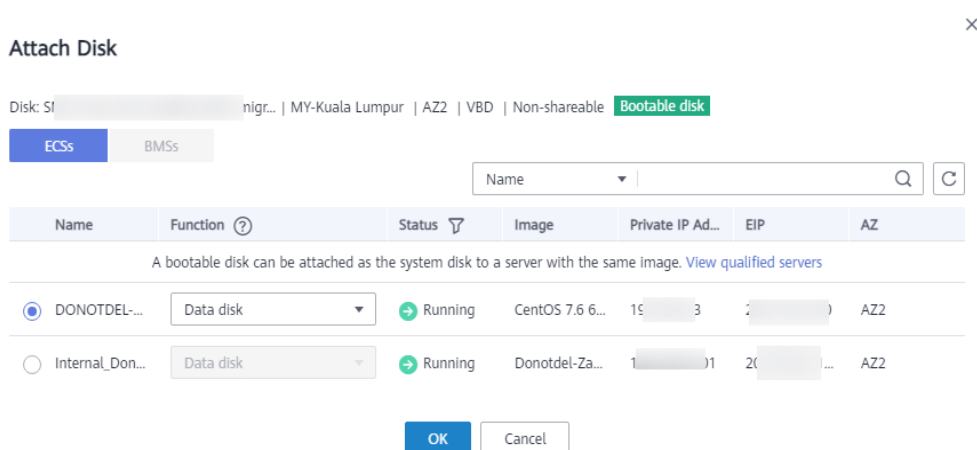
**Step 2** Click  in the upper left corner and choose **Storage > Elastic Volume Service**. The **Elastic Volume Service** page is displayed.

**Step 3** In the disk list, locate the disk 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 FAQ "How Do I Obtain My Disk Name in the ECS OS Using the Device Identifier Provided on the Console?" in the *Elastic Cloud Server FAQs*.

**Figure 4-5** Attach Disk



**Step 5** Click **OK**.

A dialog box is displayed, showing "The disk has been attached but still needs to be initialized before it can be used".

**NOTICE**



If you are attaching an EVS disk with data on it, initializing the disk will erase the existing data.

**Step 6** Click **OK** to go back to the disk list page.

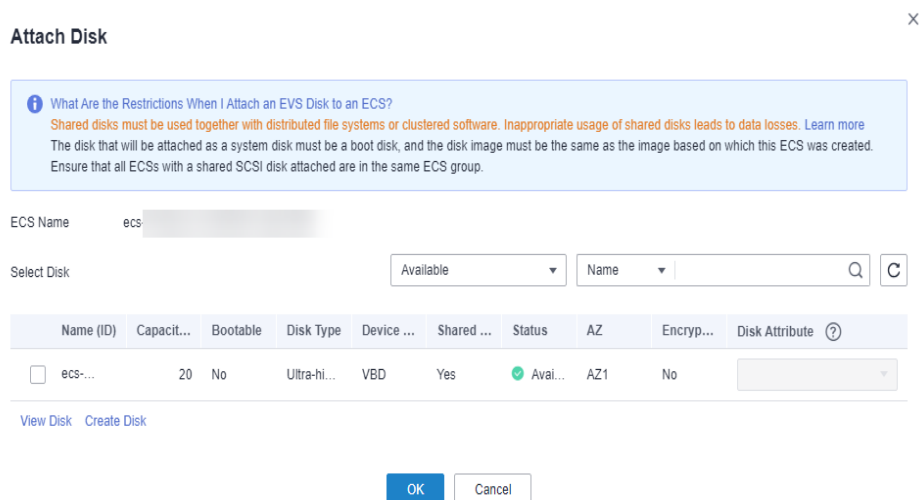
The status of the disk is **Attaching**, indicating that the disk is being attached to the server. When the disk status changes to **In-use**, the disk is successfully attached.

----End

## Attaching the Disk on the ECS Console

1. Log in to the console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  and choose **Compute > 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.
5. 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.

**Figure 4-6** Attaching an EVS disk



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`.

**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**. A dialog box is displayed, showing "The disk has been attached but still needs to be initialized before it can be used".
- After the disk is attached, you can view information about it on the **Disks** tab.

## Follow-Up Operations

- If you are attaching a new disk, you need then log in to the server and initialize the disk before it can be used. To learn how to initialize disks, see [Initializing EVS Data Disks](#).
- If you are attaching an EVS disk with data on it, you do not need to initialize it because initializing the disk will erase the existing data.

To mount a disk partition on a specific directory of the server, run the following command on the server:

```
mount Disk partition Mount point
```

## Helpful Links

If your disk cannot be attached to a server, see [Why Can't I Attach My Disk to a Server?](#)

If the attached data disk is not showing up, see [Why Can't I View the Attached Data Disk on the Server?](#)

## 4.3.2 Attaching a Shared Disk

### Scenarios

This section describes how to attach a shared EVS disk to a cloud server. Disks supporting this operation include:

- Separately created data disks
- Detached data disks

### Prerequisites

- The shared disk status is **In-use** or **Available**.
- The statuses of servers are **Running** or **Stopped**.

## Notes and 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, you need to 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 of a region.
- A shared, **In-use** disk can only be attached to servers when the maximum number of servers that the disk can be attached to has not been reached.
- A shared disk can only be attached to servers running the same type of OS (either Windows or Linux).  
For example, if you attach a shared disk to multiple Windows servers and then detach it, the shared disk cannot be attached to Linux servers later. This is because Windows and Linux support different file systems. 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.
- A frozen disk cannot be attached.

## Attaching the Disk on the EVS Console

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

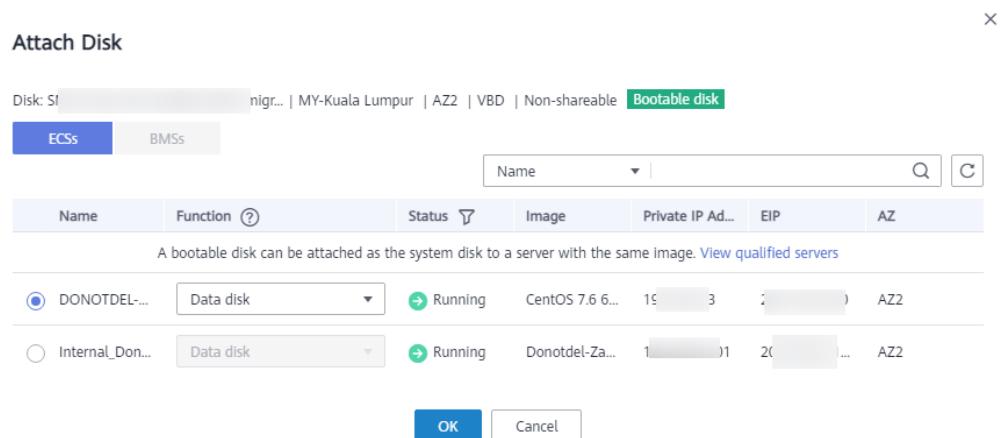
**Step 2** Click  in the upper left corner and choose **Storage > Elastic Volume Service**.

The **Elastic Volume Service** page is displayed.

**Step 3** In the disk list, locate the disk and click **Attach**.

Shared disks support batch attachment, so 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 to attach the shared disk. Ensure that the disk and servers are in the same AZ. After you select servers, the system automatically inputs **Data disk** as the disk function.

**Figure 4-7 Attach Disk****Step 5** Click **OK**.


A dialog box is displayed, showing "The disk has been attached but still needs to be initialized before it can be used".

**Step 6** Click **OK** to go back 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 console.
2. Click  in the upper left corner and select the desired region and project.
3. Choose **Compute > 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.
5. Click the name of the target ECS.  
The ECS details page 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 the system disk or a data disk but cannot specify a specific 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**. A dialog box is displayed, showing "The disk has been attached but still needs to be initialized before it can be used".

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

## Follow-Up Operations

- If you are attaching a new disk, you need then log in to the server and initialize the disk before it can be used. To learn how to initialize disks, see [Initializing EVS Data Disks](#).
- If you are attaching an EVS disk with data on it, you do not need to initialize it because initializing the disk will erase the existing data.

To mount a disk partition on a specific directory of the server, run the following command on the server:

```
mount Disk partition Mount point
```

## Helpful Links

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

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

# 4.4 Initializing EVS Data Disks

## 4.4.1 Initialization Overview

After you attach a new data disk to a server, you must initialize the disk including creating partitions, creating file systems, and mounting the partitions before you can use the disk.

### Scenarios

- **System disk**  
When a server is created, a system disk is automatically initialized with master boot record (MBR).
- **New data disk**
  - If a data disk is created together with a server, EVS automatically attaches it to the server. You only need to initialize it to make it available for use.
  - If a data disk is created explicitly, you need to first attach it to a server and then initialize it.

For detailed operation instructions, see [Table 4-4](#).

- **Existing data disk**  
An existing data disk is a disk created from a snapshot, a backup, or an image, or a disk detached from another server.
  - You can choose not to initialize the disk and use the disk existing partitions.

- In Linux, [mount the partitions on desired mount points](#) and [configure auto mount at system start](#).
  - In Windows, no further action is required. You can simply use the existing partitions.
- You can also re-initialize the data disk.

Re-partitioning a disk will erase all the existing data on the disk, so you are advised to use snapshots to back up the disk data first.

- In Linux, unmount the partitions, delete them (by running `fdisk Disk name`, entering `d` and the partition number, and entering `w`), and then re-initialize the disk.
- In Windows, delete the partitions (using the volume deletion tool) and then re-initialize the disk.

For detailed initialization operations, see [Table 4-4](#).

 NOTE

Initializing a disk does not delete the snapshots created for the disk, so you can still use snapshots to roll back data to the source disk after the disk is initialized.

## Operation Instructions

**Table 4-4** Disk initialization instructions

Disk Capacity	Partiti on Style	Partition Type	OS	Reference
Capacity ≤ 2 TiB	GPT or MBR	<ul style="list-style-type: none"> <li>• GPT partitions are not classified, and there is no limit on the number of GPT partitions.</li> <li>• MBR partitions can be:                             <ul style="list-style-type: none"> <li>– Four primary partitions</li> <li>– Three primary partitions and one extended partition</li> </ul>                             The number of logical partitions allowed in the extended partition is not limited, so theoretically you can create as many logical partitions as you want.                             <p>If you need five or more partitions, use the "primary partitions + one extended partition" model and then create logical partitions in the extended partition.</p> </li> </ul>	Linux	<a href="#">Initializing a Linux Data Disk (Less Than or Equal to 2 TiB)</a>
			Windows	<a href="#">Initializing a Windows Data Disk</a>

Disk Capacity	Partition Style	Partition Type	OS	Reference
Capacity > 2 TiB	GPT	GPT partitions are not classified, and there is no limit on the number of GPT partitions.	Linux	<a href="#">Initializing a Linux Data Disk (Greater Than 2 TiB)</a>
			Windows	<a href="#">Initializing a Windows Data Disk</a>

**NOTICE**

- The maximum disk size that MBR supports is 2 TiB, and that GPT supports is 18 EiB. If your disk is greater than 2 TiB or you may expand it to over 2 TiB later, use GPT when initializing disks.
- If you change the partition style of a disk, data on the disk will be erased. Select an appropriate partition style when initializing disks.
- In Linux, you can use either `fdisk` or `parted` to create MBR partitions, and use only `parted` to create GPT partitions.

## 4.4.2 Initializing a Linux Data Disk (Less Than or Equal to 2 TiB)

### Scenarios

This section describes how to initialize a Linux data disk manually or using a script. The operations may vary depending on the server OS. Perform initialization operations based on your server OS.

### Prerequisites

You have attached the disk to a server.

### Notes and Constraints

- A disk created from a data source does not need to be initialized. Such a disk contains the source data in the beginning. Initializing the disk may clear the initial data on it. If you need to re-initialize the disk, you are advised to back up the disk data first. To back up data using CBR, see [Backing Up EVS Disks](#). To back up data using snapshots, see [Managing EVS Snapshots](#).
- Initializing a disk does not delete the snapshots created for the disk, so you can still use snapshots to roll back data to the source disk after the disk is initialized.

## Initializing a Data Disk

### NOTE

MBR supports a maximum of four primary partitions or a maximum of three primary partitions plus one extended partition. Multiple logical partitions can be created in the extended partition.

For example, if you want to create four partitions, you have the following options:

- Create four primary partitions.
- Create one primary partition and one extended partition (three logical partitions).
- Create two primary partitions and one extended partition (two logical partitions).
- Create three primary partitions and one extended partition (one logical partition).

**The following example shows you how to use `fdisk` to create two primary MBR partitions (`/dev/vdb1`: 40 GiB; `/dev/vdb2`: 60 GiB) on the `/dev/vdb` data disk.**

### Step 1 Log 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*.

### Step 2 Create two primary partitions, `/dev/vdb1` and `/dev/vdb2` for data disk `/dev/vdb`.

1. Check that the capacity of the `/dev/vdb` data disk is 100 GiB.

#### `lsblk`

```
[root@ecs-centos76 ~]# 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 100G 0 disk
```

2. Create the first primary partition `/dev/vdb1`.

#### `fdisk /dev/vdb`

`n`

`p`

`1`

### NOTE

- Entering `p` for **Partition type** creates a primary partition, and entering `e` creates an extended partition.
- Value `1` is the primary partition number.

```
[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): n
Partition type:
  p primary (0 primary, 0 extended, 4 free)
  e extended
Select (default p): p
Partition number (1-4, default 1): 1
```

Set **First sector** to **2048** and **Last sector** to **83886079** for partition **/dev/vdb1** (40 GiB).

```
First sector (2048-209715199, default 2048): 2048
Last sector, +sectors or +size{K,M,G} (2048-209715199, default 209715199):83886079
Partition 1 of type Linux and of size 40 GB is set
```

### 3. Create the second primary partition **/dev/vdb2**.

**n**

**p**

**2**

```
Command (m for help): n
Partition type:
  p primary (0 primary, 0 extended, 4 free)
  e extended
Select (default p): p
Partition number (1-4, default 2): 2
```

Set the **First sector** to **83886080** and **Last sector** to **209715199** for partition **/dev/vdb2**.

```
First sector (83886080-209715199, default 83886080): 83886080
Last sector, +sectors or +size{K,M,G} (83886080-209715199, default 209715199):209715199
Partition 2 of type Linux and of size 60 GB is set
```

#### NOTE

First and last sectors of the partitions in this example are calculated as follows:

**Sector value = Capacity/512 bytes, 1 GiB = 1073741824 bytes**

- **First sector (2048-209715199, default 2048)** shows the sector value range of the **/dev/vdb** data disk (100 GiB).

First sector = 2048

Last sector = Sector value - 1 = (100 x 1073741824/512) - 1 = 209715200 - 1 = 209715199

- For the first partition **/dev/vdb1** (40 GiB) of the **/dev/vdb** data disk:

First sector = 2048 (The start sector of the **/dev/vdb** data disk is used.)

Last sector = Sector value - 1 = (40 x 1073741824/512) - 1 = 83886079

- For the second partition **/dev/vdb2** (60 GiB) of the **/dev/vdb** data disk:

First sector = Last sector of **/dev/vdb1** + 1 = 83886079 + 1 = 83886080

Last sector = First sector + Sector value - 1 = 83886080 + (60 x 1073741824/512) - 1 = 209715199

### Step 3 Check the sizes and partition styles of the new partitions.

#### 1. Check whether the partitioning is successful.

**p**

```
Command (m for help): p
Disk /dev/vdb: 107.4 GB, 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: 0x994727e5
```

Device	Boot	Start	End	Blocks	Id	System
/dev/vdb1		2048	83886079	41942016	83	Linux
/dev/vdb2		83886080	209715199	62914560	83	Linux

Command (m for help):



 NOTE

In case that you want to discard the changes made before, you can exit fdisk by entering **q** and press **Enter**. Then, re-create the partitions by referring to step 1.

2. Write the changes to the partition table and synchronize the new partition table to the OS.

**w**

**partprobe**

 NOTE

If error message **-bash: partprobe: command not found** is returned, the system cannot identify the command. In this case, run **yum install -y parted** to install the command. Then run the command again.

3. Confirm that the partition style is MBR.

**parted /dev/vdb**

**p**

 NOTE

If **Partition Table: msdos** is returned, the partition style is MBR.

```
[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
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 107GB
Sector size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:

Number Start End     Size  Type  File system  Flags
 1    1049kB 42.9GB 42.9GB primary
 2    42.9GB 107GB  64.4GB primary

(parted) q
[root@ecs-test-0001 ~]#
```

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

- Step 4** Create ext4 file systems for partitions **/dev/vdb1** (40 GiB) and **/dev/vdb2** (60 GiB).

**mkfs -t ext4 /dev/vdb1**

**mkfs -t ext4 /dev/vdb2**

 NOTE

It takes some time to create file systems. Do not exit before the system returns the following information:

```
[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
2621440 inodes, 10485504 blocks
524275 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
```

Check whether the file system format is ext4.

### **parted /dev/vdb**

**p**

```
[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
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 107GB
Sector size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:

Number  Start  End    Size  Type  File system  Flags
  1      1049kB 42.9GB 42.9GB primary ext4
  2      42.9GB 107GB  64.4GB primary ext4

(parted) q
[root@ecs-test-0001 ~]#
```

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

- Step 5** Create directories (mount points) and mount the new partitions on the created mount points.

```
mkdir -p /mnt/sdc
```

```
mkdir -p /mnt/sdd
```

```
mount /dev/vdb1 /mnt/sdc
```

```
mount /dev/vdb2 /mnt/sdd
```

```
lsblk
```

View the mount results.

```
[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 40G 0 part /mnt/sdc
└─vdb2 253:18 0 60G 0 part /mnt/sdd
```

You should now see that partitions **/dev/vdb1** and **/dev/vdb2** are mounted on **/mnt/sdc** and **/mnt/sdd**.

- Step 6** Use the partition UUIDs to configure auto mount at startup.

 NOTE

- Mounts become invalid after a system reboot. You can configure auto mount at startup by adding information of the new partition into the `/etc/fstab` file.
- You are advised not to use device names to identify disks in the `/etc/fstab` file because device names are assigned dynamically and may change (for example, from `/dev/vdb1` to `/dev/vdb2`) after a stop or start. This can even prevent your server from booting up.
- UUIDs are the unique character strings for identifying partitions in Linux.
- This operation will not affect the existing data on the ECS.

1. Query the partition UUIDs.

**blkid /dev/vdb1**

**blkid /dev/vdb2**

```
[root@ecs-test-0001 ~]# blkid /dev/vdb1
/dev/vdb1: UUID="0b3040e2-1367-4abb-841d-ddb0b92693df" TYPE="ext4"
/dev/vdb2: UUID="0d6769k2-1745-9dsf-453d-hgd0b34267dj" TYPE="ext4"
```

The UUIDs of partitions `/dev/vdb1` and `/dev/vdb2` are **0b3040e2-1367-4abb-841d-ddb0b92693df** and **0d6769k2-1745-9dsf-453d-hgd0b34267dj**.

2. Configure auto mount at startup.

**vi /etc/fstab**

Press **i** to enter the editing mode, move the cursor to the end of the file, press **Enter**, and add the following content:

```
UUID=0b3040e2-1367-4abb-841d-ddb0b92693df /mnt/sdc ext4 defaults 0 2
UUID=0d6769k2-1745-9dsf-453d-hgd0b34267dj /mnt/sdd ext4 defaults 0 2
```

Press **Esc**, enter **:wq**, and press **Enter** to save the settings and exit the vi editor.

**Table 4-5** Parameter description

Example Value	Description
UUID=0b3040e2-1367-4abb-841d-ddb0b92693df	The UUID of the partition.
/mnt/sdc	The mount point of the partition.
ext4	The file system format of the partition.
defaults	The partition mount option. Normally, this parameter is set to <b>defaults</b> .
0	<ul style="list-style-type: none"> <li>– The Linux dump backup option. <ul style="list-style-type: none"> <li>▪ <b>0</b>: Linux dump backup is not used. Usually, dump backup is not used, and you can set this parameter to <b>0</b>.</li> <li>▪ <b>1</b>: Linux dump backup is used.</li> </ul> </li> </ul>

Example Value	Description
2	<ul style="list-style-type: none"><li>- The fsck option, which means whether to use fsck to check the disk during startup.<ul style="list-style-type: none"><li>▪ <b>2</b>: The check starts from the partitions whose mount points are non-root directories. / is the root directory.</li><li>▪ <b>1</b>: The check starts from the partitions whose mount points are root directories.</li><li>▪ <b>0</b>: The fsck option is not used.</li></ul></li></ul>

**Step 7** Verify that auto mount takes effect.

```
umount /dev/vdb1
```

```
umount /dev/vdb2
```

```
mount -a
```

The system reloads all the content in the `/etc/fstab` file.

Query file system mounting information.

```
mount | grep /mnt/sdc
```

```
mount | grep /mnt/sdd
```

If information similar to the following is displayed, auto mount has taken effect:

```
root@ecs-test-0001 ~]# mount | grep /mnt/sdc
/dev/vdb1 on /mnt/sdc type ext4 (rw,relatime,data=ordered)
root@ecs-test-0001 ~]# mount | grep /mnt/sdd
/dev/vdb2 on /mnt/sdd type ext4 (rw,relatime,data=ordered)
```

----End

### 4.4.3 Initializing a Linux Data Disk (Greater Than 2 TiB)

#### Scenarios

When the size of a disk is greater than 2 TiB, you can only use parted to create GPT partitions. The initialization operations may vary depending on the server OS.

Partition Style	OS Requirements	File System Format	Partitioning Tool	Example Configuration
GPT	None	ext* (such as ext2, ext3, and ext4), xfs, and btrfs	parted	<ul style="list-style-type: none"> <li>• Device name: /dev/vdb</li> <li>• File system format: ext4</li> <li>• Mount point: /mnt/sdc</li> <li>• Partition name: /dev/vdb1</li> <li>• Partition style: GPT</li> <li>• Size: 3 TiB</li> </ul>

## Prerequisites

You have attached the disk to a server.

## Notes and Constraints

- A disk created from a data source does not need to be initialized. Such a disk contains the source data in the beginning. Initializing the disk may clear the initial data on it. If you need to re-initialize the disk, you are advised to back up the disk data first. To back up data using CBR, see [Backing Up EVS Disks](#). To back up data using snapshots, see [Managing EVS Snapshots](#).
- Initializing a disk does not delete the snapshots created for the disk, so you can still use snapshots to roll back data to the source disk after the disk is initialized.

## Initializing a Data Disk Greater Than 2 TiB

The following example shows you how to use parted to create a GPT partition on the /dev/vdb data disk.

**Step 1** Log 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*.

**Step 2** Create the /dev/vdb1 partition on data disk /dev/vdb.

1. Check that the capacity of the /dev/vdb data disk is 3 TiB.

### lsblk

```
[root@ecs-centos76 ~]# 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 3T 0 disk
```

2. Create the /dev/vdb1 partition.

**parted /dev/vdb**

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

**Partition Table: unknown** means that no partition style is set for the new disk.

#### NOTE

If error message **-bash: parted: command not found** is returned, the system cannot identify the command. In this case, run **yum install -y parted** to install the command. Then run the command again.

3. Set the partition style of the **/dev/vdb1** partition to GPT.

**mklabel gpt**

**unit s**

```
p
(parted) mklabel gpt
(parted) unit s
(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
(parted)
```

#### NOTE

- If your disk size is smaller than 2 TiB and you want to use parted to create an MBR partition, run **mklabel msdos**.
- If you change the partition style of a disk, data on the disk will be erased. Select an appropriate partition style when initializing disks.
- **The partition style (MBR or GPT) set here will apply to all subsequent partitions created on this EVS disk. When you create partitions on this disk later, you do not need to perform this step again.**

4. Set the name and size of the **/dev/vdb1** partition.

**mkpart /dev/vdb1 2048s 100%**

**p**

#### NOTE

- Partition **/dev/vdb1** is created, starting on **2048** and using 100% of the rest of the disk.
- If you want to create two or more partitions, calculate the first and last sectors of the partitions based on the method provided in [Step 2](#).

```
(parted) mkpart /dev/vdb1 2048s 100%
(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                /dev/vdb1
```

Enter **q** and press **Enter**. Then run **lsblk** to view the new partition **/dev/vdb1**.

```
[root@ecs-centos74 ~]# 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    3T  0 disk
├─vdb1 253:17 0    3T  0 part
```

**Step 3** Create an ext4 file system on the **/dev/vdb1** partition.

**mkfs -t ext4 /dev/vdb1**

 **NOTE**

It takes some time to create a file system. Observe the system running status and do not exit.

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

Run **parted /dev/vdb** and enter **p** to check the file system format.

```
[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
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 3299GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number Start  End    Size  File system  Name        Flags
 1     1049kB  3299GB 3299GB  ext4        /dev/vdb1

(parted) q
[root@ecs-test-0001 ~]#
```

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

**Step 4** Create a directory (mount point) and mount the new partition on the created mount point.

```
mkdir -p /mnt/sdc
```

```
mount /dev/vdb1 /mnt/sdc
```

```
lsblk
```

```
[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 3T 0 disk
|-vdb1 253:17 0 3T 0 part /mnt/sdc
```

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

**Step 5** Use the partition UUID to configure auto mount at startup.

 **NOTE**

- Mounts become invalid after a system reboot. You can configure auto mount at startup by adding information of the new partition into the **/etc/fstab** file.
- You are advised not to use device names to identify disks in the **/etc/fstab** file because device names are assigned dynamically and may change (for example, from **/dev/vdb1** to **/dev/vdb2**) after a stop or start. This can even prevent your server from booting up.
- UUIDs are the unique character strings for identifying partitions in Linux.
- This operation will not affect the existing data on the ECS.

1. Query the partition UUID.

```
blkid /dev/vdb1
```

```
[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 **0b3040e2-1367-4abb-841d-ddb0b92693df**.

2. Configure auto mount at startup.

```
vi /etc/fstab
```

Press **i** to enter the editing mode, move the cursor to the end of the file, press **Enter**, and add the following content:

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

Press **Esc**, enter **:wq**, and press **Enter** to save the settings and exit the vi editor.

**Table 4-6** Parameter description

Example Value	Description
UUID=0b3040e2-1367-4abb-841d-ddb0b92693df	The UUID of the partition.
/mnt/sdc	The mount point of the partition.
ext4	The file system format of the partition.



Example Value	Description
defaults	The partition mount option. Normally, this parameter is set to <b>defaults</b> .
0	<ul style="list-style-type: none"> <li>- The Linux dump backup option. <ul style="list-style-type: none"> <li>▪ <b>0</b>: Linux dump backup is not used. Usually, dump backup is not used, and you can set this parameter to <b>0</b>.</li> <li>▪ <b>1</b>: Linux dump backup is used.</li> </ul> </li> </ul>
2	<ul style="list-style-type: none"> <li>- The fsck option, which means whether to use fsck to check the disk during startup. <ul style="list-style-type: none"> <li>▪ <b>2</b>: The check starts from the partitions whose mount points are non-root directories. / is the root directory.</li> <li>▪ <b>1</b>: The check starts from the partitions whose mount points are root directories.</li> <li>▪ <b>0</b>: The fsck option is not used.</li> </ul> </li> </ul>

**Step 6** Verify that auto mount takes effect.

**umount /dev/vdb1**

**mount -a**

The system reloads all the content in the **/etc/fstab** file.

Query file system mounting information.

**mount | grep /mnt/sdc**

If information similar to the following is displayed, auto mount has taken effect:

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

----End

## 4.4.4 Initializing a Windows Data Disk

### Scenarios

This section uses the example configurations below to describe how to use Disk Management Tool to initialize a Windows data disk. The initialization operations

may vary depending on the server OS. Perform initialization operations based on your server OS.

## Prerequisites

You have attached the disk to a server.

## Notes and Constraints

- A disk created from a data source does not need to be initialized. Such a disk contains the source data in the beginning. Initializing the disk may clear the initial data on it. If you need to re-initialize the disk, you are advised to back up the disk data first. To back up data using CBR, see [Backing Up EVS Disks](#). To back up data using snapshots, see [Managing EVS Snapshots](#).
- Initializing a disk does not delete the snapshots created for the disk, so you can still use snapshots to roll back data to the source disk after the disk is initialized.

## Initializing a Data Disk Manually

The following example shows you how to create a 100 GiB GPT partition with an NTFS file system on a server running Windows Server 2019.

**Step 1** Log 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*.


**Step 2** On the desktop of the server, click the start icon in the lower left corner.

The **Windows Server** window is displayed.

**Step 3** Click **Server Manager**.


The **Server Manager** window is displayed.

Figure 4-8 Server Manager



Step 4 In the upper right corner, choose **Tools > Computer Management**.

Figure 4-9 Computer Management



Step 5 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.

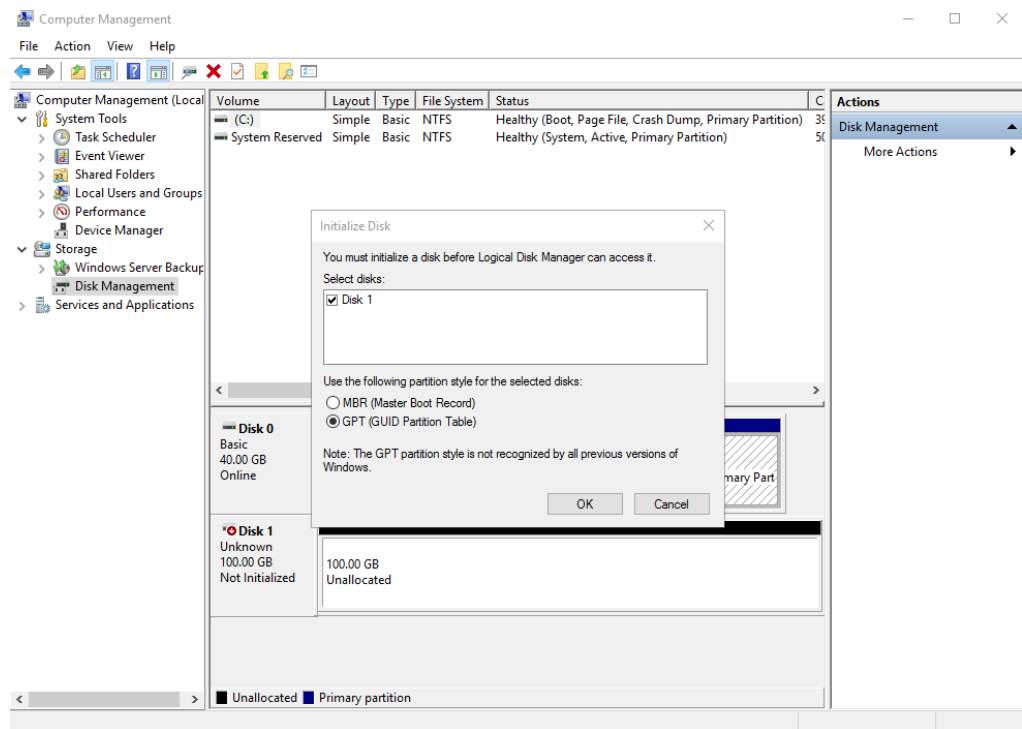
In the **Initialize Disk** dialog box, the to-be-initialized disk is selected. Select a partition style and click **OK**. In this example, **GPT (GUID Partition Table)** is selected.

**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 greater than 2 TiB.

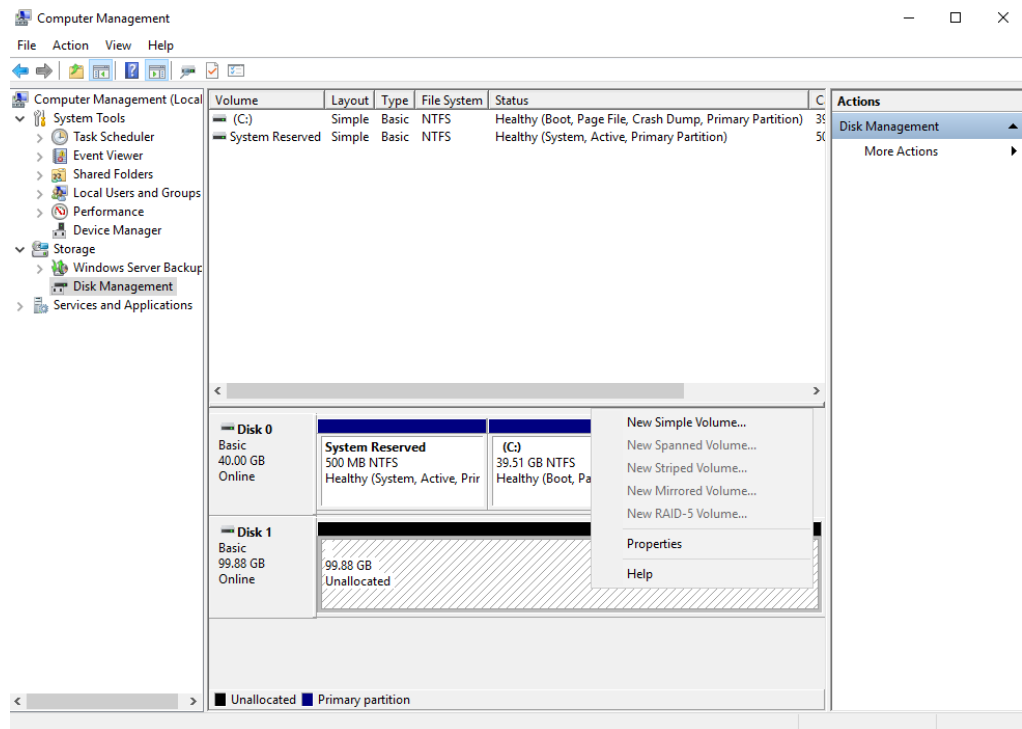
If the partition style of an in-use disk is changed, all data on the disk will be lost, so take care to select an appropriate partition style when initializing the disk. If you must change the partition style to GPT, it is recommended that you back up the disk data before the change.

**Figure 4-10** Disk list



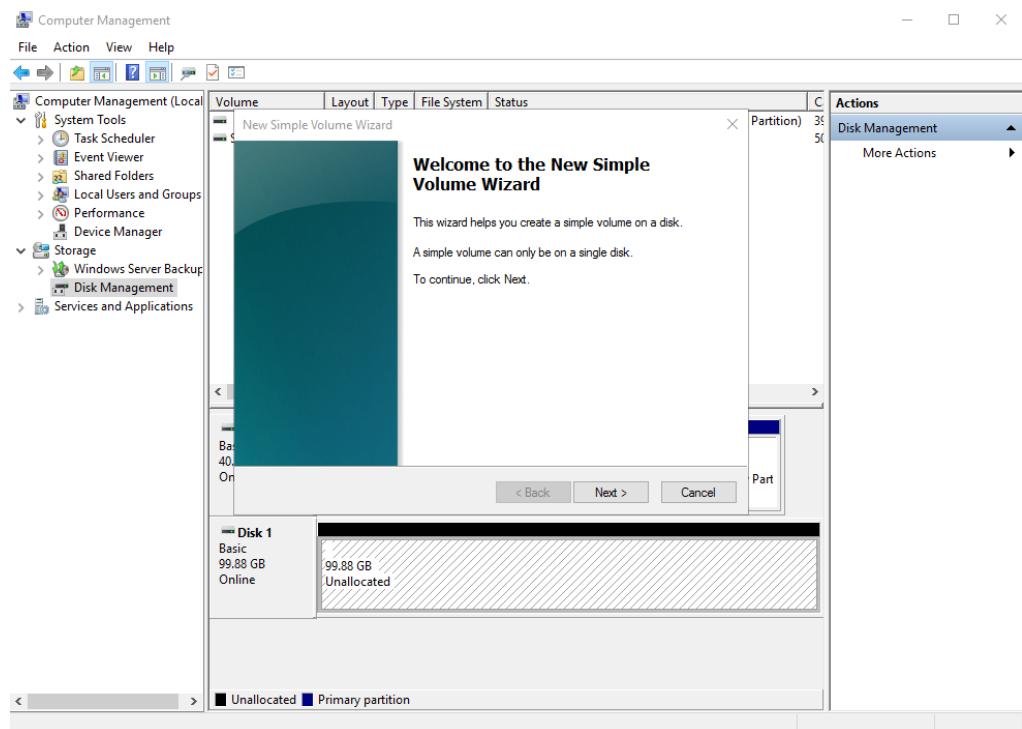
**Step 6** In the **Unallocated** area of **Disk 1**, right-click the blank area and choose **New Simple Volume**.

**Figure 4-11 Computer Management**




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

**Figure 4-12 New Simple Volume Wizard**




**Step 7** Click **Next** to go to the **Specify Volume Size** page.

Figure 4-13 Specify Volume Size



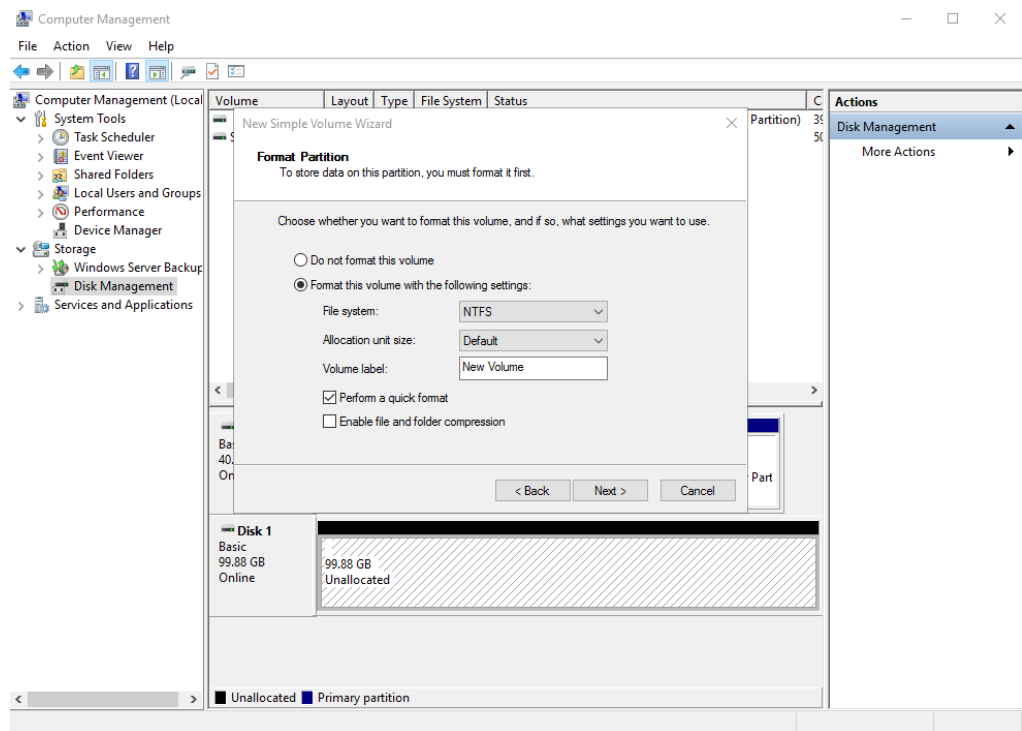
**Step 8** Specify the volume size and click **Next**. The system selects the maximum volume size by default. You can specify the volume size as required. In this example, the default setting is used.

Figure 4-14 Assign Drive 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.

Figure 4-15 Format Partition



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

Figure 4-16 Completing the New Simple Volume Wizard



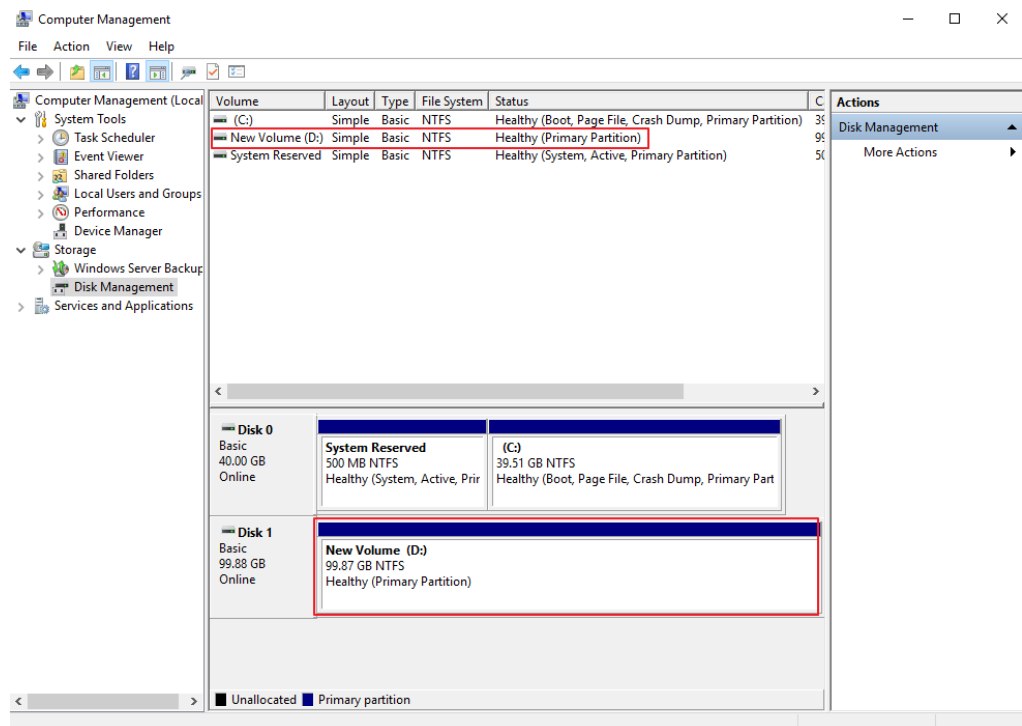
**NOTICE**


The partition sizes supported by file systems vary. Choose an appropriate file system format 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 succeeded.

**Figure 4-17** Disk initialized

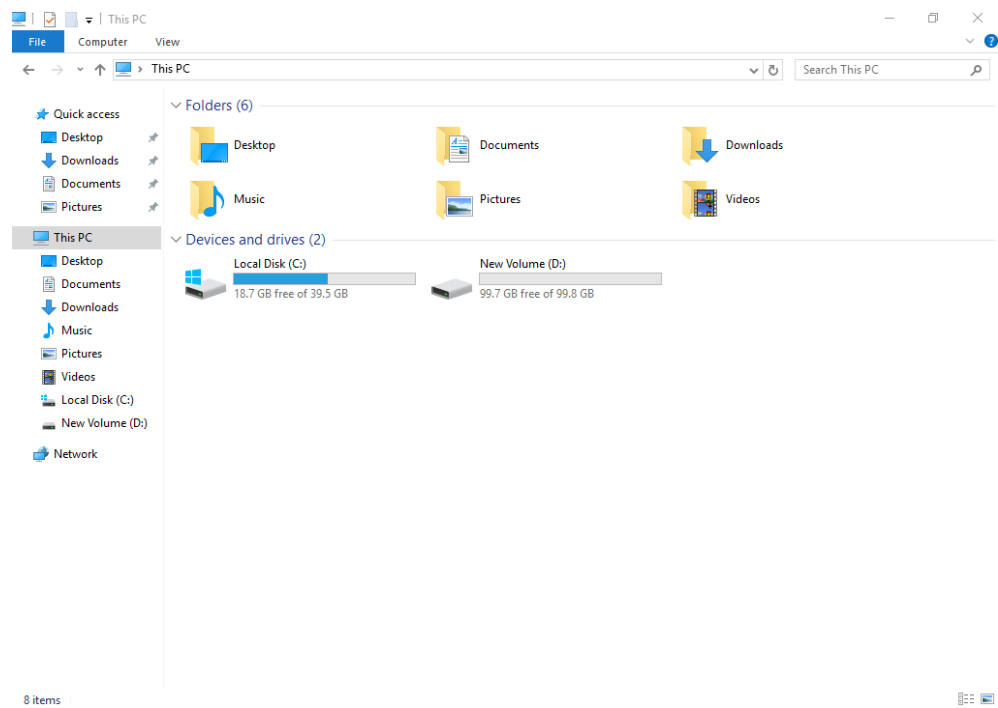


**Step 12** After the volume is created, click  on the task bar and check whether a new volume appears in the File Explorer. In this example, New Volume (D:) is the new volume.

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



**Figure 4-18** File Explorer



-----End

# 5 Viewing EVS Disk Details

## Scenarios


This section describes how to view disk details, including the disk status and specifications. Two methods are as follows:

- [Viewing Disk Details from the EVS Console](#)
- [Viewing Disk Details from the Cloud Server Console](#)

See [EVS Disk Status](#) to learn more about disk statuses.

## Viewing Disk Details from the EVS Console

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

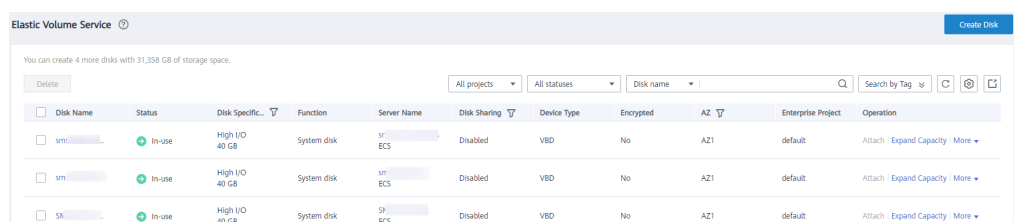
**Step 2** Click  in the upper left corner and choose **Storage > Elastic Volume Service**.

The **Elastic Volume Service** page is displayed.

**Step 3** In the disk list, view disk information including the disk status, type, size, function, and device type.

In the upper left corner of the list, you can filter disks by project, status, disk name, tag, or other properties.

**Figure 5-1** Viewing the disk list



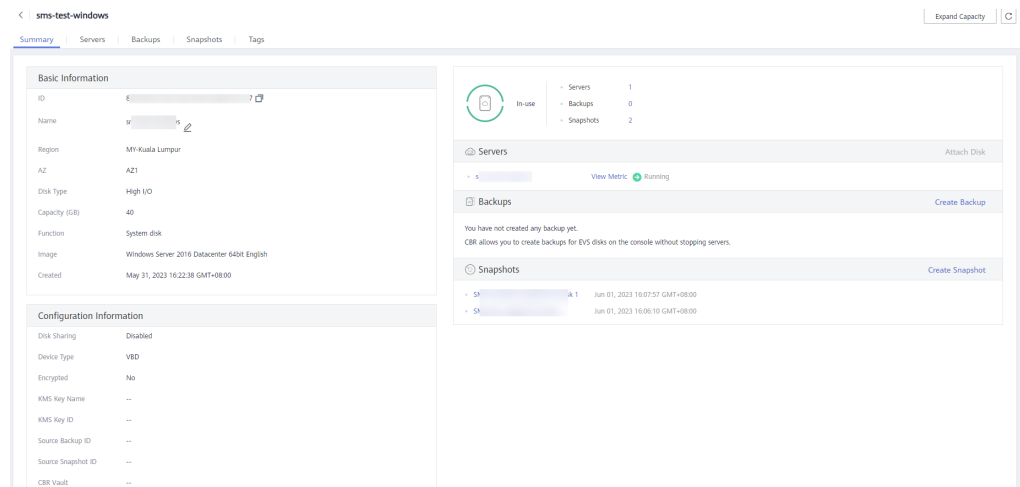
The screenshot shows the Elastic Volume Service console interface. At the top, there is a header with the service name and a 'Create Disk' button. Below the header, there is a message: 'You can create 4 more disks with 31,358 GB of storage space.' The main area contains a table with columns for Disk Name, Status, Disk Specifics, Function, Server Name, Disk Sharing, Device Type, Encrypted, AZ, Enterprise Project, and Operation. Three disks are listed in the table, all with a status of 'In-use' and a size of 40 GB.

Disk Name	Status	Disk Specifics	Function	Server Name	Disk Sharing	Device Type	Encrypted	AZ	Enterprise Project	Operation
sm-...	In-use	High I/O 40 GB	System disk	si-...	Disabled	VBD	No	AZ1	default	Attach   Expand Capacity   More
sm-...	In-use	High I/O 40 GB	System disk	sm-...	Disabled	VBD	No	AZ1	default	Attach   Expand Capacity   More
SA-...	In-use	High I/O 40 GB	System disk	si-...	Disabled	VBD	No	AZ1	default	Attach   Expand Capacity   More

**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 5-2** Disk details page



**Step 5** (Optional) Export disk information.

Click the export button in the upper right corner of the list to export disk information.

----End

## Viewing Disk Details from the Cloud Server Console


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

**Step 2** Choose **Compute > Elastic Cloud Server**.

The **Elastic Cloud Server** page is displayed.

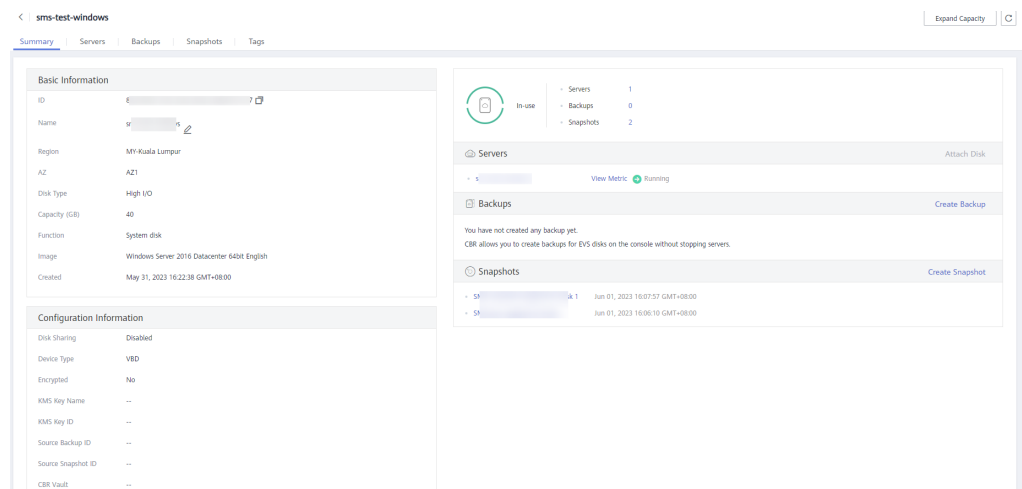
**Step 3** In the server list, locate the desired server by server name and click the name.

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 5-3** Disk details page



----End

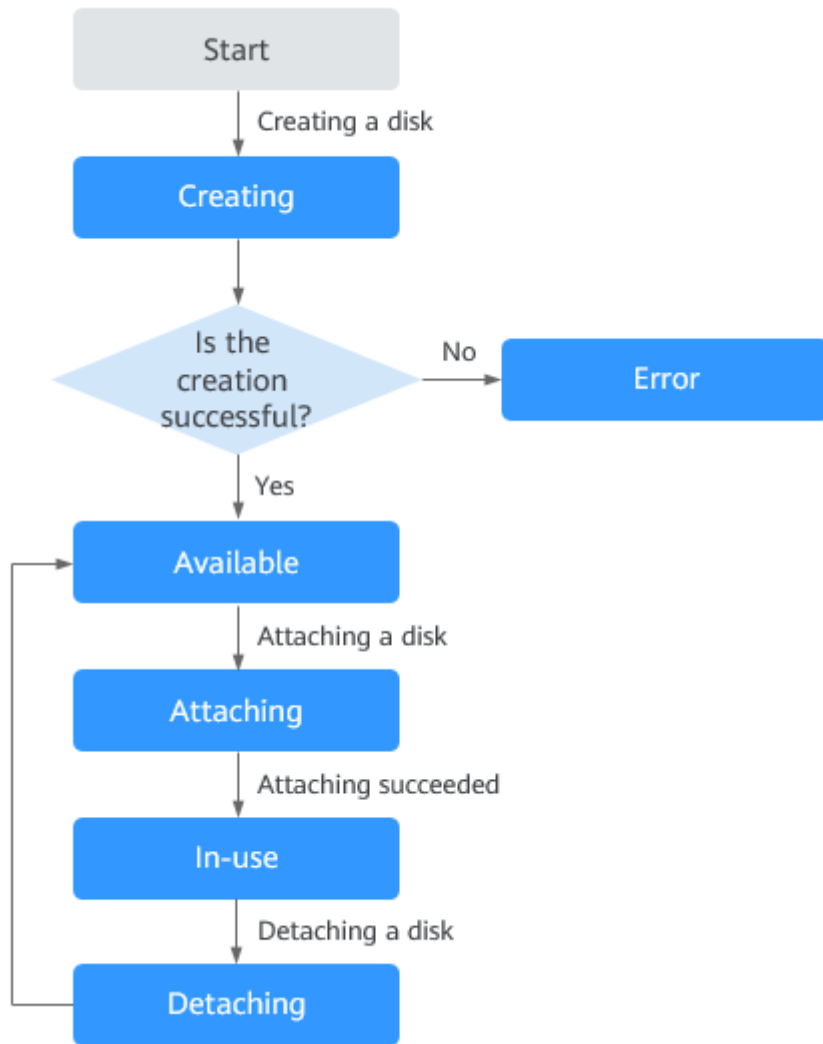
## EVS Disk Status

**Table 5-1** EVS disk status details

Status	Description	Allowed Operation
In-use	The EVS disk has been attached to a server and is in use.	<ul style="list-style-type: none"> <li>Detaching</li> <li>Expanding capacity</li> </ul>
Available	The EVS disk has not been attached to any server, so you can attach it.	<ul style="list-style-type: none"> <li>Attaching</li> <li>Expanding capacity</li> <li>Deleting</li> <li>Rolling back data to EVS disks using snapshots</li> </ul>
Creating	The EVS disk is being created.	-
Attaching	The EVS disk is being attached to a server.	-
Detaching	The EVS disk is being detached from a server.	-
Deleting	The EVS disk is being deleted.	-
Expanding	The capacity of the EVS disk is being expanded.	-
Uploading	Data on the EVS disk is being uploaded to an image. This status occurs when you create an image from a server.	-

Status	Description	Allowed Operation
Downloading	Data is being downloaded from an image to the EVS disk. This status occurs when you create a server.	-
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	<p>Data on the EVS disk is being restored from a snapshot.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• When you roll back data from a snapshot, you can only roll back data to the source EVS disk. Rollback to a specific disk is not supported.</li> <li>• A snapshot can only be used for rollback when its source disk is in the <b>Available</b> or <b>Rollback failed</b> state.</li> </ul>	-
Rollback failed	An error occurs when the EVS disk data is rolled back from a snapshot.	<ul style="list-style-type: none"> <li>• Deleting</li> <li>• Rolling back data to EVS disks using snapshots</li> </ul>

**Figure 5-4** Change between some of EVS disk statuses



**NOTE**

If an EVS disk status is **Error**, **Deletion failed**, **Expansion failed**, **Rollback failed**, or **Restoration failed**, you can rectify the error by referring to [What Should I Do If an Error Occurs on My EVS Disk?](#)

# 6 Expanding the EVS Disk Capacity

---

## 6.1 Expansion Overview

### What Is EVS Capacity Expansion?

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

### Upper Limits on Disk Capacity

The disk capacity can only be expanded, not reduced. The maximum disk capacity is as follows:

- System disk: 1 TiB
- Data disk: 32 TiB

 **NOTE**

If you detach a system disk and then attach it to another server as a data disk, the maximum capacity of this disk is still 1 TiB.

### How Do I Expand the Disk Capacity?

You can expand the capacity of an EVS disk in two steps:

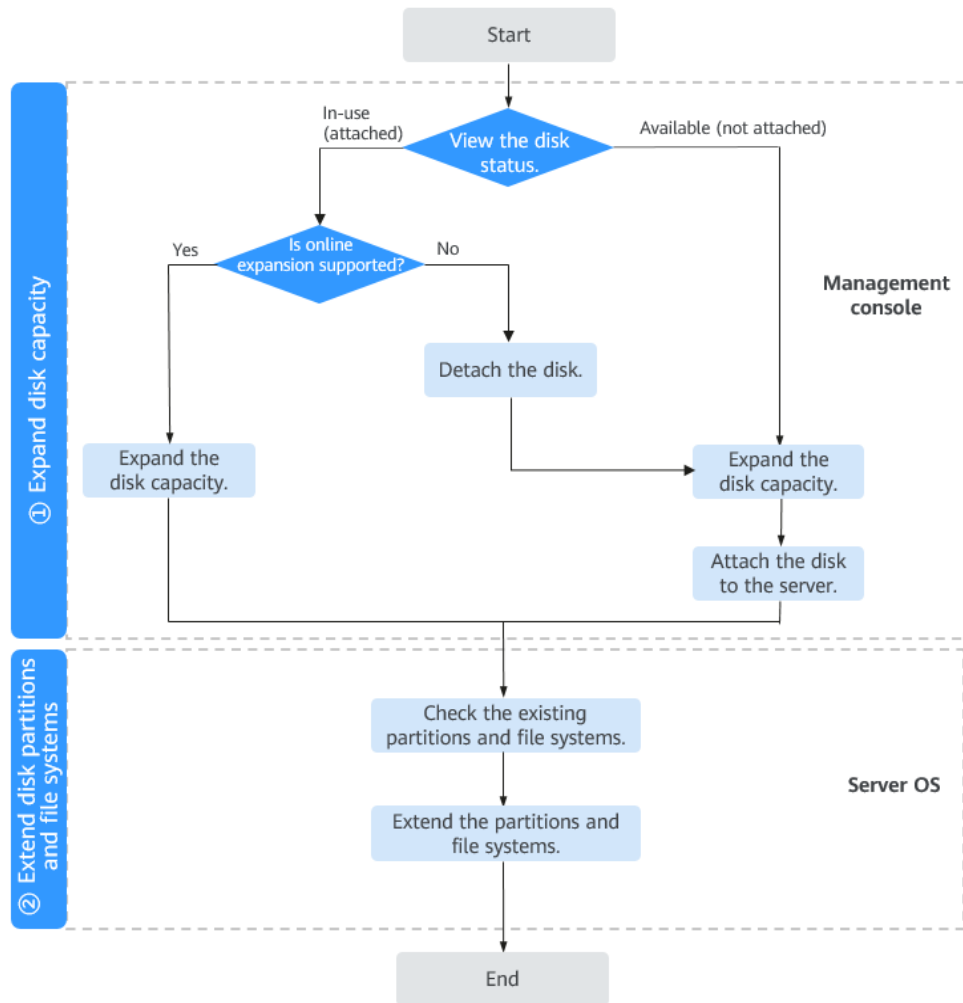
**Step 1: Expand Disk Capacity**

**Step 2: Extend Disk Partitions and File Systems**

 **NOTE**

If the server is stopped during the expansion, the additional space of a Windows system disk, Windows data disk, or Linux system disk may be automatically added to the last partition after the server is started. In this case, the additional space can be directly used. If the additional space is not automatically added, you need to extend the partition and file system according to the step 2 above.

Figure 6-1 Capacity expansion procedure



## FAQs

For more information, see [FAQs > Capacity Expansion](#).

## 6.2 Step 1: Expand Disk Capacity

### Scenarios

When your EVS disk capacity is insufficient, you can expand the disk capacity on the console to prevent any data loss that may be caused by insufficient storage space.

### Prerequisites

**NOTE**

**Ensure that the disk meets the following conditions:**

- The status of a non-shared disk is **In-use** or **Available**.



- The status of a shared disk is **Available**. If the status is **In-use**, detach the disk from all of its servers before expanding the capacity.
- The disk has been backed up using CBR or snapshots. For details, see [Backing Up EVS Disks](#) and [Managing EVS Snapshots](#) respectively.

**If the disk status is In-use, ensure that the server meets the following conditions:**

- The server status must be **Running** or **Stopped**.
- The server OS must meet the requirements described in [Related Operations](#).

If the server OS does not meet the requirements, detach the disk and then expand the capacity. Otherwise, you may need to stop and start the server to see the additional space.

## Notes and Constraints

- Disk capacity can be expanded, but cannot be reduced.
- The maximum capacity of a system disk is 1 TiB, and that of a data disk is 32 TiB. The minimum expansion increment is 1 GiB for both system disks and data disks.

## Expanding the Capacity of a Single Disk

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

**Step 2** Choose an entry to expand the capacity.

- To expand the disk on the ECS console (suitable for a disk that has been attached to an ECS):
  - a. Choose **Computing** > **Elastic Cloud Server** to go to ECS list page.
  - b. Click the name of the server where the desired disk is attached to go to the **Summary** page.
  - c. Click the **Disks** tab, locate the disk you want to expand, and click **Expand Capacity** in the **Operation** column.
- To expand the disk on the EVS console:
  - a. Choose **Storage** > **Elastic Volume Service** to go to the EVS console.
  - b. Locate the disk you want to expand and click **Expand Capacity** in the **Operation** column.

**Step 3** On the **Expand Capacity** page, set **New Capacity** and click **Next**.

**Step 4** On the **Details** page, check the disk configuration.

- Click **Submit** to start the expansion.
- Click **Previous** to change the settings, if required.

**Step 5** In the displayed dialog box, click **OK**.

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

When the disk status changes from **Expanding** to **In-use** or **Available**, and the disk capacity increases, the expansion is successful.

 **NOTE**

When a disk is in the **Expanding** state, you cannot modify the specifications of the ECS where the disk is attached.

**Step 7** (Optional) If the disk status is **Available**, attach the disk to a server by referring to [Attaching an EVS Disk](#). Skip this step if the disk status is **In-use** (attached to a server).

**Step 8** After the disk has been expanded on the console, log in to the server and extend the partition and file system, because the previous steps only enlarge the disk space.

The operations vary depending on the server OS.


- Linux: [Extending Disk Partitions and File Systems \(Linux\)](#)
- Windows: [Extending Disk Partitions and File Systems \(Windows\)](#)

----End

## Related Operations

Perform the following operations to check whether your server OS allows you to expand **In-use** disks:

1. Check your server image. Certain public images and similar private images allow you to expand **In-use** disks. You do not need to stop and then start the servers after the expansion.

To view such images, log in to the console, click  in the navigation pane on the left, and choose **Compute > Image Management Service**. On the **Public Images** tab, view the images of the **ECS system disk image** type.

2. If your server OS is not in the image list, check whether it is included in [Table 6-1](#).

If it is included in [Table 6-1](#), you can expand capacity while the disk is in use and you do not need to stop and start the server after the expansion. Otherwise, you must detach the disk and then expand its capacity, or stop and start the server after the expansion.

**Table 6-1** 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
SUSE 12	SUSE Linux Enterprise Server 12 64-bit or later

OS	Version
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 64bit
Red Hat Enterprise Linux 6	6.8 64bit
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

## 6.3 Step 2: Extend Disk Partitions and File Systems

### 6.3.1 Extending Disk Partitions and File Systems (Linux)

#### Scenarios

After a disk is expanded on the console, the disk size is enlarged, but the disk partition and file system are not extended. You must log in to the server to extend the partition and file system before you can view and use the additional space. Specifically, you can **add the additional space to an existing partition and file system** or **create a new partition and file system with the additional space**.

This section describes how to extend partitions and file systems on a system or data disk in Linux. The extension operations may vary depending on the server OS.

**Table 6-2** Operation instructions of extending partitions and file systems in Linux

Scenario	Partition Style	Disk Function	OS Requirements	File System Format	Capacity Expansion Tool	Example Configuration
	GPT or MBR	System disk Data disk	<ul style="list-style-type: none"> <li>To extend partitions and file systems of a system disk, the kernel version must be later than 3.6.0.</li> <li>To extend partitions and file systems of a data disk, there is no requirements on the OS version.</li> </ul>	ext* (such as ext2, ext3, and ext4), xfs, and btrfs	growpart	<ul style="list-style-type: none"> <li>Device name: /dev/vdb</li> <li>Existing partition: /dev/vdb1</li> <li>Space added: 50 GiB</li> </ul>

Scenario	Partition Style	Disk Function	OS Requirements	File System Format	Capacity Expansion Tool	Example Configuration
	MBR	System disk	The kernel version is earlier than 3.6.0.	ext* (such as ext2, ext3, and ext4), xfs, and btrfs	dracut-modules - growroot	<ul style="list-style-type: none"> <li>• Device name: /dev/vda</li> <li>• File system format: ext4</li> <li>• Mount point: /mnt/sda</li> <li>• Partition name: /dev/vda1</li> <li>• Space added: 60 GiB</li> <li>• Partition style: MBR</li> </ul>
	MBR	System disk Data disk	None	ext* (such as ext2, ext3, and ext4), xfs, and btrfs	<ul style="list-style-type: none"> <li>• fdisk</li> <li>• parted</li> </ul>	<ul style="list-style-type: none"> <li>• Device name: /dev/vdb</li> <li>• File system format: ext4</li> <li>• Mount points: /mnt/sdc and /mnt/sdd</li> <li>• Partition 1: /dev/vdb1 <ul style="list-style-type: none"> <li>- Size: 100 GiB</li> <li>- Partition style: MBR</li> </ul> </li> <li>• Partition 2: /dev/vdb2 <ul style="list-style-type: none"> <li>- Size: 50 GiB</li> <li>- Partition style: MBR</li> </ul> </li> </ul>

Scenario	Partition Style	Disk Function	OS Requirements	File System Format	Capacity Expansion Tool	Example Configuration
	GPT	Data disk	None	ext* (such as ext2, ext3, and ext4), xfs, and btrfs	parted	<ul style="list-style-type: none"> <li>• Device name: /dev/vdb</li> <li>• File system format: ext4</li> <li>• Mount points: /mnt/sdc and /mnt/sdd</li> <li>• Partition 1: /dev/vdb1 <ul style="list-style-type: none"> <li>- Size: 2 TiB</li> <li>- Partition style: GPT</li> </ul> </li> <li>• Partition 2: /dev/vdb2 <ul style="list-style-type: none"> <li>- Size: 1 TiB</li> <li>- Partition style: GPT</li> </ul> </li> </ul>

 **NOTE**

You can run **uname -a** to check the Linux kernel version.

## Notes and 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.
- During an expansion, the additional space is added to the end of the disk. If the disk has multiple partitions, the additional space can only be allocated to the last partition of the disk.
- If a disk uses MBR, the storage space in excess of 2 TiB cannot be used because the maximum capacity that MBR supports is 2 TiB. If your disk already uses MBR for partitioning and you require more than 2 TiB after the capacity expansion, do as follows:
  - (Recommended) Create a new EVS disk and use GPT.
  - Back up the disk data, perform the expansion, and then change the partition style from MBR to GPT. During this change, services will be interrupted and data on the disk will be erased.

## Prerequisites

- You have expanded the disk capacity and attached the disk to a server on the console. For details, see [Step 1: Expand Disk Capacity](#).
- The disk has been backed up using CBR or snapshots. For details, see [Backing Up EVS Disks](#) and [Managing EVS Snapshots](#) respectively.
- 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*.

## Extending an Existing Partition

Originally, data disk `/dev/vdb` has 100 GiB and one partition `/dev/vdb1`. Then, the data disk is expanded to 150 GiB. The following example shows how to allocate the additional 50 GiB to the existing `/dev/vdb1` partition.

**Step 1** Check whether the growpart tool is installed.

### growpart

- If the tool instructions are returned, the tool has been installed, and you do not need to install it again.

```
[root@ecs-centos76 ~]# growpart
growpart disk partition
rewrite partition table so that partition takes up all the space it can
options:
-h | --help          print Usage and exit
  --fudge F          if part could be resized, but change would be
                    less than 'F' bytes, do not resize (def ault: 1048576)
-N | --dry-run       only report what would be done, show new 'sfdisk -d'
-v | --verbose       increase verbosity / debug
-u | --update R      update the the kernel partition table info after growing
                    this requires kernel support and 'partx --update'
                    R is one of:
                    - 'auto': [default] update partition if possible
                    - 'force' : try despite- sanity checks (fail on failure)
                    - 'off'  : do not attempt
                    - 'on'   : fail if sanity checks indicate no support

Example:
- growpart /dev/sda 1
  Resize partition 1 on /dev/sda
must supply disk and part it ion-number
[root@ecs-centos76 ~]#
```

- If no tool instructions are returned, run the following command to install the tool:

### yum install cloud-utils-growpart

```
Loaded plugins: fastestmirror
Determining fastest mirrors
epel/x86_64/metalink
| 8.0 kB  00:00:00
...
Package cloud-utils-growpart-0.29-2.el7.noarch already installed and latest version
Nothing to do
```

The installation is successful.

**Step 2** Check the partitions of the `/dev/vdb` disk.

### lsblk

```
[root@ecs-centos76 ~]# 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
```

We can see that **/dev/vdb** has 150 GiB, the **/dev/vdb1** partition has 100 GiB, and the additional 50 GiB space is not allocated.

If the disk is not partitioned, you need to directly extend the file system, go to [Step 4](#).

### Step 3 Add the additional space to the **/dev/vdb1** partition.

#### growpart /dev/vdb 1

```
[root@ecs-test-0001 ~]# growpart /dev/vdb 1
CHANGED: partition=1 start=2048 old: size=209713152 end=209715200 new:
size=314570719,end=314572767
```

#### NOTE

- If the following command output is displayed:  
no tools available to resize disk with 'gpt'  
FAILED: failed to get a resizer for id "  
  
The disk uses the GPT partition style, and the **gdisk** tool is required when you use **growpart** to add the additional space. In this case, run **yum install gdisk**, enter **y** to install **gdisk**, and then run the preceding **growpart** command.
- If the following command output is displayed:  
growpart /dev/vda 1 unexpected output in sfdisk --version [sfdisk is from util-linux 2.23.2]  
  
Check whether the system character set (language environment) is **en\_US.UTF-8**. If not, run **export LC\_ALL=en\_US.UTF-8**.
- If error message "NOCHANGE:partition 1 is size xxxxxxx. it cannot be grown" or "No space left on the block device" is returned, the expansion may be failed because the server disk is full (at 100% usage). Back up the disk data and clear unnecessary files or programs.

### Step 4 Extend the file system of the **/dev/vdb1** partition.

1. Check the file system format of the **/dev/vdb1** partition.

#### parted /dev/vdb

**P**

```
[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) p
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 107GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:
```

Number	Start	End	Size	File system	Name	Flags
1	1049KB	107GB	107GB	ext4	/dev/vdb1	

(parted)

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

2. Extend the file system. As the file system format of **/dev/vdb1** is **ext4**, we use the following command.

#### resize2fs /dev/vdb1

```
[root@ecs-test-0001 ~]# resize2fs /dev/vdb1
resize2fs 1.42.9 (28-Dec-2013)
```



```
Filesystem at /dev/vdb1 is mounted on /mnt/sdc; on-line resizing required
old_desc_blocks = 13, new_desc_blocks = 19
The filesystem on /dev/vdb1 is now 39321339 blocks long.
```

**NOTE**

- If the error message "open: No such file or directory while opening /dev/vdb1" is returned, an incorrect partition is specified. Run **parted** to view disk partitions.
- If the file system format is xfs, run the following command (**/mnt/sdc** is the mount point of **/dev/vdb1**. Change it based on your actual conditions):

**sudo xfs\_growfs /mnt/sdc**

```
[root@ecs-test-0001 ~]# sudo xfs_growfs /mnt/sdc
meta-data=/dev/vdb1          isize=512  agcount=4, agsize=6553536 blks
         =                   sectsz=512  attr=2, projid32bit=1
         =                   crc=1      finobt=0  spinodes=0
data     =                   bsize=4096  blocks=26214144, imaxpct=25
         =                   sunit=0    swidth=0 blks
naming   =version 2          bsize=4096  ascii-ci=0  ftype=1
log      =internal          bsize=4096  blocks=12799, version=2
         =                   sectsz=512  sunit=0    blks, lazy-count=1
realtime =none              extsz=4096  blocks=0,  rtextents=0
data blocks changed from 26214144 to 39321339
```

**Step 5** Check the partition size after extension.**lsblk**

```
[root@ecs-centos76 ~]# 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 150G  0 part /mnt/sdc
```

We can see that the **/dev/vdb** data disk now has 150 GiB and the **/dev/vdb1** partition has 150 GiB, meaning that the extension operation is successful.

**NOTE**

If you are expanding a data disk and the OS kernel is earlier than 3.6.0, after the partition and file system are extended, you need to run **reboot** to make the additional space available for use. Restarting the OS will interrupt services. To prevent any data loss after the restart, ensure that you have backed up the disk data before the restart. To back up data using CBR, see [Backing Up EVS Disks](#). To back up data using snapshots, see [Managing EVS Snapshots](#).

----End

## Extending an Existing MBR Partition (for System Disks Whose Kernel Version Is Earlier Than 3.6.0)

Originally, system disk **/dev/vda** has 40 GiB and one partition **/dev/vda1**. Then, the disk is expanded to 100 GiB. The following example shows you how to allocate the additional 60 GiB to the existing **/dev/vda1** partition.

**NOTICE**

- If the OS kernel version is earlier than 3.6.0, you need to reboot the system after extending an existing MBR partition to make the additional space available. During the reboot, services will be interrupted. After the reboot, the additional space is automatically added to the last partition of the system disk.
- To prevent data loss after a reboot, you are advised to use [CBR](#) to back up the disk data before initializing a disk.
- If your OS kernel version is earlier than 3.6.0 and you want to create a new partition with the additional space, see [Creating a New MBR Partition](#).

**Step 1** (Optional) Install the dracut-modules-growroot tool.

**yum install dracut-modules-growroot**

```
[root@ecs-test-0002 ~]# yum install dracut-modules-growroot
Loaded plugins: fastestmirror, security
Setting up Install Process
Loading mirror speeds from cached hostfile
epel/metalink | 4.3 kB
00:00
* epel: pubmirror1.math.uh.edu
base | 3.7 kB
00:00
extras | 3.4 kB
00:00
updates | 3.4 kB
00:00
Package dracut-modules-growroot-0.20-2.el6.noarch already installed and latest version
Nothing to do
```

**NOTE**

Skip this step if the tool is already installed.

**Step 2** Regenerate the **initramfs** file.

**dracut -f****NOTE**

The **initramfs** file helps the Linux kernel to access drivers on external storage devices.

**Step 3** Check the information of the **/dev/vda** disk.

**lsblk**

```
[root@ecs-test-0002 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda 253:0 0 100G 0 disk
├vda1 253:1 0 40G 0 part /
vdb 253:16 0 100G 0 disk
├vdb1 253:17 0 100G 0 part /mnt/sdc
```

We can see that the **/dev/vda** system disk has the **/dev/vda1** partition, then the disk is expanded to 100 GiB, and the additional space is not allocated. So, **/dev/vda** has 100 GiB, and **/dev/vda1** has 40 GiB.

**Step 4** Restart the server.

**reboot**

Reconnect to the server after it is restarted.

**Step 5** Check the information of the `/dev/vda` disk.

#### lsblk

```
[root@ecs-test-0002 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda 253:0 0 100G 0 disk
├vda1 253:1 0 100G 0 part /
vdb 253:16 0 100G 0 disk
├vdb1 253:17 0 100G 0 part /mnt/sdc
```

We can now see that `/dev/vda` has 100 GiB and `/dev/vdb1` also has 100 GiB.

----End

## Creating a New MBR Partition

Originally, data disk `/dev/vdb` has 100 GiB and one partition `/dev/vdb1`, and then the disk is expanded to 150 GiB. The following example shows you how to use `fdisk` to allocate the additional 50 GiB to a new partition (`/dev/vdb2`).

**Step 1** Check the information of the `/dev/vdb` disk.

1. Check disk partition sizes.

#### lsblk

```
[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
```

We can see that the `/dev/vdb` data disk has the `/dev/vdb1` partition, then 50 GiB is added to the disk, and the additional 50 GiB is not allocated.

So, `/dev/vdb` has 150 GiB, and `/dev/vdb1` has 100 GiB.

2. Check the disk partition style.

#### parted /dev/vdb

#### p

```
[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
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 161GiB
Sector size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:
```

Number	Start	End	Size	File system	Name	Flags
1	1049kB	107GiB	107GiB	ext4	/dev/vdb1	

(parted)

In this example, the disk uses MBR.

Enter `q` and press **Enter** to exit parted.

 NOTE

- If **Partition Table: msdos** is returned, the partition style is MBR.
- If **Partition Table: gpt** is returned, the partition style is GPT.
- If **Partition Table: loop** is returned, the disk is not partitioned (the entire disk is partitioned into one partition), and only a file system is created.

**Step 2** Use the additional space to create a second primary partition **/dev/vdb2** on the **/dev/vdb** disk.

1. Create the partition.

**fdisk /dev/vdb**

**n**

**p**

```
[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): n
```

```
Partition type:
```

```
  p primary (0 primary, 0 extended, 4 free)
  e extended
```

```
Select (default p): p
```

```
Partition number (2-4, default 2):
```

**Partition type** shows that there are two types of partitions. Choosing **p** creates a primary partition and choosing **e** creates an extended partition.

**Partition number** indicates the serial number of the primary partition. Because partition number **1** has been used, the value ranges from **2** to **4**.

 NOTE

MBR supported up to four primary partitions or three primary partitions plus one extended partition.

The number of logical partitions allowed in the extended partition is not limited, so theoretically you can create as many logical partitions as you want. If you need five or more partitions, use the "primary partitions + one extended partition" model and then create logical partitions in the extended partition.

2. Enter **2** as the primary partition number and view the first sector range.

```
Partition number (2-4, default 2): 2
```

```
First sector (83886080-209715199, default 83886080):
```

**First sector** shows the first sector range. The value ranges from **83886080** to **209715199**, and the default value is **83886080**.

3. Press **Enter** to use the default first sector and then press **Enter** to use the default last sector.

```
First sector (83886080-209715199, default 83886080):
```

```
using default value 83886080
```

```
Last sector, +sectors or +size{K,M,G} (83886080-209715199, default 209715199):
```

```
using default value 209715199
```

```
Partition 2 of type Linux and of size 40 GB is set
```

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

**Last sector** shows the last sector range. The value ranges from **83886080** to **209715199**, and the default value is **209715199**.

 **NOTE**

If you want to create two or more partitions, calculate the first and last sectors of the partitions as follows:

Assume that the `/dev/vdb` data disk has 100 GiB, and you are going to partition it into two primary partitions, first primary partition `/dev/vdb1` (40 GiB) and second primary partition `/dev/vdb2` (60 GiB).

Based on the facts that **Capacity = Sectors value x 512 bytes** and **1 GiB = 1073741824 bytes**, the sector value can be calculated using **Sectors value = Capacity/512 bytes**.

- Sector value of the data disk `/dev/vdb` (100 GiB) is **209715200** ( $100 \times 1073741824/512$ ), so the disk's last sector is **209715199** ( $209715200 - 1$ ).

In the preceding example, **First sector (2048-209715199, default 2048)** means that the first sector of the disk ranges from **2048** to **209715199**.

- Sector value of the first primary partition `/dev/vdb1` (40 GiB) is **83886080** ( $40 \times 1073741824/512$ ), so the partition's last sector is **83886079** ( $83886080 - 1$ ).

In this example, the default first sector is used as first sector of this partition, which is **2048**.

- Sector value of the second primary partition `/dev/vdb2` (60 GiB) is **125829120** ( $60 \times 1073741824/512$ ), so the partition's last sector is **125829119** ( $125829120 - 1$ ).

First sector = Last sector of `/dev/vdb1` + 1 =  $83886079 + 1 = 83886080$

Last sector = First sector + Sector value - 1 =  $83886080 + 125829120 - 1 = 209715199$

**Step 3** Check the size and partition style of the new partition.

1. Enter **p** and press **Enter** to print details of the `/dev/vdb2` partition.

Command (m for help): p

```
Disk /dev/vdb: 107.4 GB, 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: 0x994727e5
```

Device	Boot	Start	End	Blocks	Id	System
/dev/vdb1		2048	83886079	41942016	83	Linux
/dev/vdb2		83886080	209715199	62914560	83	Linux

Command (m for help):

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

 **NOTE**

In case that you want to discard the changes made before, you can exit fdisk by entering **q** and press **Enter**. Then, re-create the partition.

3. Synchronize the new partition table to the OS.

**partprobe**

 NOTE

If error message **-bash: partprobe: command not found** is returned, the system cannot identify the command. In this case, run **yum install -y parted** to install the command. Then run the command again.

- 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 extend the disk, the end of the disk changes accordingly. In this case, enter **Fix** to move the backup file of the information to the 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

After you enter **Fix**, the system automatically sets the GPT partition style for the additional space.

#### Step 4 Create an ext4 file system on the **/dev/vdb2** partition.

```
mkfs -t ext4 /dev/vdb2
```

 NOTE

It takes some time to create a file system. Observe the system running status and do not exit.

```
[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
2621440 inodes, 10485504 blocks
524275 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
```

Run **parted /dev/vdb** and enter **p** to check the file system format.

```
[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
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 107GiB
Sector size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:
```

```

Number Start  End   Size  Type  File system  Flags
1    1049kB 42.9GB 42.9GB primary ext4
2    42.9GB 107GB 64.4GB primary ext4

(parted) q
[root@ecs-test-0001 ~]#

```

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

An ext4 file system is created for the **/dev/vdb2** partition.

**Step 5** Create a directory (mount point) and mount the new partition on the created mount point.

```
mkdir -p /mnt/sdd
```

```
mount /dev/vdb2 /mnt/sdd
```

```
lsblk
```

View the mount results.

```

[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
└─vdb2 253:18  0  50G  0 part /mnt/sdd

```

You should now see that partition **/dev/vdb2** is mounted on **/mnt/sdd**.

**Step 6** Use the partition UUID to configure auto mount at startup.

#### NOTE

- If device names are used to identify disks in the **/etc/fstab** file, your server may fail to run after reboot. This is because device names are assigned dynamically and may change (for example, from **/dev/vdb1** to **/dev/vdb2**) after a stop or start.
- UUIDs are the unique character strings for identifying partitions in Linux.

1. Query the UUID of the **/dev/vdb2** partition.

```
blkid /dev/vdb2
```

```

[root@ecs-test-0001 ~]# blkid /dev/vdb2
/dev/vdb2: UUID="0b3040e2-1367-4abb-841d-ddb0b92693df" TYPE="ext4"

```

The UUID of **/dev/vdb2** is **0b3040e2-1367-4abb-841d-ddb0b92693df**.

2. Configure auto mount at startup.

```
vi /etc/fstab
```

Press **i** to enter editing mode, move the cursor to the end of the file, press **Enter**, and add the partition information.

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

Press **Esc**, enter **:wq**, and press **Enter** to save the settings and exit the vi editor.

**Table 6-3** Parameter description

Example Value	Description
UUID=0b3040e2-1367-4abb-841d-ddb0b92693df	The UUID of the partition.

Example Value	Description
/mnt/sdc	The mount point of the partition.
ext4	The file system format of the partition.
defaults	The partition mount option. Normally, this parameter is set to <b>defaults</b> .
0	<ul style="list-style-type: none"> <li>- The Linux dump backup option. <ul style="list-style-type: none"> <li>▪ <b>0</b>: Linux dump backup is not used. Usually, dump backup is not used, and you can set this parameter to <b>0</b>.</li> <li>▪ <b>1</b>: Linux dump backup is used.</li> </ul> </li> </ul>
2	<ul style="list-style-type: none"> <li>- The fsck option, which means whether to use fsck to check the disk during startup. <ul style="list-style-type: none"> <li>▪ <b>2</b>: The check starts from the partitions whose mount points are non-root directories. / is the root directory.</li> <li>▪ <b>1</b>: The check starts from the partitions whose mount points are root directories.</li> <li>▪ <b>0</b>: The fsck option is not used.</li> </ul> </li> </ul>

**Step 7** Verify that auto mount takes effect.

```
umount /dev/vdb2
```

```
mount -a
```

The system reloads all the content in the **/etc/fstab** file.

Query file system mounting information.

```
mount | grep /mnt/sdd
```

If information similar to the following is displayed, auto mount has taken effect:

```
root@ecs-test-0001 ~]# mount | grep /mnt/sdd
/dev/vdb2 on /mnt/sdd type ext4 (rw,relatime,data=ordered)
```

----End

## Creating a New GPT Partition

Originally, data disk **/dev/vdb** has 2 TiB and one partition **/dev/vdb1**, and then the disk is expanded to 3 TiB. The following example shows you how to



use `parted` to allocate the additional 1 TiB to a new GPT partition (`/dev/vdb2`).

**Step 1** Check the information of the `/dev/vdb` disk.

1. Check disk partition sizes.

#### `lsblk`

```
[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 3T 0 disk
└─vdb1 253:17 0 2T 0 part /mnt/sdc
```

We can see that the `/dev/vdb` data disk has the `/dev/vdb1` partition, then 1 TiB is added to the disk, and the additional space is not allocated. So, `/dev/vdb` has 3 TiB, and `/dev/vdb1` has 2 TiB.

2. Check the disk partition style.

#### `parted /dev/vdb`

##### `p`

```
[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
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 3299GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number Start End Size File system Name Flags
1 1049kB 2199GB 2199GB ext4 /dev/vdb1

(parted)
```

In this example, the disk uses GPT.

Enter `q` and press **Enter** to exit `parted`.

#### NOTE

- If **Partition Table: msdos** is returned, the partition style is MBR.
- If **Partition Table: gpt** is returned, the partition style is GPT.
- If **Partition Table: loop** is returned, the disk is not partitioned (the entire disk is partitioned into one partition), and only a file system is created.

**Step 2** Create a new partition `/dev/vdb2` on the `/dev/vdb` disk.

1. Create the `/dev/vdb2` partition.

#### `parted /dev/vdb`

##### `unit s`

##### `p`

```
[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) unit s
(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    4294965247s    4294963200s    ext4        /dev/vdb1
(parted)
```

Take note of the last sector of the **/dev/vdb1** partition, which is **4294965247s**.

#### NOTE

- If error message **-bash: parted: command not found** is returned, the system cannot identify the command. In this case, run **yum install -y parted** to install the command. Then run the command again.
- 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)?  
Fix/Ignore/Cancel?

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 extend the disk, the end of the disk changes accordingly. In this case, enter **Fix** to move the backup file of the information to the 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?

After you enter **Fix**, the system automatically sets the GPT partition style for the additional space.

2. Set the partition name and size.

```
mkpart /dev/vdb2 4294965248s 100%
```

```
p
```

#### NOTE

In the command, **4294965248s** is the first sector of this partition, which is the last sector of the **/dev/vdb1** partition plus one, and **100%** sets the last sector of this partition, which means to use 100% of the disk remaining space for **/dev/vdb2**.

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

3. Check the **/dev/vdb2** partition.

```
lsblk
```

```
[root@ecs-centos74 ~]# 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 3T 0 disk
├─vdb1 253:17 0 2T 0 part /mnt/sdc
└─vdb2 253:18 0 1T 0 part
```

- Step 3** Create an ext4 file system on the **/dev/vdb2** partition.

```
mkfs -t ext4 /dev/vdb2
```

#### NOTE

It takes some time to create a file system. Observe the system running status and do not exit.

```
[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  
67108864 inodes, 268435456 blocks  
13421772 blocks (5.00%) reserved for the super user  
First data block=0  
Maximum filesystem blocks=2415919104  
8192 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  
  
Allocating group tables: done  
Writing inode tables: done  
Creating journal (32768 blocks): done  
Writing superblocks and filesystem accounting information: done  
[root@ecs-test-0001 ~]#
```

Run **parted /dev/vdb** and enter **p** to check the file system format.

```
[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  
Model: Virtio Block Device (virtblk)  
Disk /dev/vdb: 3299GB  
Sector size (logical/physical): 512B/512B  
Partition Table: gpt  
Disk Flags:  
  
Number  Start  End      Size    File system  Name      Flags  
1       1049kB 2199GB  2199GB  ext4         /dev/vdb1  
2       2199GB 3299GB  1100GB  ext4         /dev/vdb2  
  
(parted) q  
[root@ecs-test-0001 ~]#
```

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

**Step 4** Create a directory (mount point) and mount the new partition on the created mount point.

```
mkdir -p /mnt/sdc
```

```
mount /dev/vdb1 /mnt/sdc
```

```
lsblk
```

```
[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 3T 0 disk  
└vdb1 253:17 0 2T 0 part /mnt/sdc  
└vdb2 253:18 0 1T 0 part /mnt/sdd
```

You should now see that partition **/dev/vdb2** is mounted on **/mnt/sdd**.

**Step 5** Use the partition UUID to configure auto mount at startup.

 **NOTE**

- If device names are used to identify disks in the `/etc/fstab` file, your server may fail to run after reboot. This is because device names are assigned dynamically and may change (for example, from `/dev/vdb1` to `/dev/vdb2`) after a stop or start.
- UUIDs are the unique character strings for identifying partitions in Linux.

1. Query the UUID of the `/dev/vdb2` partition.

**blkid /dev/vdb2**

```
[root@ecs-test-0001 ~]# blkid /dev/vdb2
/dev/vdb2: UUID="0b3040e2-1367-4abb-841d-ddb0b92693df" TYPE="ext4"
```

The UUID of `/dev/vdb2` is **0b3040e2-1367-4abb-841d-ddb0b92693df**.

2. Configure auto mount at startup.

**vi /etc/fstab**

Press **i** to enter editing mode, move the cursor to the end of the file, press **Enter**, and add the partition information.

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

Press **Esc**, enter **:wq**, and press **Enter** to save the settings and exit the vi editor.

**Table 6-4** Parameter description

Example Value	Description
UUID=0b3040e2-1367-4abb-841d-ddb0b92693df	The UUID of the partition.
/mnt/sdc	The mount point of the partition.
ext4	The file system format of the partition.
defaults	The partition mount option. Normally, this parameter is set to <b>defaults</b> .
0	<ul style="list-style-type: none"> <li>- The Linux dump backup option. <ul style="list-style-type: none"> <li>▪ <b>0</b>: Linux dump backup is not used. Usually, dump backup is not used, and you can set this parameter to <b>0</b>.</li> <li>▪ <b>1</b>: Linux dump backup is used.</li> </ul> </li> </ul>

Example Value	Description
2	<ul style="list-style-type: none"> <li>- The fsck option, which means whether to use fsck to check the disk during startup. <ul style="list-style-type: none"> <li>▪ <b>2</b>: The check starts from the partitions whose mount points are non-root directories. / is the root directory.</li> <li>▪ <b>1</b>: The check starts from the partitions whose mount points are root directories.</li> <li>▪ <b>0</b>: The fsck option is not used.</li> </ul> </li> </ul>

**Step 6** Verify that auto mount takes effect.

```
umount /dev/vdb2
```

```
mount -a
```

The system reloads all the content in the `/etc/fstab` file.

Query file system mounting information.

```
mount | grep /mnt/sdd
```

If information similar to the following is displayed, auto mount has taken effect:

```
root@ecs-test-0001 ~]# mount | grep /mnt/sdd
/dev/vdb2 on /mnt/sdd type ext4 (rw,relatime,data=ordered)
```

```
----End
```

## 6.3.2 Extending Disk Partitions and File Systems (Windows)

### Scenarios

After a disk is expanded on the console, the disk size is enlarged, but the disk partition and file system are not extended. You must log in to the server to extend the partition and file system before you can view and use the additional space. Specifically, you can **add the additional space to an existing partition and file system** or **create a new partition and file system with the additional space**.

This section describes how to extend partitions and file systems on a system or data disk in Windows. The extension operations may vary depending on the server OS.

- Extending an Existing Partition
- Creating a New Partition

### Notes and 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.

- During an expansion, the additional space is added to the end of the disk. If the disk has multiple partitions, the additional space can only be allocated to the last partition of the disk.
- If a disk uses MBR, the storage space in excess of 2 TiB cannot be used because the maximum capacity that MBR supports is 2 TiB. If your disk already uses MBR for partitioning and you require more than 2 TiB after the capacity expansion, do as follows:
  - (Recommended) Create a new EVS disk and use GPT.
  - Back up the disk data, perform the expansion, and then change the partition style from MBR to GPT. During this change, services will be interrupted and data on the disk will be erased.

## Prerequisites

- You have expanded the disk capacity and attached the disk to a server on the console. For details, see [Step 1: Expand Disk Capacity](#).
- The disk has been backed up using CBR or snapshots. For details, see [Backing Up EVS Disks](#) and [Managing EVS Snapshots](#) respectively.
- 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*.

## Extending an Existing Partition

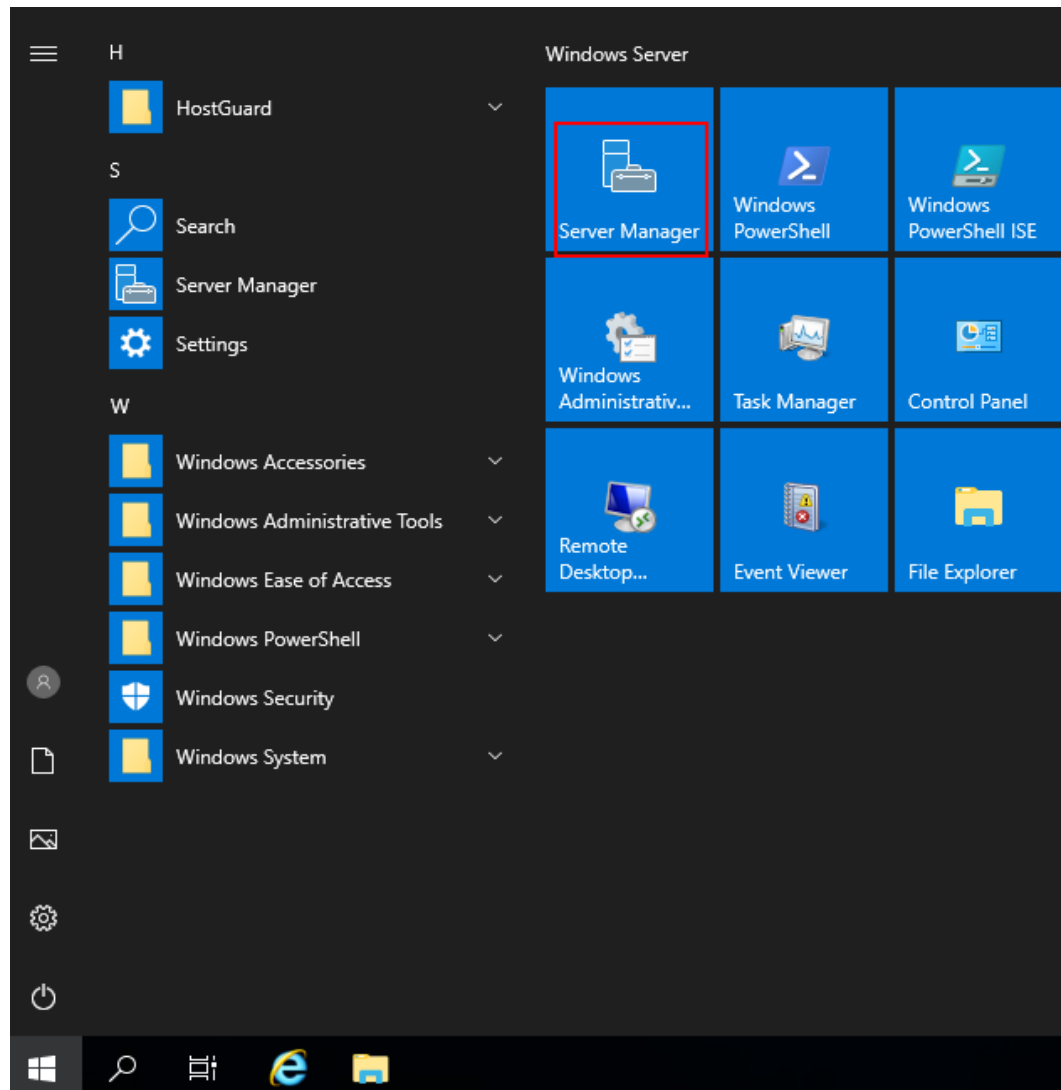
**Originally, the D drive in the Windows Server 2019 has 60 GiB, and then 30 GiB is added to the disk. The following example shows how to allocate the additional 30 GiB to the D drive.**

**Step 1** Log in to the server. On the server desktop, click the start icon in the lower left corner.

The **Windows Server** window is displayed.


**Step 2** Click **Server Manager** to open **Server Manager**.

**Figure 6-2** Server Manager



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

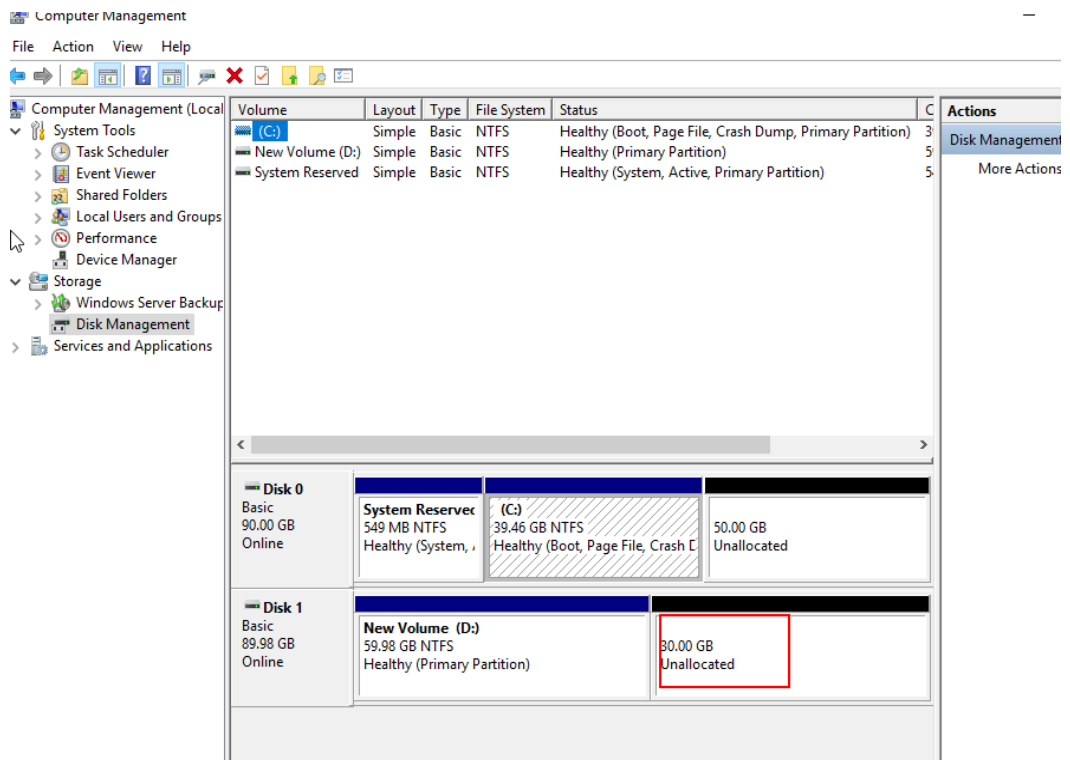
**Figure 6-3** Computer Management



**Step 4** Choose **Storage > Disk Management** to go to the disk list page.

The **Unallocated** area shows the newly added disk space, which is not added to any partition or file system. Now we will perform the following steps to **add the additional space to an existing partition and a file system**.

**Figure 6-4** Disk expanded but additional space not allocated





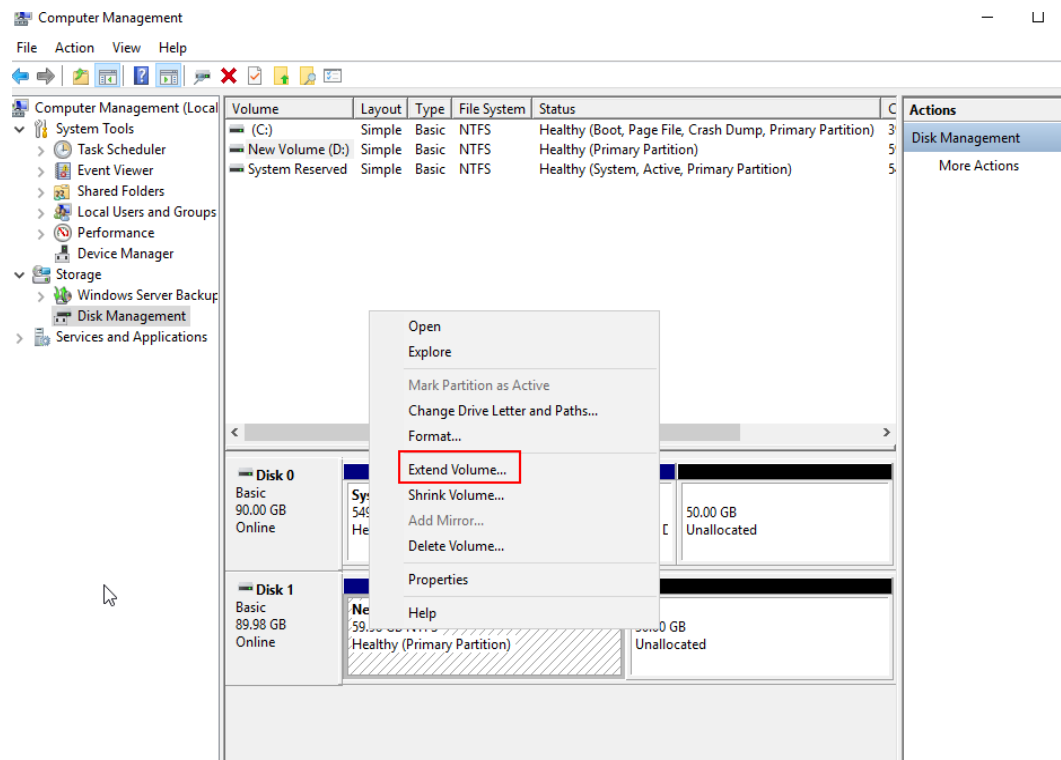
 **NOTE**

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

**Step 5** On the **Disk Management** page, find the disk and volume that you want to extend. Check the size and unallocated space.

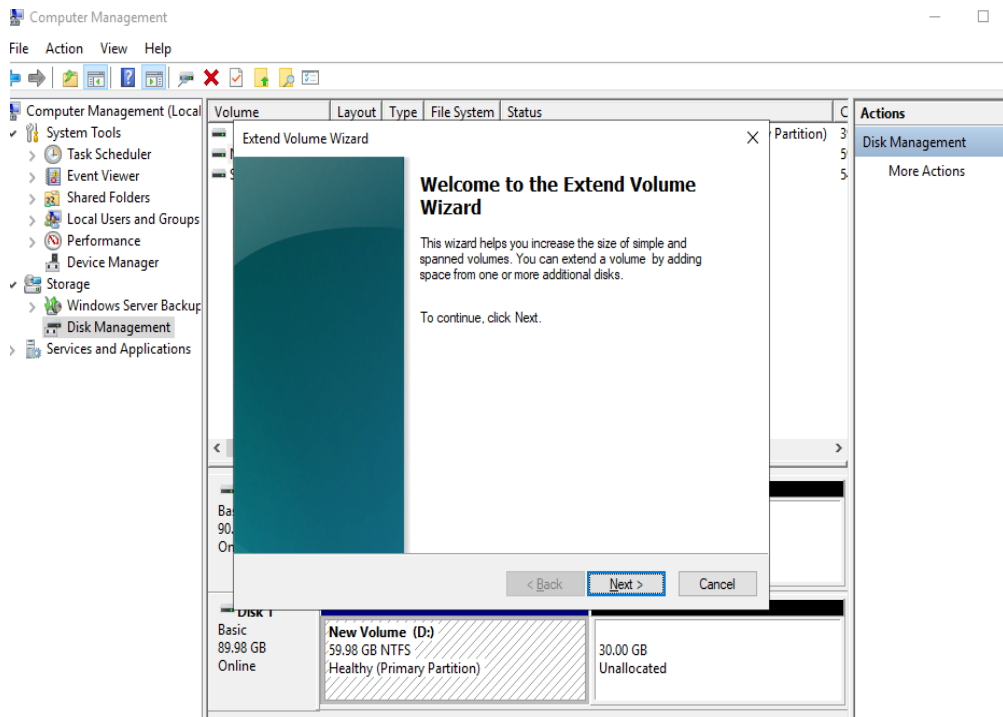
**Step 6** Right-click the volume and choose **Extend Volume** from the shortcut menu. In this example, right-click **New Volume (D:)**.

**Figure 6-5** Choosing Extend Volume



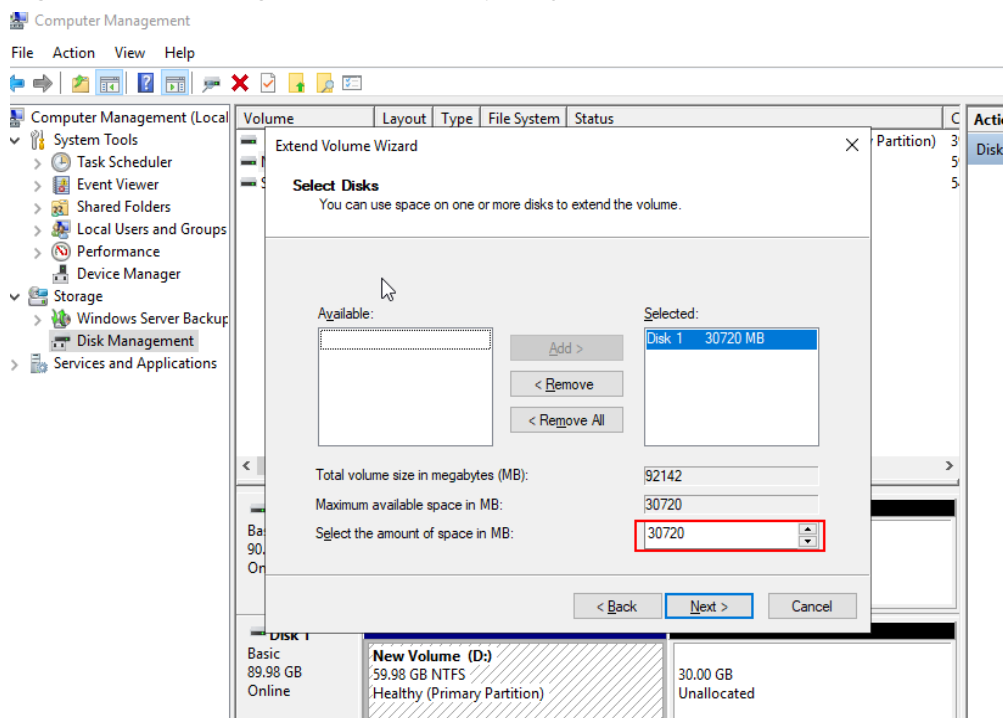
On the displayed **Extend Volume Wizard** windows, click **Next**.

**Figure 6-6** Extend Volume Wizard



**Step 7** In the text box to the right of **Select the amount of space in MB**, enter the amount of space you want to add and click **Next**. The default setting is used in this example.


**Figure 6-7** Selecting the amount of space you want to add



**Step 8** Click **Finish**.

After the extension succeeded, the volume size is greater than the original size.

**Figure 6-8** Extension succeeded



----End

## Creating a New Partition

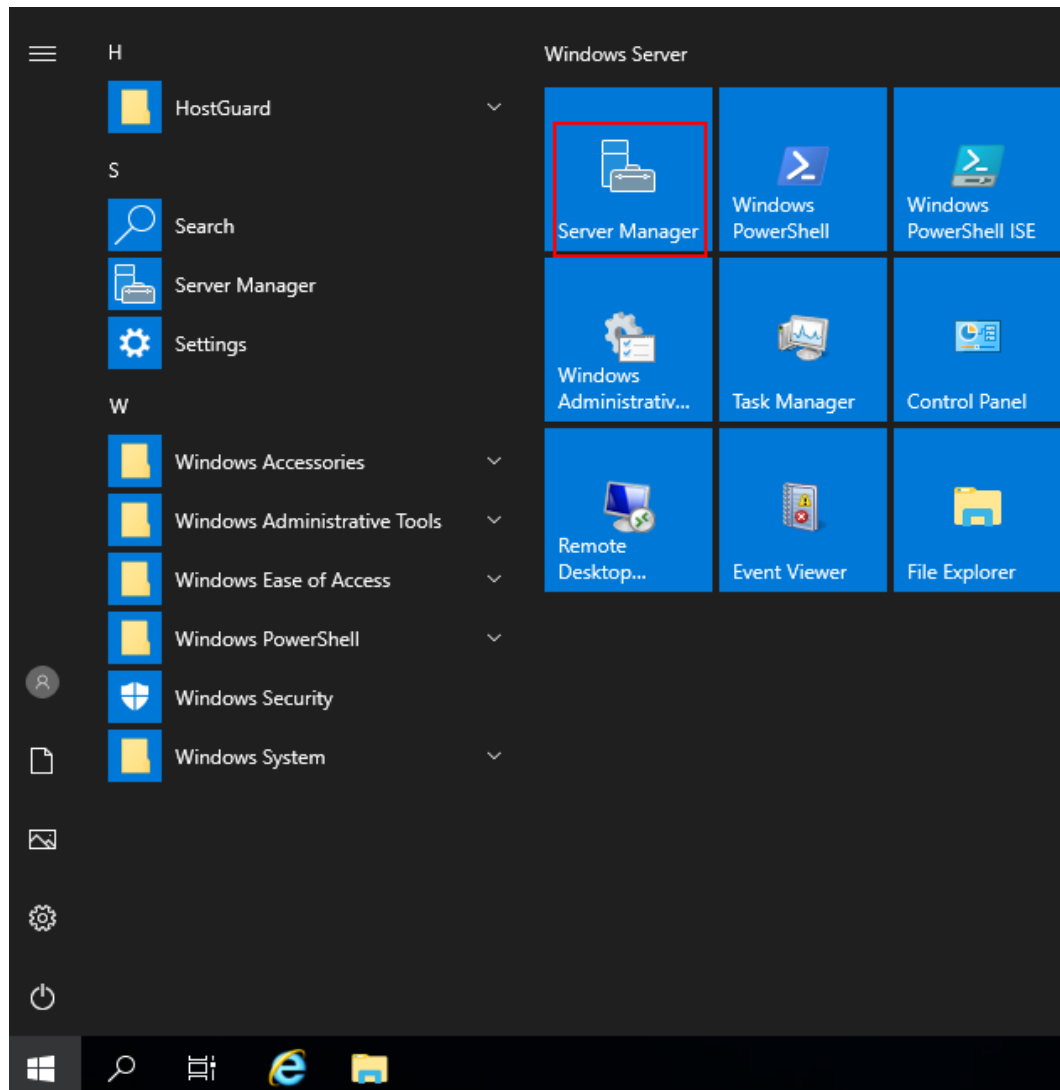
The following example shows you how to create a GPT partition with an NTFS file system on a server running Windows Server 2019.

**Step 1** Log in to the server. On the server desktop, click the start icon in the lower left corner.

The **Windows Server** window is displayed.

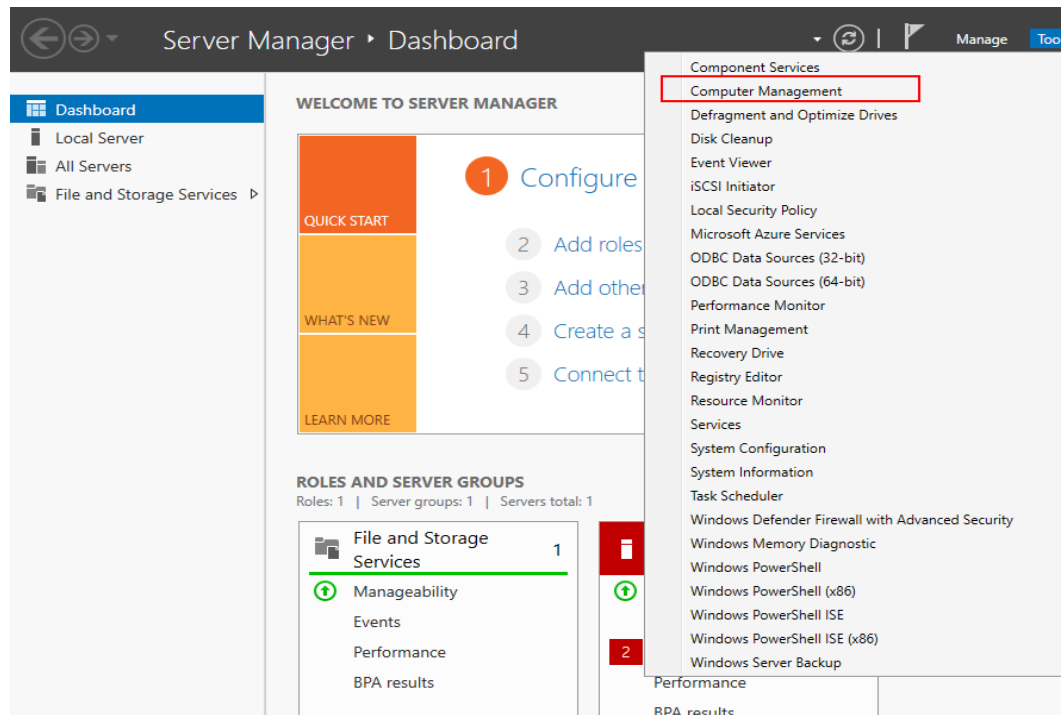
**Step 2** Click **Server Manager** to open **Server Manager**.

**Figure 6-9** Server Manager




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

Figure 6-10 Computer Management



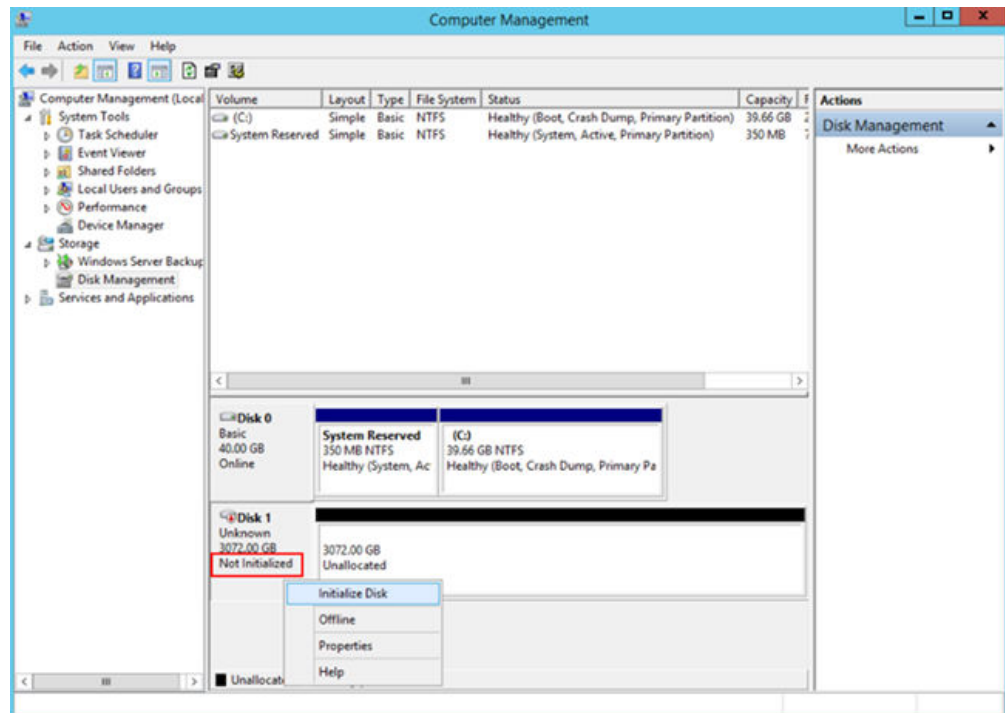
- Step 4** Choose **Storage > Disk Management** to go to the disk list page.
- Disks are displayed in the right pane. If there is a disk that is not initialized, the system will prompt you with the **Initialize Disk** dialog box.

Figure 6-11 Disk list



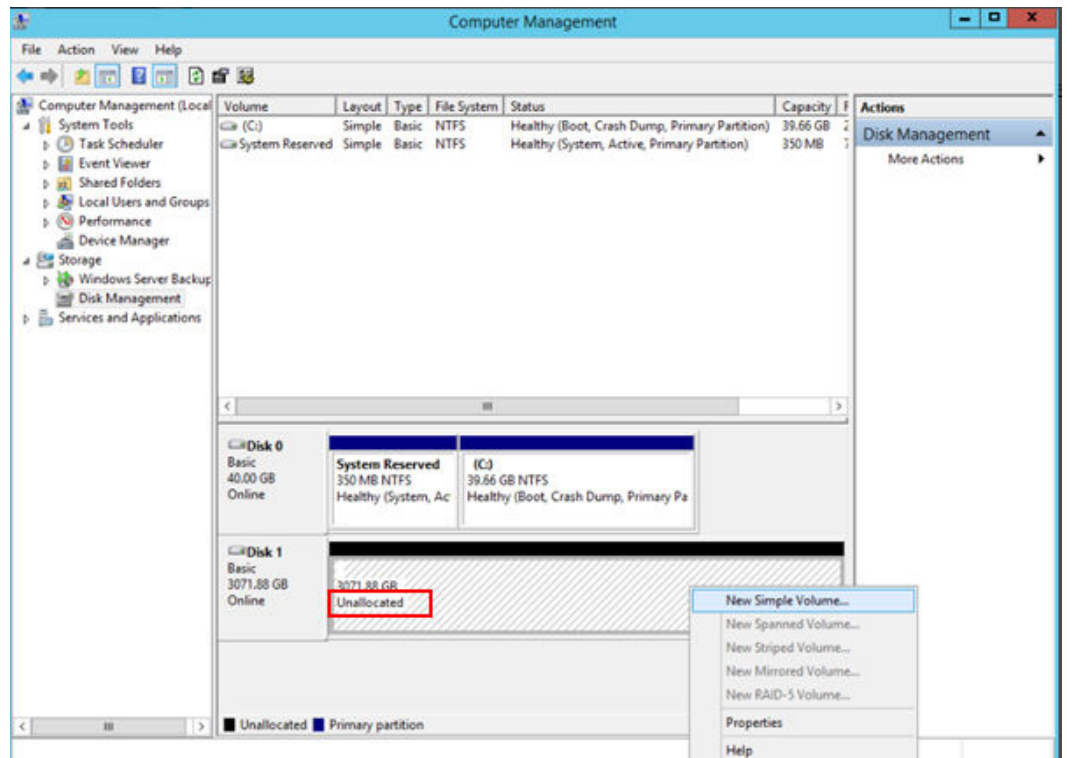
- If the **Initialize Disk** dialog box is not prompted up and the disk has no partitions (entire disk shown as **Unallocated**), right-click the area where the to-be-initialized disk is and choose **Initialize Disk** from the shortcut menu.

Figure 6-12 Initialize Disk



- If the **Initialize Disk** dialog box is not prompted up but the disk has a partition (primary partition) and unallocated space, the disk has been expanded. Now you need to extend the partition and file system by either **creating a new partition and file system with the additional space** or **adding the additional space to an existing partition and file system**.
  - To create a new partition and file system, go to [Step 5](#) and subsequent steps.
  - To allocate the additional space to an existing partition and file system, go to [Extending an Existing Partition](#).

**Figure 6-13** Disk expanded but additional space not allocated



**NOTE**


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

**Step 5** On the **Initialize Disk** dialog box, select **GPT (GUID Partition Table)** and click **OK** to go back to **Computer Management**.

**NOTE**


If your disk size is greater than 2 TiB or you may expand it to more than 2 TiB, select **GPT (GUID Partition Table)**.

Figure 6-14 Selecting GPT



**Step 6** In the **Unallocated** area of **Disk 1**, right-click and choose **New Simple Volume** from the shortcut menu and initialize the disk as prompted.

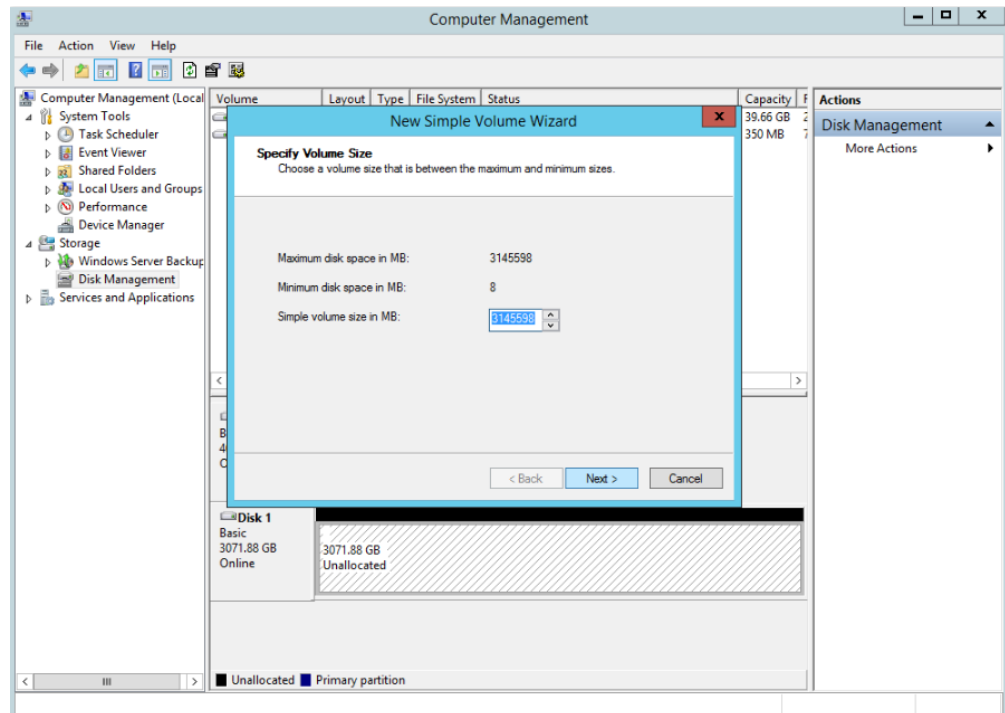
Figure 6-15 New Simple Volume





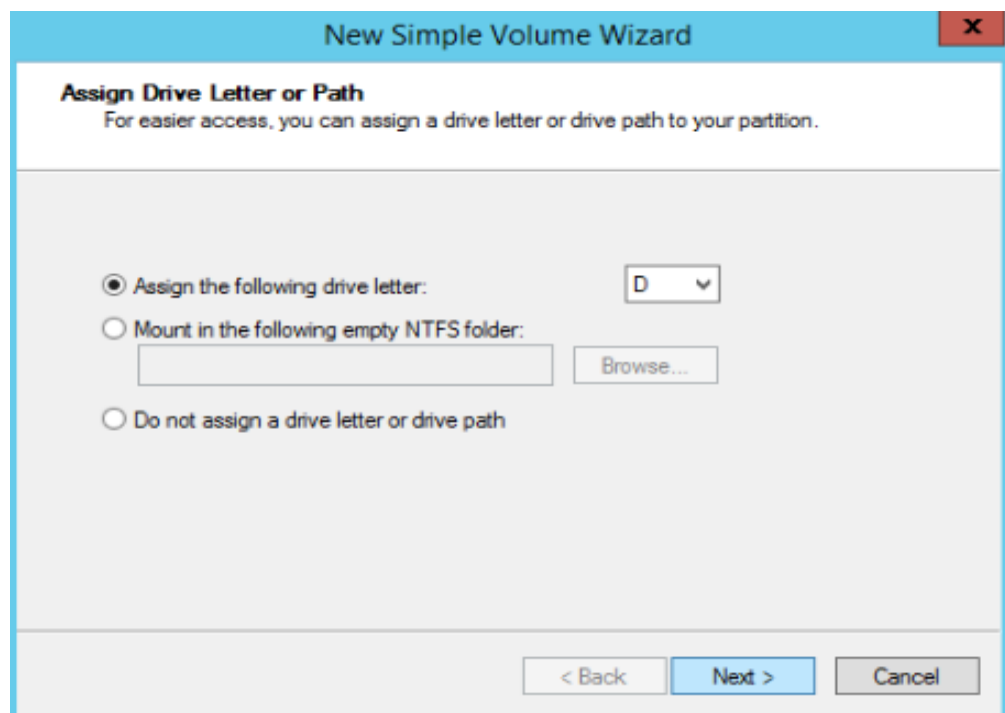
1. On the **Specify Volume Size** page, retain the default settings and click **Next**. The system uses the maximum disk space as the default volume size. You can specify a volume size as needed.

**Figure 6-16** Specify Volume Size



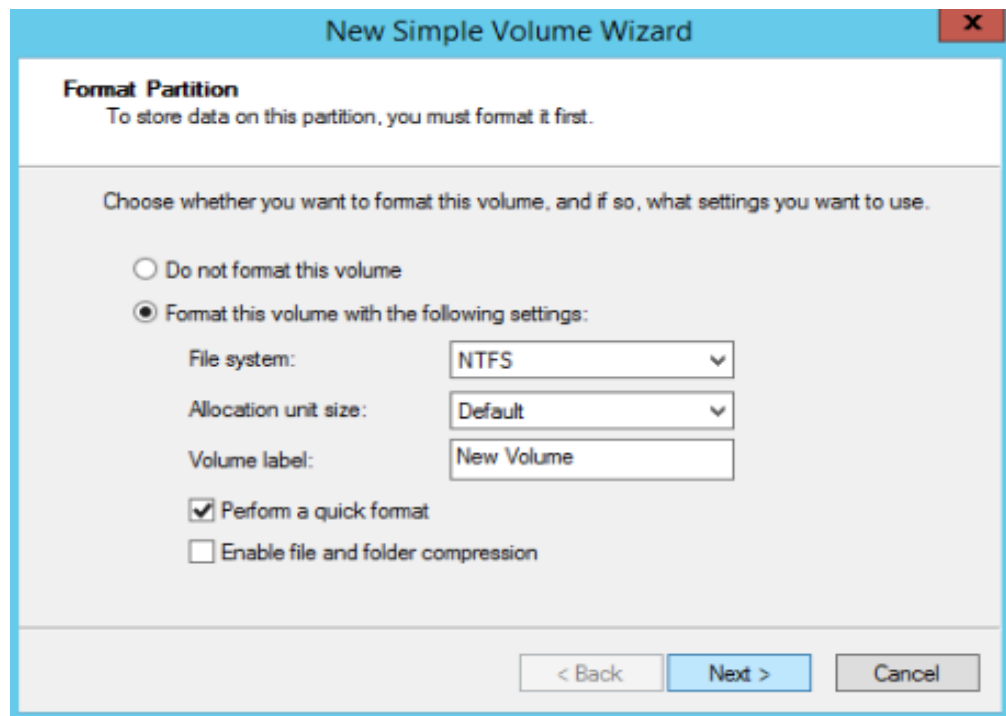
2. On the **Assign Drive Letter or Path** page, retain the default settings and click **Next**.

**Figure 6-17** Assign Drive Letter or Path



3. On the **Format Partition** page, retain the default settings and click **Next**.  
The default file system format is NTFS. You can set other parameters based on your need.

**Figure 6-18** Format Partition



---

**NOTICE**

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

4. On the **Completing the New Simple Volume Wizard** page, click **Finish**.  
Wait for the initialization to complete. When the volume status changes to **Healthy**, the initialization has succeeded.

Figure 6-19 Completing the New Simple Volume Wizard


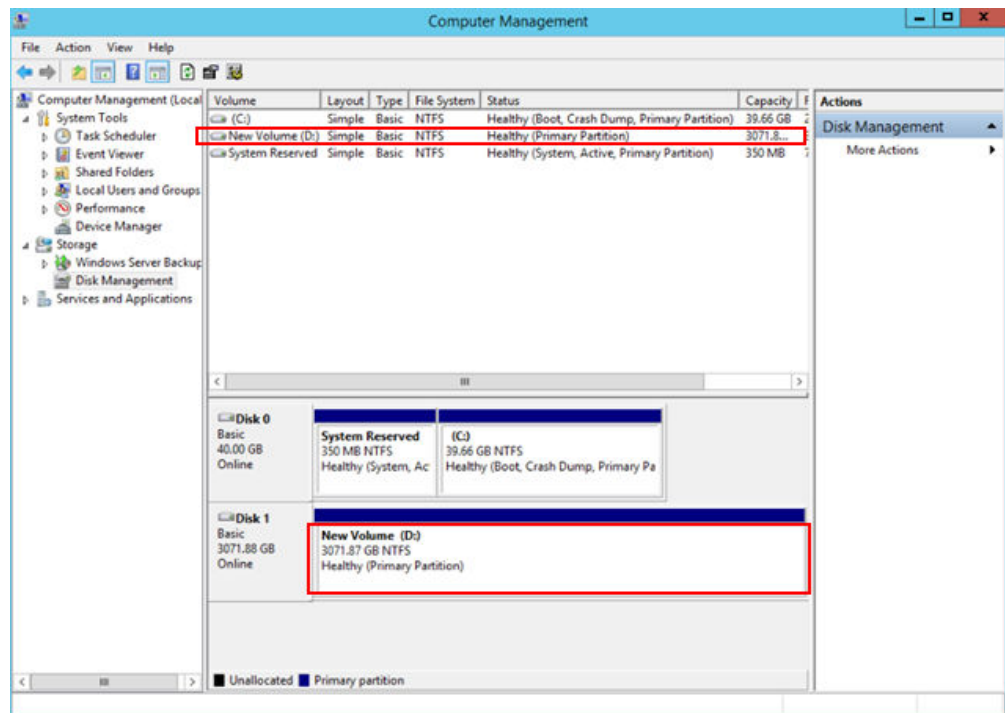



Figure 6-20 Viewing the initialization results

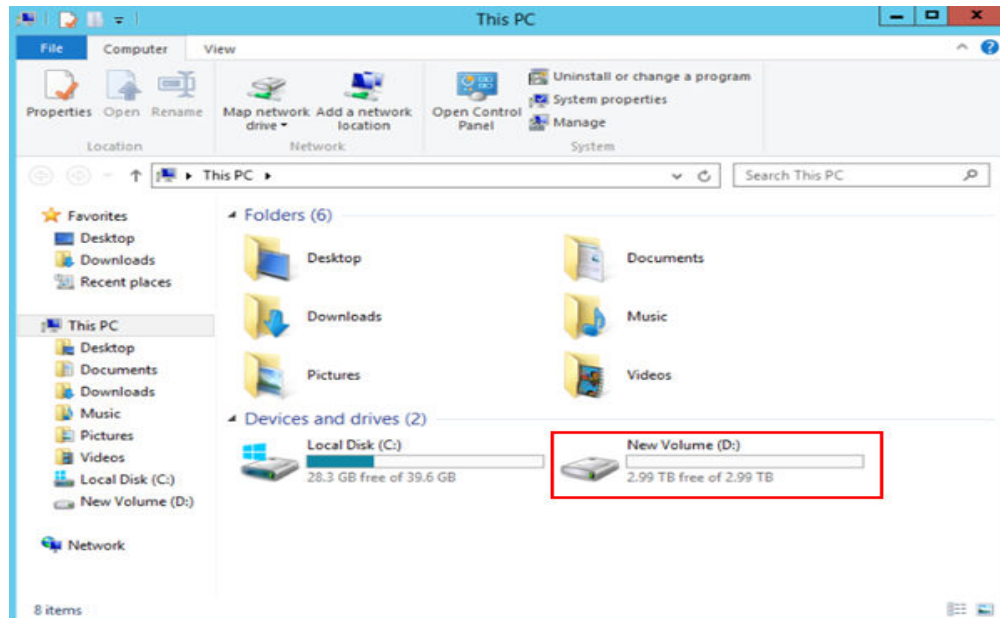


**Step 7** (Optional) Alternatively, choose **Server Manager > File and Storage Services > Volumes > Disks** to view the disk status, capacity, and partition style.

**Step 8** After the volume is created, click  on the task bar and check whether a new volume appears in the File Explorer. 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 6-21** File Explorer



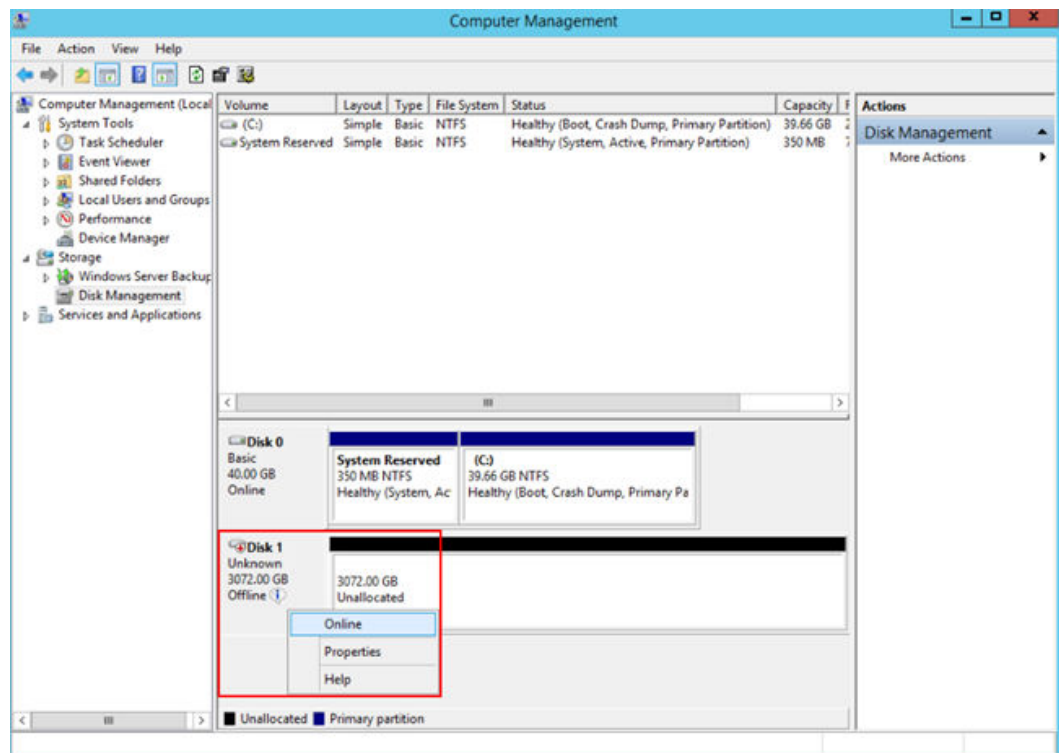
----End

## Related Operations

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

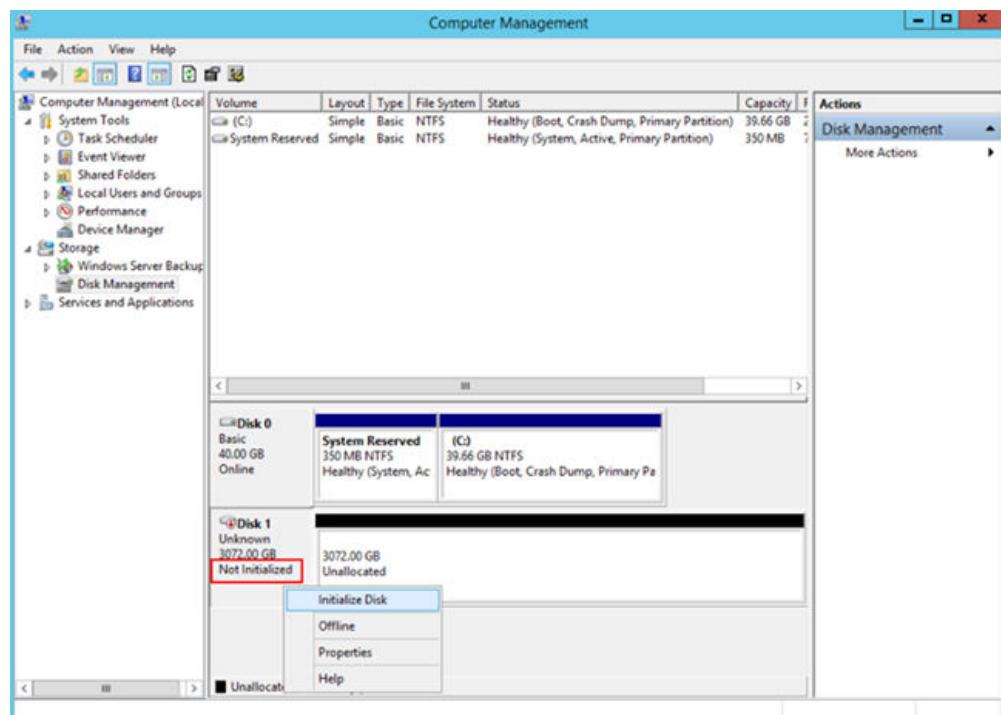
**Step 1** In the **Disk 1** area, right-click and choose **Online** from the shortcut menu.

Figure 6-22 Online



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

Figure 6-23 Disk brought online



----End

# 7 Detaching and Deleting an EVS Disk

## 7.1 Detaching an EVS Disk

### Scenarios

Disk Function	Server Status	Scenarios
System disk	Only offline detachment is supported. You can only detach a system disk when the server status is <b>Stopped</b> .	<ul style="list-style-type: none"><li>• If the file system on your system disk is damaged and the server cannot be started, you can detach the system disk and attach it to another server as a data disk. After the file system is fixed, you can re-attach the disk to the original server as the system disk.</li><li>• If you no longer need a system disk or want to replace it with a new one, you can detach it.</li></ul>
Data disk	Both online detachment and offline detachment are supported. You can detach a data disk when the server status is <b>Stopped</b> or <b>Running</b> .	<ul style="list-style-type: none"><li>• If you want to use a data disk on another server in the same region and AZ, you can detach it and then attach it to that server.</li><li>• If a data disk is no longer required, you can detach it and then delete it.</li></ul>

 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 the system disk is detached, 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 re-attached to a server to be used as a system disk or data disk depending on the disk function selected.
- 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 is detached, the disk function remains unchanged, and the status changes to **Available**. For a shared disk, the status changes to **Available** only after it is detached from all its servers.

## Notes and Constraints

- You can attach SCSI disks to BMSs and use them as data disks.
- After a system disk is detached, some operations cannot be performed on the original server and the system disk. The restrictions are as follows:
  - Server: starting the server, remote login, resetting the password, changing server billing mode, changing server specifications, changing the OS, reinstalling the OS, creating images, creating backups, adding disks, changing the security group, and changing the VPC
- A shared data disk can be detached from ECSs in a batch.

## Prerequisites

- Before detaching an EVS disk from a running Windows server, 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 server, you must log in to the server and run the **umount** command to cancel the association between the disk and the file system, and ensure that no programs are reading data from or writing data to the disk. Otherwise, you will not be able to detach the disk.

## Detaching a System Disk

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

**Step 2** Choose **Compute > 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.

----End

## Detaching a Non-Shared Data Disk

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

**Step 2** Click  in the upper left corner and choose **Storage > Elastic Volume Service**.

The **Elastic Volume Service** page is displayed.

**Step 3** Choose a way to detach the disk by determining whether you want to check the 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 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.

In the disk list, the disk status is **Detaching**, indicating that the disk is being detached from the server.

When the status changes to **Available**, the disk has been detached.

----End

## Detaching a Shared Data Disk

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

**Step 2** Choose **Storage > Elastic Volume Service**.

The **Elastic Volume Service** page is displayed.

**Step 3** Choose a way to detach the disk by determining whether you want to check the 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 disk has been attached.
  - c. Click  to select servers 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:
  - 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  to select servers.  
Shared EVS disks support batch detachment so that you can select multiple servers at a time.
  - c. Click **Yes** to detach the disk.

In the disk list, the disk status is **Detaching**, indicating that the disk is being detached from the server.

If a shared disk has been attached to multiple servers and you only detach it from some of the servers, the disk status will go back to **In-use** after the disk has been detached from the servers. The disk status changes to **Available** only after the disk has been detached from all the servers.

----End

## Helpful Links

To check out more detachment FAQs, see [Detachment](#).

## 7.2 Deleting an EVS Disk

### Scenarios

If an EVS disk is no longer used, you can delete the disk to release the virtual resources. When a disk is deleted, EVS immediately destroys the metadata to ensure 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 requests to the disk.

## Notes and Constraints

- 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.

---

### 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.

---

## Deleting EVS Disks

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

**Step 2** Click  in the upper left corner and choose **Storage > Elastic Volume Service**.  
The **Elastic Volume Service** page is displayed.

**Step 3** In the disk list, locate the target disk and choose **More > Delete** in the **Operation** column.

**Step 4** (Optional) If multiple disks are to be deleted, select  in front of each target disk and click **Delete** in the upper left area of the list.

**Step 5** In the displayed dialog box, confirm the information and click **Yes**.

----End

## Helpful Links

For more deletion FAQs, see [Deletion](#).

# 8 Managing EVS Snapshots

## 8.1 EVS Snapshot Overview

### Overview

An EVS snapshot is a complete copy or image of the disk data taken at a specific time. Snapshot is a major disaster recovery (DR) approach, and you can use a snapshot to restore disk data to the time when the snapshot was created.

**Table 8-1** Snapshot-related operations

Operation	Description	Reference
Creating snapshots	You can create a snapshot to save the disk data at a specified time. <b>NOTE</b> Snapshots are read-only. After snapshots are created, data in the snapshots cannot be modified.	<a href="#">Creating an EVS Snapshot</a>
Rolling back data	If data on a disk is incorrect or damaged, you can roll back data from a snapshot to the source disk.	<a href="#">Rolling Back Disk Data from a Snapshot</a>
Creating disks from a snapshot	You can create disks from a snapshot to quickly copy the snapshot data to disks.	<a href="#">Creating a Disk from a Snapshot</a>
Checking snapshot information	You can check the snapshot details, including the region and AZ, source disk information, and tags.	<a href="#">Checking EVS Snapshot Details</a>
Deleting snapshots	If you no longer require certain snapshots or the snapshot quantity reaches the maximum allowed, you can delete the snapshots.	<a href="#">Deleting an EVS Snapshot</a>

## Snapshot Usage 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 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 data from snapshots. However, you can create disk B using an existing snapshot of system disk A and attach disk B to a properly running server, for example server B. In this case, server B 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 source 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. 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.

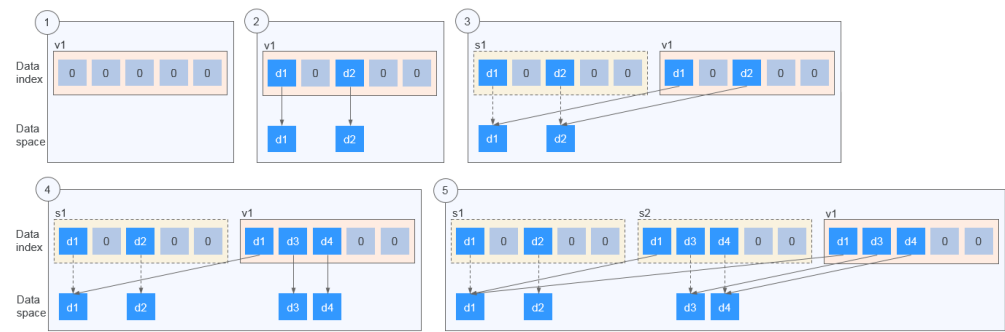
## Snapshot Principles

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. For details, see [Differences Between Disk Backups and Disk Snapshots](#).

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

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

**Figure 8-1** Snapshot principles



## Differences Between Disk Backups and Disk Snapshots

Both disk backups and disk snapshots provide redundancies for improved disk data reliability. [Table 8-2](#) lists the differences between them.

**Table 8-2** Differences between backups and snapshots

Item	Storage Solution	Data Synchronization	DR Range	Service Recovery
Backup	Backups are stored in OBS, instead of disks. This ensures data restoration upon disk damage or corruption.	A backup is a copy of a disk taken at a given time and is stored in a different location. Automatic backup can be performed based on backup policies. Deleting a disk will not delete its backups.	A backup and its source disk reside in different AZs.	You can use a backup to roll back data to its source disk or create a new disk. The data durability is high.

Item	Storage Solution	Data Synchronization	DR Range	Service Recovery
Snapshot	<p>Snapshots are stored on the same disk as the source data.</p> <p><b>NOTE</b> Creating a backup requires a certain amount of time because data needs to be transferred to OBS. Creating a snapshot or rolling back data from a snapshot consumes less time than creating a backup.</p>	<p>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.</p>	<p>A snapshot and its source disk reside in the same AZ.</p>	<p>You can use a snapshot to roll back data to its source disk or create a new disk.</p>

## 8.2 Using EVS Snapshots

### 8.2.1 Creating an EVS Snapshot

#### Scenarios

You can create EVS snapshots to save disk data at specific time points. Before you perform any critical operation, such as a data rollback, software upgrade, or data migration, you are advised to create snapshots to back up data. This ensures that your data is not affected even if an exception occurred during the operation.

 **NOTE**

During the snapshot creation, disk I/Os are affected, so you may experience slow reads or writes at some points. It is recommended that you create snapshots at off-peak hours.

#### Prerequisites


Snapshots can only be created for **Available** or **In-use** disks.

## Notes and Constraints

- Snapshots can be created for both system disks and data disks.
- Snapshots of encrypted disks are stored encrypted, and those of non-encrypted disks are stored non-encrypted.
- You can manually create a maximum of seven snapshots for a disk.
- The enterprise project of a snapshot is the same as that of the snapshot's source disk.

## Creating a Snapshot on the Disks Page

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

**Step 2** Click  in the upper left corner and choose **Storage > Elastic Volume Service**.  
The **Elastic Volume Service** page is displayed.

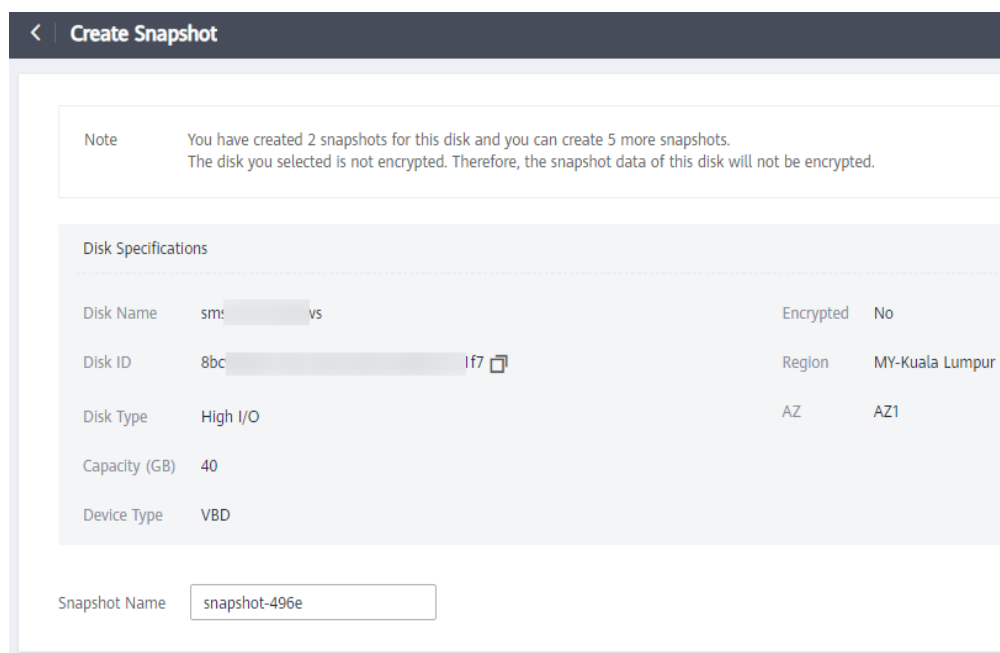
**Step 3** In the disk list, locate the target disk and click **Create Snapshot** in the **Operation** column.

Configure the snapshot parameter according to [Table 8-3](#).

**Table 8-3** Snapshot parameter

Parameter	Description	Example Value
Snapshot Name	Mandatory The name can contain a maximum of 64 characters.	snapshot-01

**Figure 8-2** 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

## Creating a Snapshot on the Snapshots Page

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

**Step 2** Click  in the upper left corner and choose **Storage > Elastic Volume Service**.

The **Elastic Volume Service** page is displayed.

**Step 3** In the navigation pane on the left, choose **Elastic Volume Service > Snapshots**.

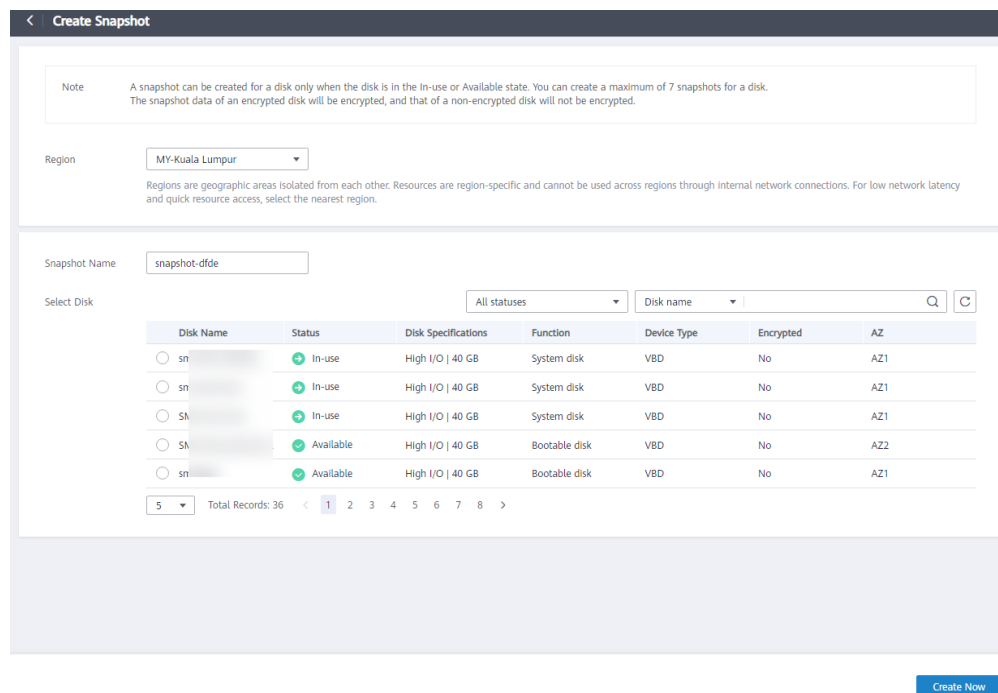
On the **Snapshots** page, click **Create Snapshot**.

Configure the snapshot parameters according to [Table 8-4](#).

**Table 8-4** 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 8-3** 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

## 8.2.2 Rolling Back Disk 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.

### Notes and Constraints

- Snapshot data can only be rolled back to source EVS disks. Rollback to a different disk is not possible.
- You can only roll back disk data from a snapshot when the source disk status is **Available** (not attached to any server) or **Rollback failed**. If the source disk is attached, detach the disk first.
- If a snapshot is being created, it cannot be used to roll back disk data.
- A snapshot whose name starts with **autobk\_snapshot\_vbs\_**, **manualbk\_snapshot\_vbs\_**, **autobk\_snapshot\_csbs\_**, or **manualbk\_snapshot\_csbs\_** is automatically generated during backup. Such a snapshot can only be viewed. It cannot be used to roll back the disk data.

## Rolling Back Disk Data from a Snapshot

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

**Step 2** Click  in the upper left corner and choose **Storage > Elastic Volume Service**.

The **Elastic Volume Service** page is displayed.

**Step 3** In the navigation pane on the left, choose **Elastic Volume Service > Snapshots**.

The **Snapshots** page is displayed.

**Step 4** In the snapshot list, locate 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

## 8.2.3 Creating a Disk from a Snapshot

### Scenarios


This section describes how to create an EVS disk on the **Snapshots** page. 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 [Creating an EVS Disk](#).

### Notes and Constraints

- Batch disk creation from a snapshot is not supported.
- A disk created from a snapshot has the same device type (SCSI or VBD), encryption attribute, AZ, region, and disk type as the snapshot's source disk.
- A snapshot whose name starts with **autobk\_snapshot\_vbs\_**, **manualbk\_snapshot\_vbs\_**, **autobk\_snapshot\_csbs\_**, or **manualbk\_snapshot\_csbs\_** is automatically generated during backup. Such a snapshot can only be viewed. It cannot be used to create new disks.

### Creating an EVS Disk from a Snapshot

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

**Step 2** Click  in the upper left corner and choose **Storage > Elastic Volume Service**.

The **Elastic Volume Service** page is displayed.

**Step 3** In the navigation pane on the left, choose **Elastic Volume Service > Snapshots**.

The **Snapshots** page is displayed.

**Step 4** In the snapshot list, locate 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 [Creating an EVS Disk](#).

 **NOTE**

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 a 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** Confirm the configuration and click **Submit**.

**Step 8** In the disk list, view the disk status.

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

----End

## 8.2.4 Checking EVS Snapshot Details

### Scenarios

You can check the snapshot details, including the region and AZ, source disk information, and tags.

### Checking Snapshot Details

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

**Step 2** Choose **Storage > Elastic Volume Service**.

**Step 3** In the navigation pane on the left, choose **Elastic Volume Service > Snapshots**.

The **Snapshots** page is displayed.

**Step 4** In the snapshot list, locate the target snapshot.

Select the snapshot and check the snapshot details at the bottom of the page.

----End

### Snapshot Statuses

An EVS snapshot has several statuses. [Table 8-5](#) lists the EVS snapshot statuses, the meaning of each status, and the operations a snapshot in each status allows.

**Table 8-5** 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.	<ul style="list-style-type: none"><li>• Creating EVS disks using snapshots</li><li>• Deleting snapshots</li><li>• Rolling back data to EVS disks using snapshots</li></ul>
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. <b>NOTE</b> <ul style="list-style-type: none"><li>• When you roll back from a snapshot, you can only roll back data to the source EVS disk. Rollback to a specified disk is not supported.</li><li>• A snapshot can only be used for rollback when its source disk is in the <b>Available</b> or <b>Rollback failed</b> state.</li></ul>	No operations are allowed.

## 8.2.5 Deleting an EVS Snapshot

### Scenarios

If you no longer require certain snapshots or the snapshot quantity reaches the maximum allowed, you can delete the snapshots.

### Prerequisites

- The snapshot status must be **Available** or **Error**.


### Notes and Constraints

- If a snapshot is deleted, disks rolled back or created from this snapshot are not affected.
- If a snapshot's source disk is deleted, all snapshots of this disk will also be deleted.
- If you reinstall or change the server OS, snapshots of the system disk are automatically deleted. Those of the data disks can be used as usual.
- A snapshot whose name starts with **autobk\_snapshot\_vbs\_**, **manualbk\_snapshot\_vbs\_**, **autobk\_snapshot\_csbs\_**, or

**manualbk\_snapshot\_csbs\_** is automatically generated during backup. You can only check details of such snapshots and cannot delete them.

## Procedure

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

**Step 2** Click  in the upper left corner and choose **Storage > Elastic Volume Service**.  
The **Elastic Volume Service** page is displayed.

**Step 3** In the navigation pane on the left, choose **Elastic Volume Service > Snapshots**.  
The **Snapshots** page is displayed.

**Step 4** In the snapshot list, locate the target snapshot and click **Delete** in the **Operation** column.

**Step 5** 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

# 9 Managing Encrypted EVS Disks

---

## What Is EVS Disk Encryption?

EVS enables you to encrypt data on newly created disks as required. Keys used to encrypt EVS disks are provided by the Key Management Service (KMS) of Data Encryption Workshop (DEW), which is secure and convenient. You do not need to establish and maintain the key management infrastructure.

---

### 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 [Creating an EVS Disk](#).

---

## Keys Used for EVS Encryption

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

When an encrypted disk is attached, EVS accesses KMS, and KMS sends the data key (DK) to the host memory for use. The disk uses the DK plaintext to encrypt and decrypt disk I/Os. The DK plaintext is only stored in the memory of the host housing the ECS and is not stored persistently on the media. If a 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 can no longer 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 9-1](#) for more information.

**Table 9-1** Impact of CMK unavailability

CMK Status	Impact	How to Restore
Disabled	<ul style="list-style-type: none"> <li>For an encrypted disk already attached: Reads and writes to the disk are normal. If the disk is detached, it cannot be attached again.</li> </ul>	Enable the CMK. For details, see "Managing CMKs" > "Enabling One or More CMKs" in the <i>Data Encryption Workshop User Guide</i> .
Scheduled deletion	<ul style="list-style-type: none"> <li>For an encrypted disk not attached: The disk cannot be attached anymore.</li> </ul>	Cancel the scheduled deletion for the CMK. For details, see "Managing CMKs" > "Canceling the Scheduled Deletion of One or More CMKs" in the <i>Data Encryption Workshop User Guide</i> .
Deleted		Data on the disks can never be restored.

## 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 you select when creating the server.

**Table 9-2** Relationship between encrypted images and system disk encryption

Creating Server Using Encrypted Image	Whether System Disk Will Be Encrypted	Description
Yes	Yes	For details, see "Managing Private Images" > "Encrypting Images" in the <i>Image Management Service User Guide</i> .
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 9-3** Relationship between backups, snapshots, images, and data disk encryption

Disk Created On	Method of Creation	Whether Data Disk Will Be Encrypted	Description
ECS 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 <b>Getting Started &gt; Creating an ECS &gt; Step 1: Configure Basic Settings</b> in the <i>Elastic Cloud Server User Guide</i> .
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	<ul style="list-style-type: none"> <li>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.</li> <li>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.</li> </ul>
	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.

Disk Created On	Method of Creation	Whether 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	-

## Notes and Constraints

**Table 9-4** Constraints on disk encryption

Item	Description
Types of disks supporting encryption	All disk types
Constraints on encrypted disks	<p>The encryption attribute of a disk cannot be changed after the disk is created, meaning that:</p> <ul style="list-style-type: none"> <li>• An encrypted disk cannot be changed to a non-encrypted disk.</li> <li>• A non-encrypted disk cannot be changed to an encrypted disk.</li> </ul>
Constraints on user permissions	<p>When a user uses the encryption function, the condition varies depending on whether the user is the first one ever in the current region or project to use this function.</p> <ul style="list-style-type: none"> <li>• If the user is the first user, the user needs to follow the prompt to create an agency, which grants EVS Administrator permissions to EVS. Then, the user can create and obtain keys to encrypt and decrypt disks.</li> <li>• If the user is not the first user, the user can use encryption directly.</li> </ul>

Item	Description
Constraints on encrypted images	<ul style="list-style-type: none"><li>• Encrypted images cannot be replicated across regions.</li><li>• Encrypted images cannot be changed to non-encrypted images.</li><li>• Encrypted images cannot be exported.</li></ul>

## 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 [Creating 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 be impossible 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 [Keys Used for EVS Encryption](#).

For details about how to detach an encrypted disk, see [Detaching an EVS Disk](#).

---

# 10 Managing Shared EVS Disks

---

## What Is Disk Sharing?

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


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

---

### NOTICE

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

---

**Figure 10-1** Application scenario of shared EVS disks

## Advantages

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

## Specifications and Performance

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

## How Do I 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 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 reads/writes to the disk from multiple ECSs.
- ECS groups and SCSI reservations have the following relationship: A SCSI reservation on a single EVS disk cannot differentiate multiple ECSs on the same physical host. For that reason, if multiple ECSs that use the same shared EVS disk are running on the same physical host, SCSI reservations will not work properly. So you are advised to use SCSI reservations only on ECSs that are in the same ECS group, thus having a working anti-affinity policy.

## Constraints on Shared Disks

- A shared disk can be attached to a maximum of 16 servers.
- The sharing attribute of a disk cannot be changed after the disk is created.
- Shared disks can only be used as data disks, not system disks.
- A shared file system or cluster management system must be set up before you can properly use a shared disk. If you simply attach a shared disk to multiple servers, data cannot be shared between those servers and may be overwritten.
- When a shared disk is attached to multiple servers, the total performance of the disk on all servers cannot exceed the maximum allowed on a single disk.

## 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 an EVS Disk](#).

## Data Sharing Principles and Common Usage Mistakes

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

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

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

- Data inconsistency caused by read/write conflicts

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

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

- Data inconsistency caused by data caching

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

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

Before you buy a shared EVS disk, determine its device type (VBD or SCSI) based on the applications that will use the shared disk. Shared SCSI EVS disks support SCSI reservations. Before using SCSI reservations, you need to install a driver in the server OS and ensure that the OS image is included in the compatibility list.

**NOTICE**

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

---

## Helpful Links

For more disk sharing FAQs, see [Sharing](#).



# 11 Managing EVS Disk Backups

---

## 11.1 CBR Overview

### What Is CBR?

Cloud Backup and Recovery (CBR) enables you to easily back up cloud servers and cloud disks. In case of a virus attack, accidental deletion, or software or hardware fault, you can restore data to any point when the data was backed up.

CBR protects your workloads by ensuring the security and consistency of your data.

### 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 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.
- Cloud disk backup: provides snapshot-based backups for EVS disks.

- **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.

Vaults can be either backup vaults or replication vaults. Backup vaults store resource backups, and replication vaults store backup replicas.

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

- **Policy**

There are backup policies and replication policies.

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

## Backup Mechanism

The first backup is a full backup and backs up all used data blocks.

For example, if a disk size is 100 GB and 40 GB has been used, only the 40 GB of data is backed up.

Subsequent backups are incremental backups. An incremental backup backs up only the data changed since the last backup to save the storage space and backup time.

When a backup is deleted, data blocks will not be deleted if they are depended on by other backups, ensuring that other backups can still be used for restoration. Both a full backup and an incremental backup can be used to restore data to a given backup point in time.

When creating a backup of a disk, CBR also creates a snapshot for it. CBR keeps only the latest snapshot. Every time it creates a new snapshot during backup, it deletes the old snapshot.

CBR stores backups in OBS to ensure data security.

## Backup Options

CBR supports one-off backup and periodic backup. A one-off backup task is manually created by users and is executed only once. Periodic backup tasks are automatically executed based on a user-defined backup policy.

**Table 11-1** One-off backup and periodic backup

Item	One-Off Backup	Periodic Backup
Backup policy	Not required	Required
Number of backup tasks	One manual backup task	Periodic tasks driven by a backup policy
Backup name	User-defined backup name, which is <b>manualbk_XXXX</b> by default	System-assigned backup name, which is <b>autobk_XXXX</b> by default
Backup mode	Full backup for the first time and incremental backup subsequently, by default	Full backup for the first time and incremental backup subsequently, by default

Item	One-Off Backup	Periodic Backup
Application scenario	Executed before patching or upgrading the OS or upgrading an application on a resource. A one-off backup can be used to restore the resource to the original state if the patching or upgrading fails.	Executed for routine maintenance of a resource. The latest backup can be used for restoration if an unexpected failure or data loss occurs.

## 11.2 Backing Up EVS Disks

### Scenarios

EVS disk backups are created using the CBR service. For details, see section "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.

### Notes and Constraints

- Backups can be created only when the disks are in the **Available** or **In-use** state.
- Backup data can only be restored to original disks.
- Only users with the CBR FullAccess permissions can use the cloud disk backup function. If the user does not have the permissions, contact the account administrator to grant the permissions first.

### Creating a Disk Backup Vault and Applying a Backup Policy

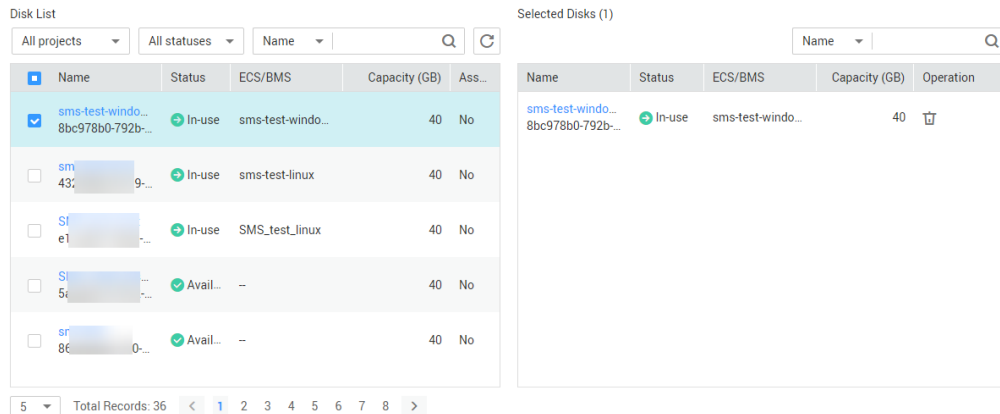
**Step 1** Log in to the CBR console.

1. Log in to the console.
2. Click  in the upper left corner of the page and 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 11-1** Selecting disks



**NOTE**

- Only **Available** and **In-use** disks can be selected.
- You can also associate disks with the vault you are creating 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 policy. After the vault is created, CBR will apply the policy to this 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 makes it easy to manage projects and groups of cloud resources and users. Use the **default** enterprise project or create one.

**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 cloud disk backup page and view the created vault in the vault list.

You can associate disks to the new vault or create backups for the disks. For details, see section "Vault Management" in the *Cloud Backup and Recovery User Guide*.

----End

# 12 Managing EVS Transfers

---

## Scenarios

EVS transfer allows you to transfer disks from one account to another. After a transfer succeeds, the ownership of the disk belongs 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*.

## Notes and Constraints

- Encrypted EVS disks cannot be transferred.
- 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.

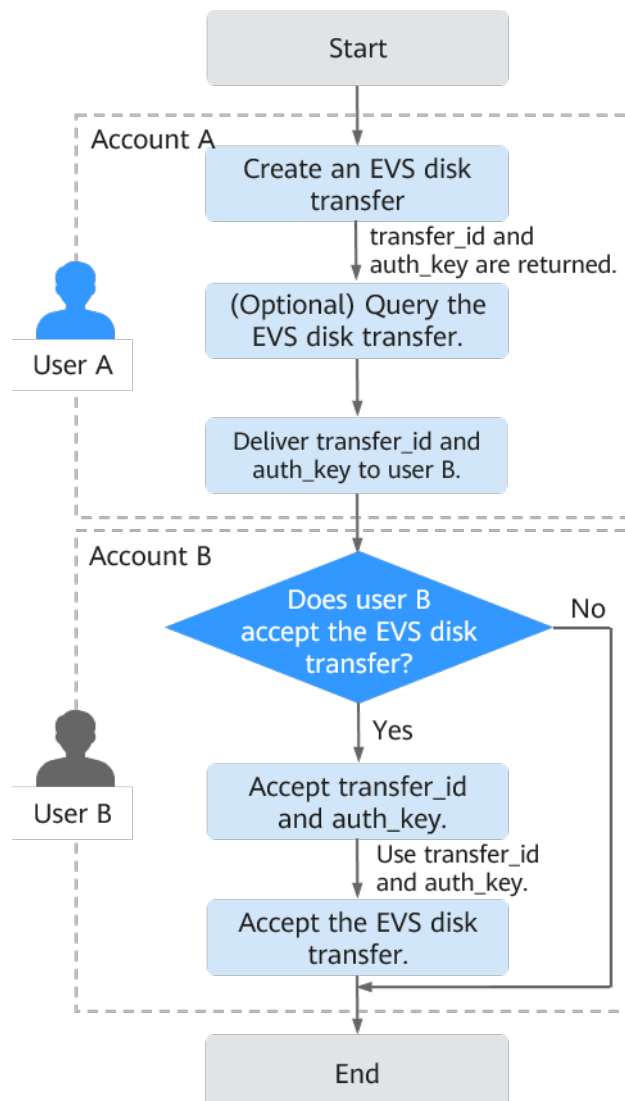
## Procedure

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 12-1](#) 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 12-1** 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

# 13 Managing EVS Tags

## 13.1 Tag Overview

Tags identify EVS resources for purposes of easy categorization and quick search.

If your organization has enabled the tag policy type for EVS and has a tag policy attached, you must comply with the tag policy rules when creating disks, otherwise disks may fail to be created. Contact the organization administrator to learn more about tag policies.

**Table 13-1** Tag overview

Operation	Scenario
<a href="#">Adding a Tag</a>	Add tags for existing disks or during disk creations.
<a href="#">Modifying a Tag</a>	Change tag values for existing disks. Tag keys of existing disks cannot be changed.
<a href="#">Deleting a Tag</a>	Delete tags that are no longer needed for existing disks.
<a href="#">Searching for Disks by Tag</a>	After tags are added, search for disks by tags.

## 13.2 Adding a Tag

### Scenarios

You can add tags for an existing EVS disk. You can also add tags when creating a disk. For details, see [Creating an EVS Disk](#).

### Tag Rules

A tag consists of a tag key and a tag value. Tag rules are described as follows: (Tag rules vary depending on regions. See the rules displayed on the console.)



First set of rules:

- A tag key can contain a maximum of 36 characters. It can contain only letters, digits, special characters (.-\_), and Unicode characters.
- A tag value can contain a maximum of 43 characters. It can contain only letters, digits, special characters (.-\_), and Unicode characters.

Second set of rules:

- A tag key can contain a maximum of 36 characters. It cannot contain special characters (=\*<>\\,|/) or start or end with spaces.
- A tag value can contain a maximum of 43 characters. It cannot contain special characters (=\*<>\\,|/) or start or end with spaces.

Third set of rules:

- A tag key can contain a maximum of 128 characters. It cannot contain special characters (\*<>\\,|), start with **\_sys\_**, or start or end with spaces.
- A tag value can contain a maximum of 255 characters. It cannot contain special characters (\*<>\\,|) or start or end with spaces.

## Notes and Constraints

- You can add a maximum of 10 tags for a single EVS disk.
- Tag keys of the same EVS disk must be unique.

## Procedure

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

**Step 2** Click  in the upper left corner and choose **Storage > Elastic Volume Service**.

The **Elastic Volume Service** 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**.

- Tag key: This parameter is mandatory.
- Tag value: This parameter is optional.

The **Tags** tab is displayed, and you can view the newly added tag.

----End

## 13.3 Modifying a Tag

### Scenarios

You can change the value of a tag for an existing disk, but cannot change the key of a tag.

### Tag Rules

A tag consists of a tag key and a tag value. Tag rules are described as follows: (Tag rules vary depending on regions. See the rules displayed on the console.)

First set of rules:

- A tag key can contain a maximum of 36 characters. It can contain only letters, digits, special characters (.-\_), and Unicode characters.
- A tag value can contain a maximum of 43 characters. It can contain only letters, digits, special characters (.-\_), and Unicode characters.

Second set of rules:

- A tag key can contain a maximum of 36 characters. It cannot contain special characters (=\*<>\\,|/) or start or end with spaces.
- A tag value can contain a maximum of 43 characters. It cannot contain special characters (=\*<>\\,|/) or start or end with spaces.

Third set of rules:

- A tag key can contain a maximum of 128 characters. It cannot contain special characters (\*<>\\,|), start with `_sys_`, or start or end with spaces.
- A tag value can contain a maximum of 255 characters. It cannot contain special characters (\*<>\\,|) or start or end with spaces.

### Notes and Constraints

- You can add a maximum of 10 tags for a single EVS disk.
- Tag keys of the same EVS disk must be unique.

### Procedure

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

**Step 2** Choose **Storage > Elastic Volume Service**.

The **Elastic Volume Service** 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


## 13.4 Deleting a Tag

### Scenarios

If an existing tag is no longer needed, you can delete it.

### Procedure

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

**Step 2** Click  in the upper left corner and choose **Storage > Elastic Volume Service**.

The **Elastic Volume Service** 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 **OK**.

The **Tags** tab is displayed, and the deletion is complete.

----End


## 13.5 Searching for Disks by Tag

### 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 console.


**Step 2** Click  in the upper left corner and choose **Storage > Elastic Volume Service**.

The **Elastic Volume Service** 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.

You can add a maximum of 10 tags to search for disks. If you add more than one tag, only the disks containing all specified tags will be returned.

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**

# 14 Managing EVS Quotas

---



## 14.1 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 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

## 14.2 Increasing EVS Resource Quotas

### Scenarios

If any resource quota no longer meets your service requirements, you can apply for a higher quota.

## How Do I Apply for a Higher Quota?

1. Log in to the management console.
2. In the upper right corner of the page, choose **Resources > My Quotas**.  
The **Service Quota** page is displayed.
3. Click **Increase Quota** in the upper right corner of the page.
4. On the **Create Service Ticket** page, configure parameters as required.  
In the **Problem Description** area, fill in the content and reason for adjustment.
5. After all necessary parameters are configured, select **I have read and agree to the Ticket Service Protocol and Privacy Statement** and click **Submit**.

# 15 Cloud Eye Monitoring

## 15.1 Viewing Basic 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 15-1 EVS metrics

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period
disk_device_read_bytes_rate	Disk Read Bandwidth	Number of bytes read from the monitored disk per second Unit: Bytes/s	$\geq 0$ bytes/s	EVS disk	5 minutes in average
disk_device_write_bytes_rate	Disk Write Bandwidth	Number of bytes written to the monitored disk per second Unit: Bytes/s	$\geq 0$ bytes/s	EVS disk	5 minutes in average

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period
disk_device_read_requests_rate	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
disk_device_write_requests_rate	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	<i>Server ID-drive letter</i> , for example, <b>6f3c6f91-4b24-4e1b-b7d1-a94ac1cb011d-vda</b> (vda is the drive letter) <i>Server ID-volume- Volume ID</i> , for example, <b>6f3c6f91-4b24-4e1b-b7d1-a94ac1cb011d-volume-31f45764-38b3-44ad-aaca-4015c83371e6</b>

## Viewing Monitoring Data

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

**Step 2** Choose **Storage > Elastic Volume Service**.

The **Elastic Volume Service** 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



## 15.2 Viewing EVS Monitoring Data Included in OS Metrics (with Agent Installed)

### 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 15-2 EVS-related metrics

Metric	Name	Description	Value Range	Monitored Object	Monitoring Period
mountPointPrefix_disk_free	(Agent) Available Disk Space	Available disk space on the monitored object Unit: GiB <ul style="list-style-type: none"><li>Linux: Run the <b>df -h</b> command to check the value in the <b>Avail</b> 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 (~).</li><li>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 (~).</li></ul>	≥0 GiB	ECS	5 minutes in average

Metric	Name	Description	Value Range	Monitored Object	Monitoring Period
mountPointPrefix_disk_usedPercent	(Agent) Disk Usage	<p>Percentage of total disk space that is used, which is calculated as follows: <b>Disk Usage = Used Disk Space / Disk Capacity.</b></p> <p>Unit: Percent</p> <ul style="list-style-type: none"> <li>Linux: Obtain the metric value using the following formula: <b>Disk Usage = Used Disk Space / Disk Capacity.</b> 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 (~).</li> <li>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 (~).</li> </ul>	0-100%	ECS	5 minutes in average

Metric	Name	Description	Value Range	Monitored Object	Monitoring Period
mountPointPrefix_disk_ioUtils and volumePrefix_disk_ioUtils	(Agent) Disk I/O Usage	<p>Percentage of the time that the disk has had I/O requests queued to the total disk operation time</p> <p>Unit: Percent</p> <ul style="list-style-type: none"> <li>Linux: The disk I/O usage is calculated based on the data changes in the thirteenth column of the corresponding device in file <b>/proc/diskstats</b> in a collection period. <p>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 (~).</p> </li> <li>Windows does not support this metric.</li> </ul>	0-100%	ECS	5 minutes in average

## Dimensions

Key	Value
instance_id	Specifies the ECS ID.

# 16 Recording EVS Operations Using CTS

---

## 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 16-1** 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 section "Querying Real-Time Traces" in the *Cloud Trace Service User Guide*.

# 17 FAQ

---

## 17.1 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?](#)
- [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 a Capacity Expansion?](#)

- [Why 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 a Different 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 the Available Capacity of My Disk?](#)
- [Can I Transfer the Data Disk Capacity to a System Disk?](#)

## Performance

- [How Do 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 or Increased I/Os?](#)

## 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 Data from a Snapshot After Reinstalling the OS or Formatting a 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 Disk Backup?](#)
- [Can I Back Up and Restore My EVS Disk to a Different Region?](#)

## 17.2 General

### 17.2.1 How Do I Start Using a Newly Created Disk?

A newly created disk must be attached to a server and then initialized in the server OS before you can use it.

For more information, see [Getting Started](#).

### 17.2.2 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.

### 17.2.3 Can EVS Disks Be Used Alone?

No.

EVS disks must be attached to servers before you can use them.

### 17.2.4 How Can I View My Disk Details?

To do so, perform the following operations:

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

**Step 2** Under **Storage**, click **Elastic Volume Service**.

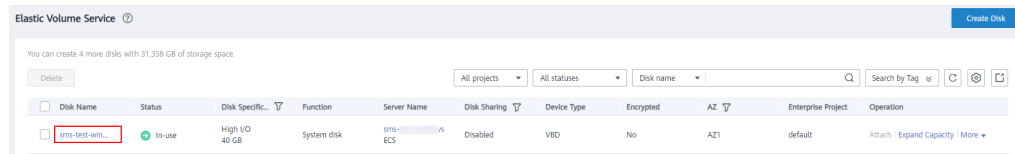
The disk list page is displayed.

**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 17-1** Viewing disk details



Disk Name	Status	Disk Specifics	Function	Server Name	Disk Sharing	Device Type	Encrypted	AZ	Enterprise Project	Operation
sms-test-will	In-use	High I/O 40 GB	System disk	ecs-*	Disabled	VBD	No	AZ1	default	Attach Expand Capacity More

----End

## 17.2.5 How Do I Change the Name of My Disk?

### 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.




### Procedure

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

**Step 2** Choose **Storage > 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 **Name/ID** 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  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



## 17.2.6 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.

## 17.2.7 Can I Change the Disk Type, Device Type, or Sharing Attribute of My Disk?

The following table describes whether the disk type, device type, encryption, and sharing attributes of a disk can be changed.


**Table 17-1** EVS disk change description

Attribute	Allow Change	Change Direction
Disk type	No	The EVS disk type cannot be changed.
Sharing	No	<ul style="list-style-type: none"> <li>From shared to non-shared</li> <li>From non-shared to shared</li> </ul>
Device type	No	<ul style="list-style-type: none"> <li>From SCSI to VBD</li> <li>From VBD to SCSI</li> </ul>
Encryption	No	<ul style="list-style-type: none"> <li>From non-encrypted to encrypted</li> <li>From encrypted to non-encrypted</li> </ul>

**However, you can:**

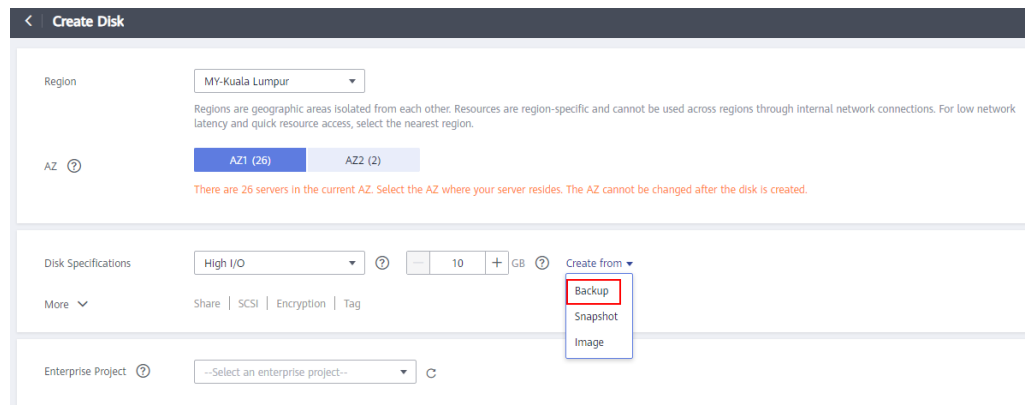
1. Create a backup for the disk.

**Figure 17-2** 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 17-3** Create from backup



## 17.2.8 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 17-2** Solutions for disk errors

Error Status	Handling Suggestion
Error	Delete the disk in the <b>Error</b> 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.
Restoration 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.

## 17.2.9 How Can I Obtain ECS NIC Information?

### Scenarios

Obtain the ECS NIC information, such as the ECS IP address.

### Procedure

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

**Step 2** Choose **Compute > 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 **Network Interfaces** tab.

The network interfaces page is displayed.

**Step 5** On the **Network Interfaces** tab, click  to view information about an ECS network interface.

**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 network interface information.

----End

## 17.2.10 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 17-4** shows the query result.

**Figure 17-4** Queried WWN information

WWN:  6888603000038430fa17a17502223655

- If the SCSI EVS disk was created before the WWN feature rollout, the disk WWN will fail to be obtained.

**Figure 17-5** shows the query result.

**Figure 17-5** No WWN information

WWN:  --

## 17.2.11 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*.

## 17.2.12 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.

## 17.2.13 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**
2. View the data disks previously attached to Windows.  
**parted -l**
3. Mount the data disks.  
**mount -t ntfs-3g *Data disk path* *Mount point***

## 17.2.14 What Are the Differences Between MBR and GPT Partition Styles?

**Table 17-3** lists the common disk partition styles. In Linux, different partition styles require different partitioning tools.

**Table 17-3** Disk partition styles

Disk Partition Style	Maximum Disk Capacity Supported	Maximum Number of Partitions Supported	Linux Partitioning Tool
Master Boot Record (MBR)	2 TiB	<ul style="list-style-type: none"> <li>• 4 primary partitions</li> <li>• 3 primary partitions and 1 extended partition</li> </ul> <p>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:</p> <ul style="list-style-type: none"> <li>• 3 primary partitions and 1 extended partition, with the extended partition divided into 3 logical partitions</li> <li>• 1 primary partition and 1 extended partition, with the extended partition divided into 5 logical partitions</li> </ul>	<p>You can use either of the following tools:</p> <ul style="list-style-type: none"> <li>• fdisk</li> <li>• parted</li> </ul>
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 greater 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.

### 17.2.15 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.

## 17.2.16 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.

## 17.2.17 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 "How Do I Upload Files to My ECS?" in the *Elastic Cloud Server User Guide*.

## 17.2.18 What EVS Disk Types Are Available?

EVS disk types are classified based on I/O performance. The following table describes the details of each type. EVS disks differ in performance and price. You can choose whichever disk type that is the best fit for your applications by referring to [Disk Types and Performance](#).

**Table 17-4** EVS performance data

Parameter	High I/O	Ultra-high I/O
IOPS per GiB/EVS disk	6	50
Max. IOPS/EVS disk	5,000	33,000
Baseline IOPS/EVS disk	1,200	1,500
Disk IOPS	Min. [5,000, 1,200 + 6 x Capacity (GiB)]	Min. [33,000, 1,500 + 50 x Capacity (GiB)]
IOPS burst limit/EVS disk	5,000	16,000
Max. throughput	150 MiB/s	350 MiB/s
API name	SAS	SSD
<b>NOTE</b> This API name is the value of the <b>volume_type</b> parameter in the EVS API. It does not represent the type of the underlying hardware device.		

Parameter	High I/O	Ultra-high I/O
Typical scenarios	Mainstream applications requiring high performance and high reliability, such as large-scale development and test environments, web server logs, and enterprise applications. Typical enterprise applications include SAP applications, Microsoft Exchange, and Microsoft SharePoint.	Read/write-intensive workloads that demand 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.

## 17.3 Attachment

### 17.3.1 Why Can't I View the Attached Data Disk on the Server?

#### Troubleshooting

Table 17-5 Possible causes

OS	Possible Cause	Solution
Linux	<ul style="list-style-type: none"> <li>New data disks are not formatted and partitioned by default. An unformatted disk will not be listed in the command output. You must manually initialize the disk.</li> <li>If a data disk cannot be found after the server is restarted, automatic partition mounting at system start may not be configured.</li> </ul>	<a href="#">Linux Data Disk</a>
Windows	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.	<a href="#">Windows Data Disk</a>

#### Linux Data Disk

**Symptom:** A data disk has been attached to a Linux server on the 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
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/vda1      ext4      43G  1.9G  39G   5% /
devtmpfs       devtmpfs  2.0G    0  2.0G   0% /dev
tmpfs          tmpfs     2.0G    0  2.0G   0% /dev/shm
tmpfs          tmpfs     2.0G  9.1M  2.0G   1% /run
tmpfs          tmpfs     2.0G    0  2.0G   0% /sys/fs/cgroup
tmpfs          tmpfs     398M    0  398M   0% /run/user/0
/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 about how to initialize data disks, see [Initializing EVS Data Disks](#).

- **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. Mount the data disk.

**mount** *Disk partition Mount point*

In this example, run the following command:

```
mount /dev/vdb1 /mnt/sdc
```

Perform the following steps to configure auto mount at system start:

- b. 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"
```

The UUID of the **/dev/vdb1** partition is displayed.

- c. 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
```

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 **b**.
- **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**: The fsck option is not used.
    - 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.
- e. Press **Esc**, enter **:wq**, and press **Enter**.

The system saves the configurations and exits the vi editor.

Verify that the disk is auto-mounted at startup.

- i. Unmount the partition.

**umount** *Disk partition*

In this example, run the following command:

**umount /dev/vdb1**

- ii. Reload all the content in the **/etc/fstab** file.

**mount -a**

- iii. 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, auto mount has taken effect:

```
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 console, but the disk cannot be viewed on the server. For example, Volume (D:) was not shown in **This PC** of a server running Windows Server 2012. Normally, Volume (D:) appears, as shown in [Figure 17-6](#).

**Figure 17-6** 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 about how to initialize data disks, see [Initializing EVS Data Disks](#).

### 17.3.2 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 17-6** Troubleshooting

Symptom	Solution
The target server cannot be found on the <b>Attach Disk</b> page.	<ul style="list-style-type: none"> <li>Go to <a href="#">Check Whether the Disk and Server Are in the Same AZ</a>.</li> <li>Cloud servers created from ISO images are only used for OS installation. They have limited functions and cannot have EVS disks attached.</li> </ul>
The <b>Attach</b> button is grayed out.	<ul style="list-style-type: none"> <li>Go to <a href="#">Maximum Number of Disks That Can Be Attached to the Server Has Been Reached</a>.</li> </ul>
An incorrect OS type warning is displayed when a shared disk is attached.	Go to <a href="#">Check Whether the Servers Attached with the Shared Disk Are Running the Same Type of OS</a> .


## 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.

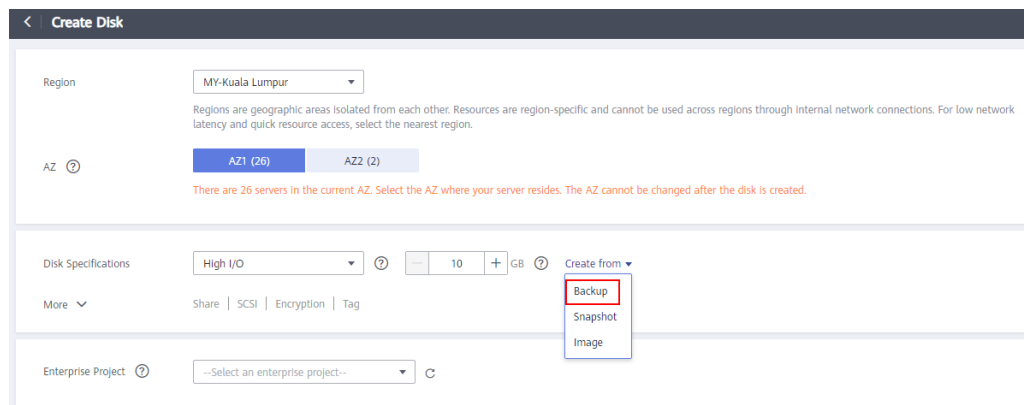
- If the disk data is no longer needed, delete the disk, and then create a new disk in the AZ where your target server is located.
- If the disk data is still required, create a new disk with the same data in the target AZ. The procedure is as follows:
  - a. Create a backup for the disk.

**Figure 17-7** Create Backup



- b. 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 17-8** 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 **In-use**, 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.

 **NOTE**

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?](#)

## Check Whether the Disk Has Been Added to a Replication Pair

**Symptom:** The **Attach** button is grayed out. When you hover the mouse over the **Attach** button, message "This operation cannot be performed on a disk in a replication pair" is displayed.

**Solution:** Delete the replication pair and attach the disk again.

## 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.

### 17.3.3 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.
- A shared file system or cluster management system must be set up before you can properly use a shared disk. If you simply attach a shared disk to multiple servers, data cannot be shared between those servers and may be overwritten.

### 17.3.4 Can I Attach a Disk to a Server in a Different AZ?

No.

Disks and the server you attach the disks to must be in the same AZ. The same is true for shared disks. A shared disk can only be attached to the servers in the same AZ.

### 17.3.5 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 about how to initialize data disks, see [Initializing EVS Data Disks](#).

### 17.3.6 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.

### 17.3.7 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.

### 17.3.8 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.

- System disk created along with a server: You can detach the system disk and then re-attach it to the same server as a system disk or data disk.

 NOTE

Each server can only have 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.


## 17.3.9 How Do I Obtain My Disk Device Name in the ECS OS Using the Device Identifier Provided on the Console?

### Scenarios

You find that the device name displayed in 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 in an ECS OS according to the device identifier on the console.

For details about how to attach disks, see [Attaching an EVS Disk](#).

### Obtaining the Disk ID of an ECS on the Console

1. Log in to the management console.
2. Under **Computing**, choose **Elastic Cloud Server**.
3. Click the target ECS name in the ECS list.  
The ECS details page is displayed.
4. Click the **Disks** tab and then click  to expand the disk information.
5. Check the device type and ID of the disk.

 NOTE

If **Device Identifier** is not displayed on the page, stop the ECS and restart it.

– KVM ECS

- 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 device name, see [Using a Serial Number to Obtain the Disk Device 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 BDF to Obtain a Disk Device Name \(Linux\)](#). (BDF cannot be used to obtain the disk device name of Windows ECSs.)

- If **Device Type** is **SCSI**, use a WWN to obtain the disk device 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 Device 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
```

#### NOTE

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:

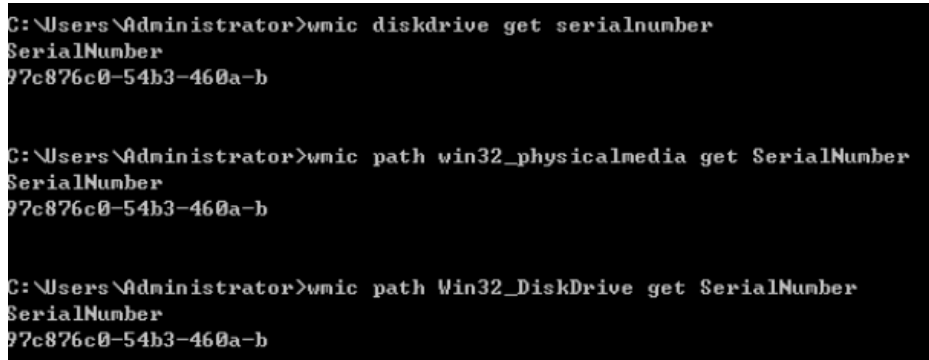
```
wmic diskdrive get serialnumber
```

```
wmic path win32_physicalmedia get SerialNumber
```

```
wmic path Win32_DiskDrive get SerialNumber
```

Information similar to the following is displayed:

**Figure 17-9** Obtaining the disk serial number



```
C:\Users\Administrator>wmic diskdrive get serialnumber
SerialNumber
97c876c0-54b3-460a-b

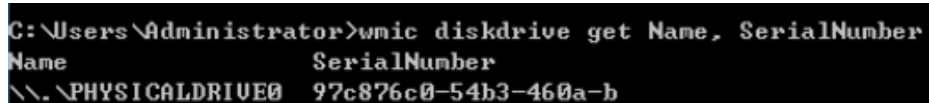
C:\Users\Administrator>wmic path win32_physicalmedia get SerialNumber
SerialNumber
97c876c0-54b3-460a-b

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 17-10** Checking the disk corresponding to the serial number



```
C:\Users\Administrator>wmic diskdrive get Name, SerialNumber
Name                SerialNumber
\\.\PHYSICALDRIVE0 97c876c0-54b3-460a-b
```

### PowerShell

1. 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
  - i. 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 17-11](#), 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 17-11](#), the disk is **Disk 0**.

**Figure 17-11** Viewing the disk on which the logical disk is created

```
PS C:\Users\Administrator> Get-CimInstance -ClassName Win32_LogicalDiskToPartition |select Antecedent, Dependent | fl
Antecedent : Win32_DiskPartition (DeviceID = "Disk #0, Partition #1")
Dependent  : Win32_LogicalDisk (DeviceID = "C:")

PS C:\Users\Administrator> Get-Disk |select Number, SerialNumber
Number SerialNumber
-----
0 97c876c0-54b3-460a-b
1 dswfa16520d39517815206127
```

- 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/*
```

### 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 BDF 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 ID/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)

1. 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 68886030000**3252ffa**16520d39517815.

- a. Obtain the 21st to 17th digits that are counted backwards (**3252f**).
- b. Convert a hexadecimal (**3252f**) to a decimal (**206127**).

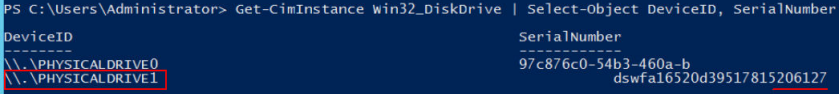
3. Start PowerShell as an administrator in a Windows OS.

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 17-12** Disk with the serial number ending with **206127**



```
PS C:\Users\Administrator> Get-CimInstance Win32_DiskDrive | Select-Object DeviceID, SerialNumber
DeviceID                               SerialNumber
-----                               -
\\.\PHYSICALDRIVE0                     97c876c0-54b3-460a-b
\\.\PHYSICALDRIVE1                     dswfa16520d39517815206127
```

## 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 WWN|grep 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-36888603000008b32fa16688d09368506 -> ../../sda
```

## 17.4 Capacity Expansion

### 17.4.1 Can I Reduce or Temporarily Expand the Disk Capacity?

No. The disk capacity can only be expanded, and temporary capacity expansion is not supported.

### 17.4.2 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.

### 17.4.3 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. You are advised to back up the disk data before expanding capacity.

- For details about disk backups, see the [Backing Up EVS Disks](#).
- For details about disk snapshots, see [Creating an EVS Snapshot](#).

### 17.4.4 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 from 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 \(Linux\)](#)

## 17.4.5 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 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 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 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.

## 17.4.6 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 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 capacity expansion of **In-use** disks.

### NOTE

Only some server OSs support capacity expansion of **In-use** disks. For details, see [How Do I Extend the Disk Partition and File System in Windows Server 2008?](#)

2. Log in to the server and create a new partition or allocate the additional space to one that is already there.

The unmount operation is not required, either for Windows or Linux.

## 17.4.7 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).

## 17.4.8 How Can I Allocate Newly Added Space to a New Partition?

### Windows

In Windows, if you already have a D drive and want to add an E drive, refer to the following:

The "Creating a New Partition" part in [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, refer to the following:

The "Creating a New MBR Partition" and "Creating a New GPT Partition" parts in [Extending Disk Partitions and File Systems \(Linux\)](#)

## 17.4.9 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 drive, refer to the following:

The "System Disk: Add the Additional Space to the Original Volume" and "Data Disk: Add the Additional Space to the Original Volume" parts in [Extending Disk Partitions and File Systems \(Windows\)](#)

### Linux

In Linux, if your disk already has, for example, partition `/dev/vdb1` and you want to add space to this partition, refer to the following:

The "Extending an Existing Partition" part in [Extending Disk Partitions and File Systems \(Linux\)](#)

## 17.4.10 Why Did My Disk Capacity Remain Unchanged on the Server After a Capacity Expansion?

After expanding disk capacity on the 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 \(Linux\)](#)
- [Extending Disk Partitions and File Systems \(Windows\)](#)

## 17.4.11 Why 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 17-7** Possible causes

Possible Cause	Solution	Reference
A shared disk is still attached to a server.	Detach the disk from all the servers.	<a href="#">Shared Disk Is Still Attached to a Server</a>

### 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.

## 17.4.12 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
Filesystem Type Size Used Avail Use% Mounted on
/dev/vda1 ext4 43G 1.9G 39G 5% /
devtmpfs devtmpfs 2.0G 0 2.0G 0% /dev
tmpfs tmpfs 2.0G 0 2.0G 0% /dev/shm
tmpfs tmpfs 2.0G 9.1M 2.0G 1% /run
tmpfs tmpfs 2.0G 0 2.0G 0% /sys/fs/cgroup
tmpfs tmpfs 398M 0 398M 0% /run/user/0
/dev/vdb ext4 64G 55M 61G 1% /mnt/sdc
```

----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
         =                  crc=1      finobt=0, spinodes=0
data     =                  bsize=4096  blocks=2621440, imaxpct=25
         =                  sunit=0     swidth=0  blks
naming   =version2         bsize=4096  ascii-ci=0  ftype=1
log      =internal        bsize=4096  blocks=2560, version=2
         =                  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
/dev/vda1  ext4     40G  2.3G  35G   7% /
devtmpfs  devtmpfs 1.9G   0  1.9G   0% /dev
tmpfs     tmpfs    1.9G   0  1.9G   0% /dev/shm
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

## 17.4.13 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:

```
sdf      8:80   0    40G   0 disk
├─sdf1   8:81   0    500M  0 part /boot
├─sdf2   8:82   0     5G   0 part [SWAP]
├─sdf3   8:83   0   34.5G 0 part /
└─sdf4   8:84   0     64M 0 part
```

## 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:

```
sdf      8:80    0   140G  0 disk
├─sdf1   8:81    0    500M  0 part /boot
├─sdf2   8:82    0     5G   0 part [SWAP]
├─sdf3   8:83    0   34.5G  0 part /
└─sdf4   8:84    0     64M  0 part
```

The system disk (**sdf**) has been expanded from 40 GiB to 140 GiB. The **sdf4** partition (64 MiB) is the configdriver partition 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
```

```
[root@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:



```
sdf      8:80    0    140G    0 disk
├─sdf1   8:81    0    500M    0 part  /boot
├─sdf2   8:82    0     5G    0 part  [SWAP]
├─sdf3   8:83    0   34.5G   0 part  /
└─sdf4   8:84    0   97.7M   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
├─sdf1   8:81    0    500M    0 part  /boot
├─sdf2   8:82    0     5G    0 part  [SWAP]
├─sdf3   8:83    0   34.5G   0 part  /
└─sdf4   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 ~]# growpart /dev/sdf 3
CHANGED: partition=3 start=11511808 old: size=72243200 end=83755008 new: size=28
1889471,end=293401279
```

Run the **lsblk** command to view the new root partition size.

```
sdf      8:80    0    140G    0 disk
├─sdf1   8:81    0    500M    0 part  /boot
├─sdf2   8:82    0     5G    0 part  [SWAP]
├─sdf3   8:83    0  134.4G   0 part  /
└─sdf4   8:84    0   97.7M   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 ~]# 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
Filesystem      Size  Used Avail Use% Mounted on
/dev/sdf3       133G  1.9G  125G   2% /
devtmpfs        63G   0    63G   0% /dev
tmpfs           63G   0    63G   0% /dev/shm
tmpfs           63G  9.0M   63G   1% /run
tmpfs           63G   0    63G   0% /sys/fs/cgroup
/dev/sdf1       477M  104M  344M  24% /boot
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

## 17.4.14 How Do I Check the Disk Partition Style in Linux?

You can use either `fdisk` or `parted` to check the disk partition style.

- [Method 1: Check Partition Style and File System Format Using `fdisk`](#)
- [Method 2: Check Partition Style and File System Format Using `parted`](#)

### 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.

 **NOTE**

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?](#) 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
/dev/vda1  *            2048     83886079     41942016   83  Linux

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
```

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 possible reason may be that existing partitions use GPT and there is unallocated disk space. In this case, you cannot query all the partitions using `fdisk -l`. 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 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:

```
[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.

### NOTE

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?](#) 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

## 17.4.15 How Do I Extend the Disk Partition and File System in Windows Server 2008?

### Scenarios

After a disk is expanded on the 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 the Additional Space to the Original Volume](#).
- 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 a New Volume with the 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: Create a New Volume Using the Available Space Shrunk from the Original Volume](#).
- 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 the Additional Space to the Original Volume](#).
  - 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 a New Volume with the 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**

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 [Backing Up EVS Disks](#). For details about using snapshots, see [Creating an EVS Snapshot](#).

---

## Notes and 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.
- During an expansion, the additional space is added to the end of the disk. If the disk has multiple partitions, the additional space can only be allocated to the last partition of the disk.
- If a disk uses MBR, the storage space in excess of 2 TiB cannot be used because the maximum capacity that MBR supports is 2 TiB. If your disk already uses MBR for partitioning and you require more than 2 TiB after the capacity expansion, do as follows:
  - (Recommended) Create a new EVS disk and use GPT.

- Back up the disk data, perform the expansion, and then change the partition style from MBR to GPT. During this change, services will be interrupted and data on the disk will be erased.

## Prerequisites

- You have expanded the disk capacity and attached the disk to a server on the console.
- You have logged in to the ECS.
  - 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 the Additional Space to the Original Volume

In this example, the system disk has 50 GiB originally, and 22 GiB is added on the 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.

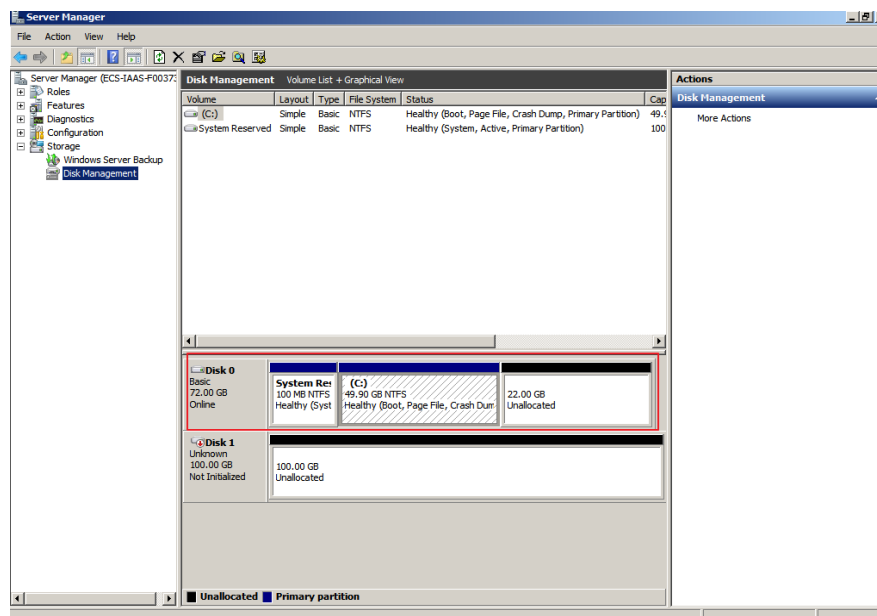
**Step 1** On the desktop of the server, right-click **Computer** and choose **Manage** from the shortcut menu.

The **Server Manager** window is displayed.

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

The **Disk Management** window is displayed.

Figure 17-13 Disk Management (system disk)

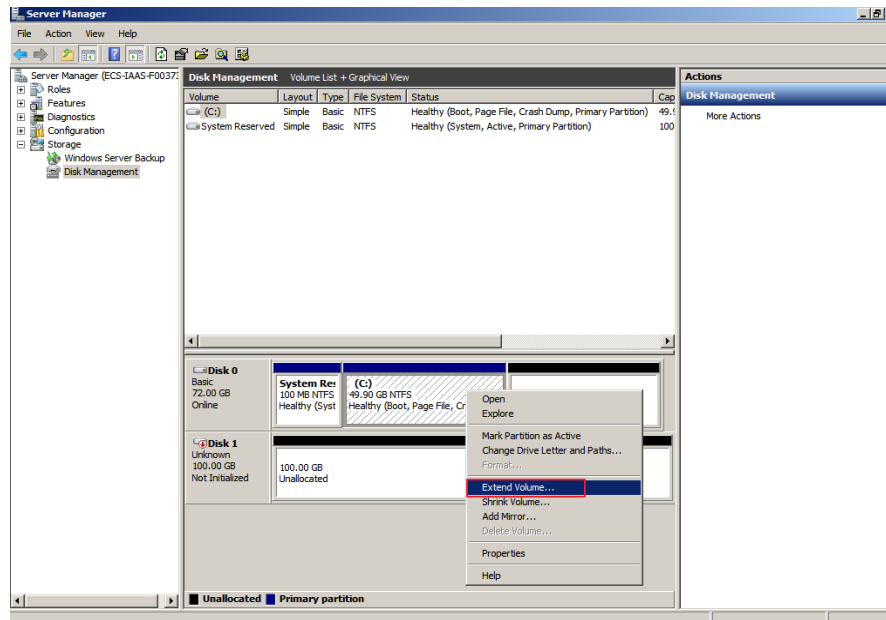


### NOTE

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

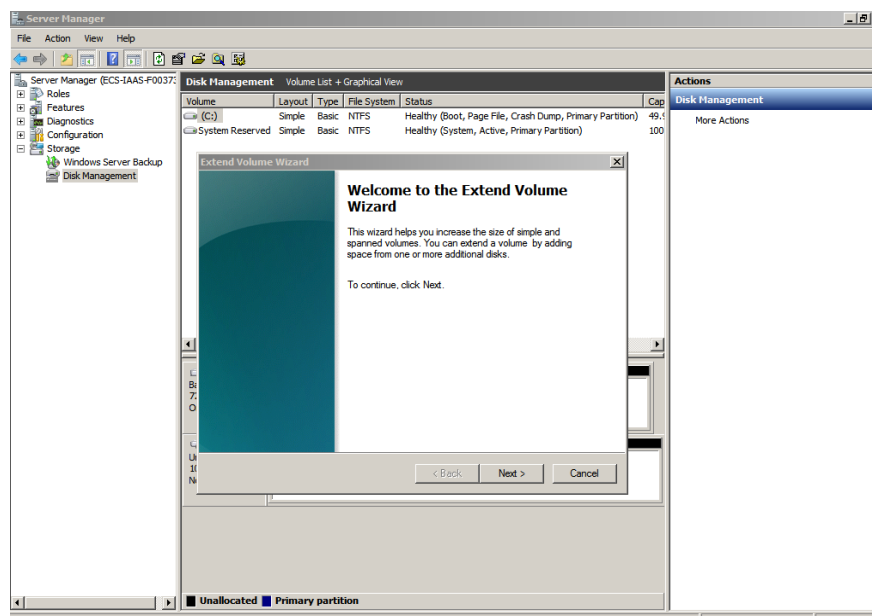
- Step 3** On the **Disk Management** page, select the disk and volume that you want to extend. The current volume size and unallocated space are displayed.
- Step 4** Right-click the target volume and choose **Extend Volume**.

**Figure 17-14** Choosing **Extend Volume**



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

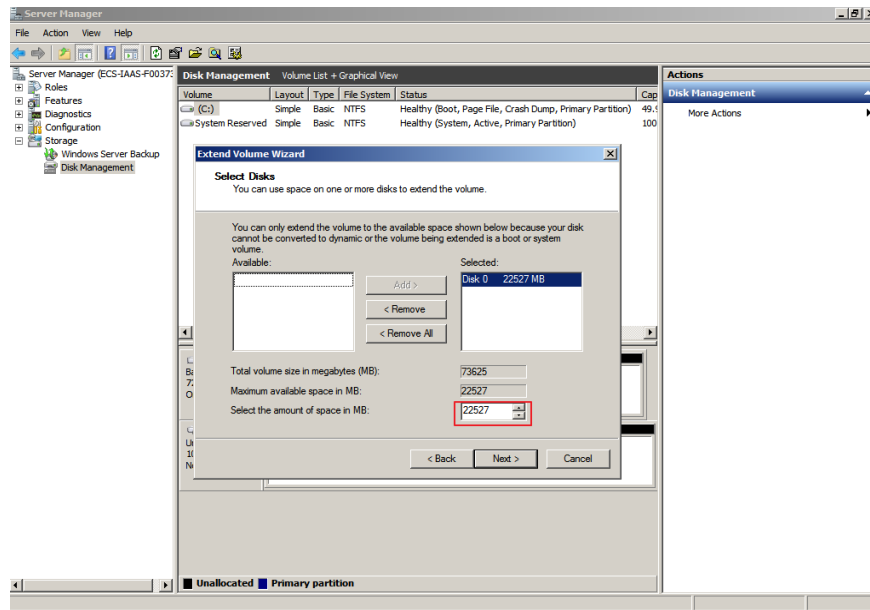
**Figure 17-15** Extend Volume Wizard



- 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**.



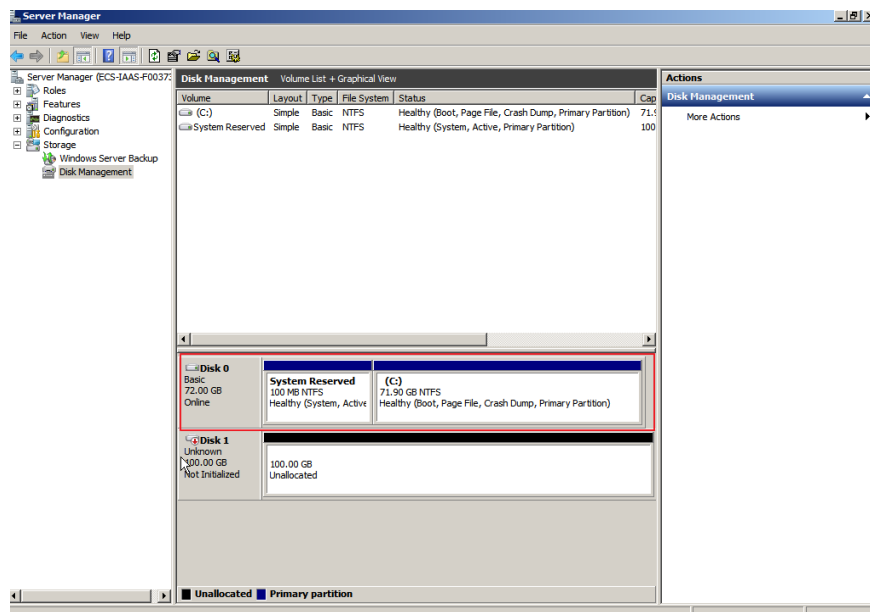
**Figure 17-16** Selecting space



**Step 7** Click **Finish**.

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

**Figure 17-17** Capacity expansion succeeded



----End

## System Disk: Create a New Volume with the Additional Space

In this example, the system disk has 40 GiB originally, and 60 GiB is added on the 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.

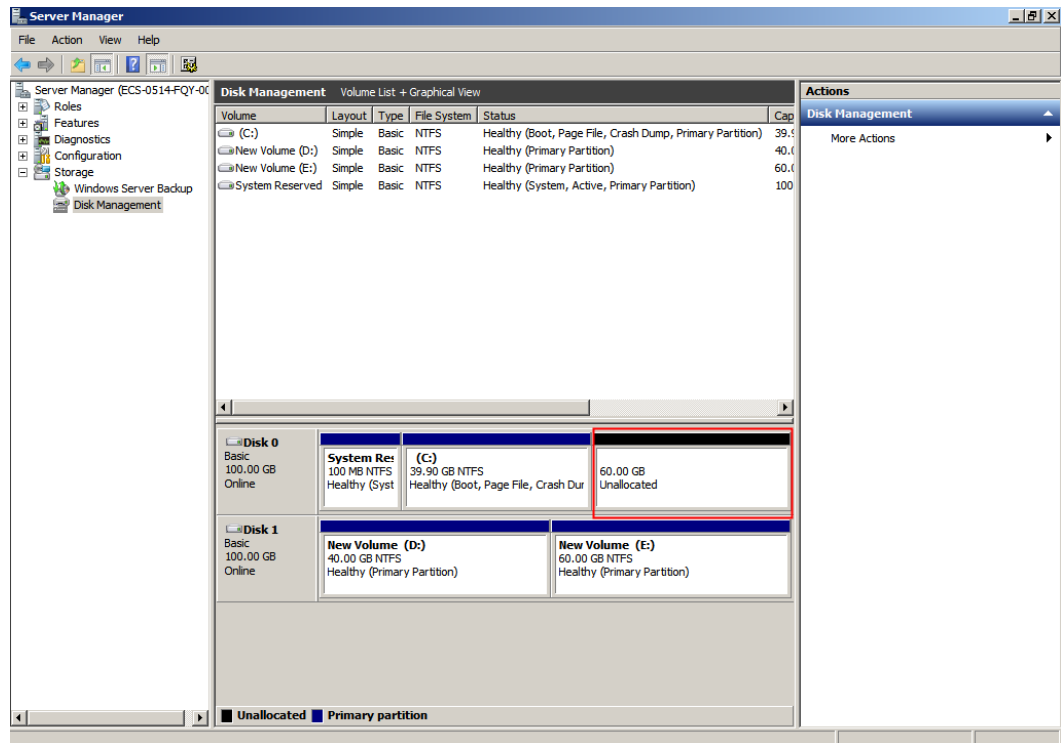
**Figure 17-18** Refresh (system disk)



**Step 3** If you cannot see 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 17-19** Unallocated disk space (system disk)



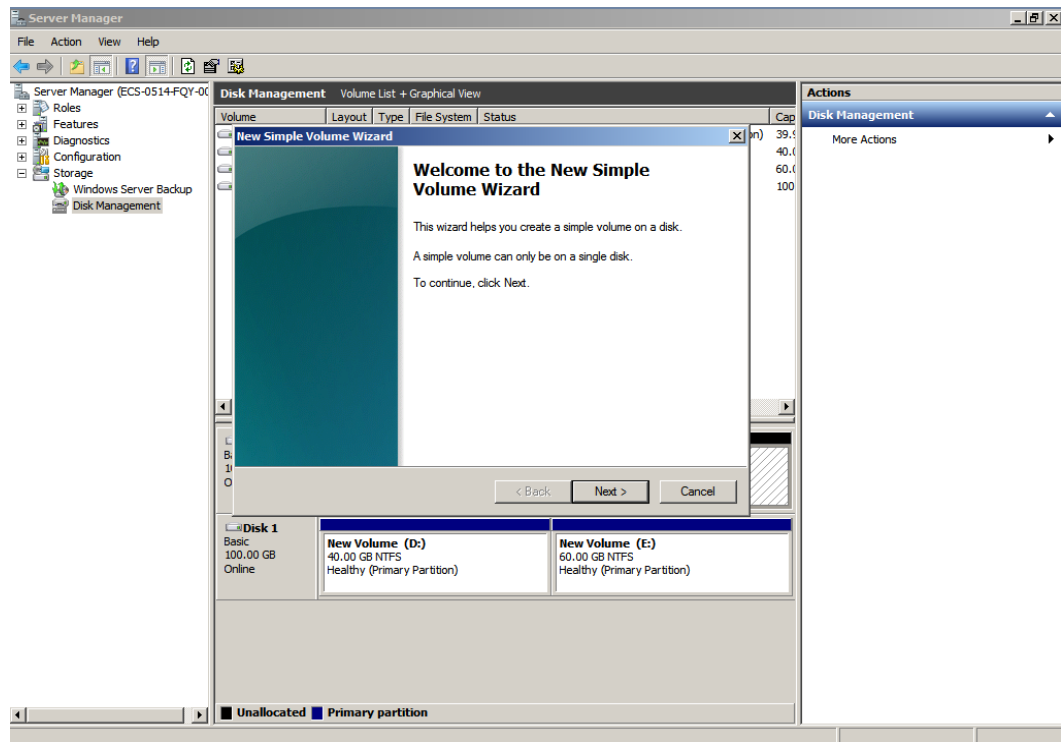
**Step 4** In the **Unallocated** area of **Disk 0**, right-click the blank area and choose **New Simple Volume**.

**Figure 17-20** New Simple Volume (system disk)



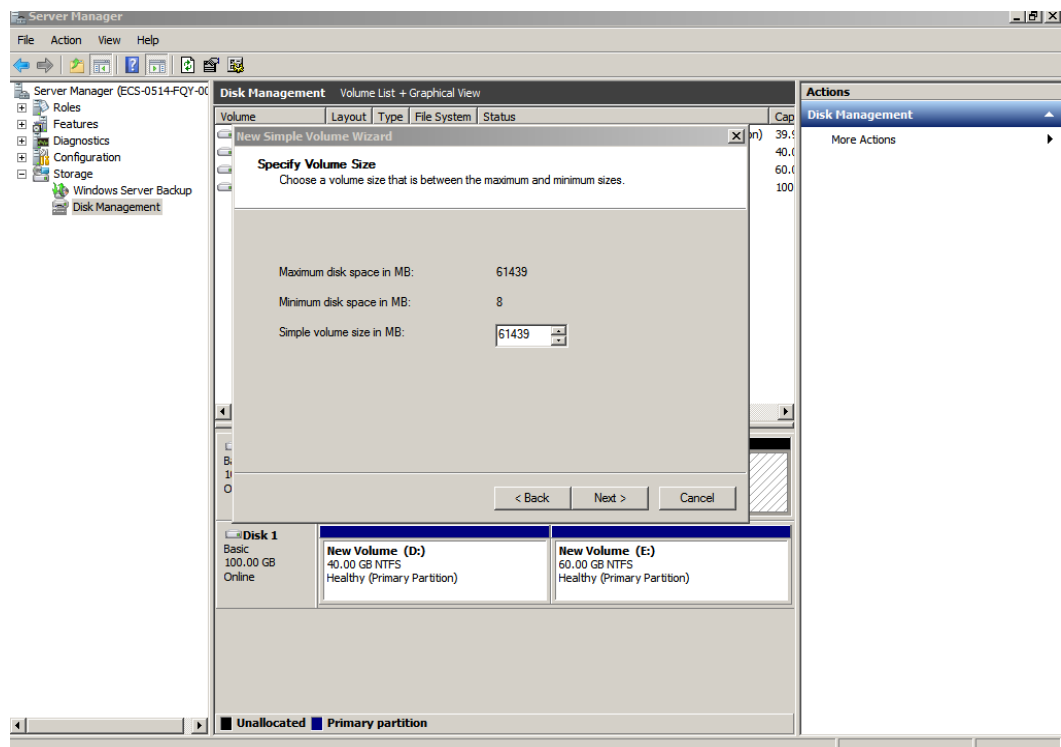
**Step 5** On the displayed **New Simple Volume Wizard** window, click **Next**.

Figure 17-21 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.

Figure 17-22 Specify Volume Size (system disk)




**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.

**Figure 17-23** Assign Drive Letter or Path (system 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**.

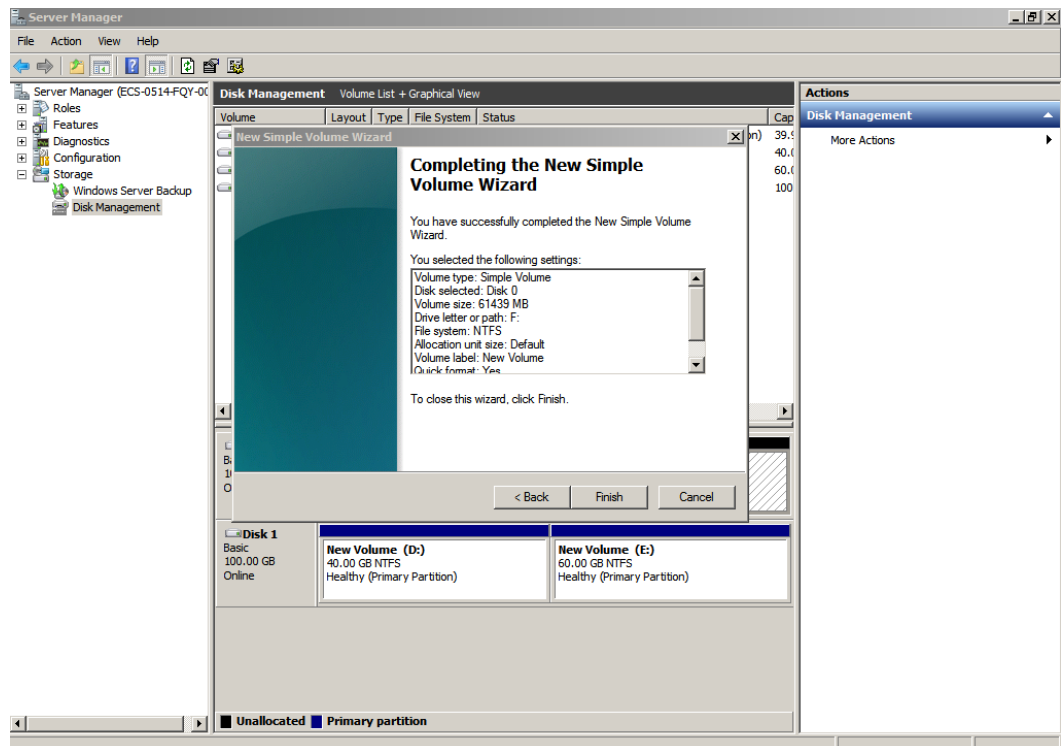
Figure 17-24 Format Partition (system disk)



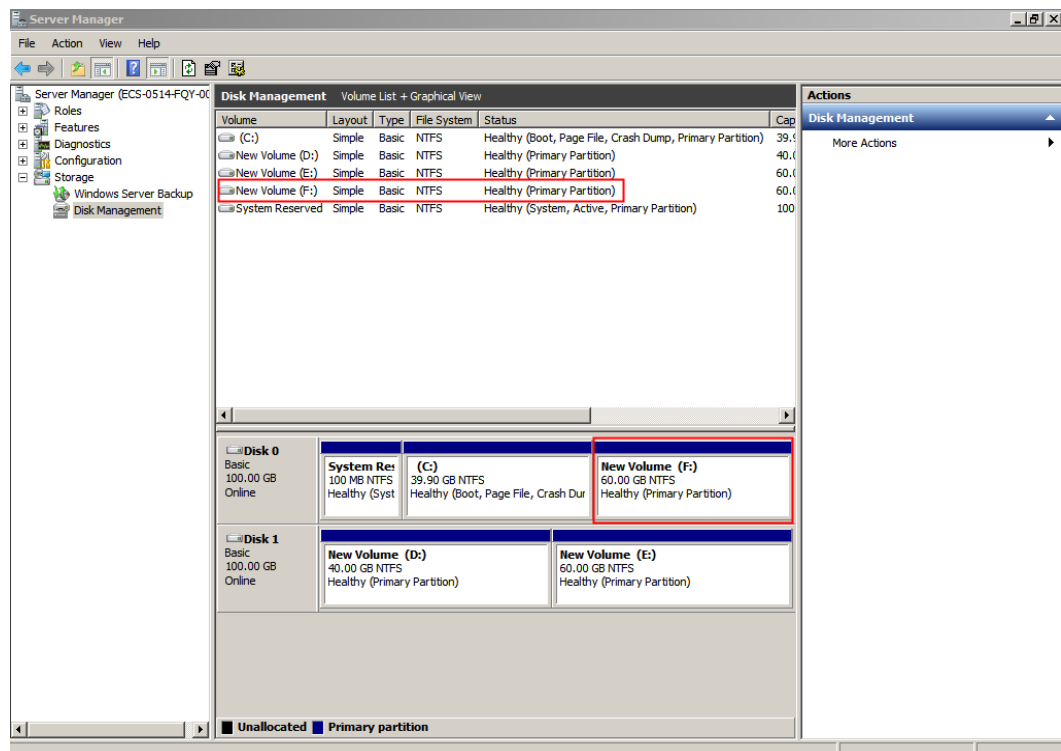
Step 9 Click **Finish**.

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

Figure 17-25 Completing the New Simple Volume Wizard (new volume F:)



**Figure 17-26 New Volume (F:)**



-----End

## System Disk: Create a New Volume Using the Available Space Shrunk from the Original Volume

In this example, the system disk has 40 GiB originally, and 60 GiB is added on the 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.

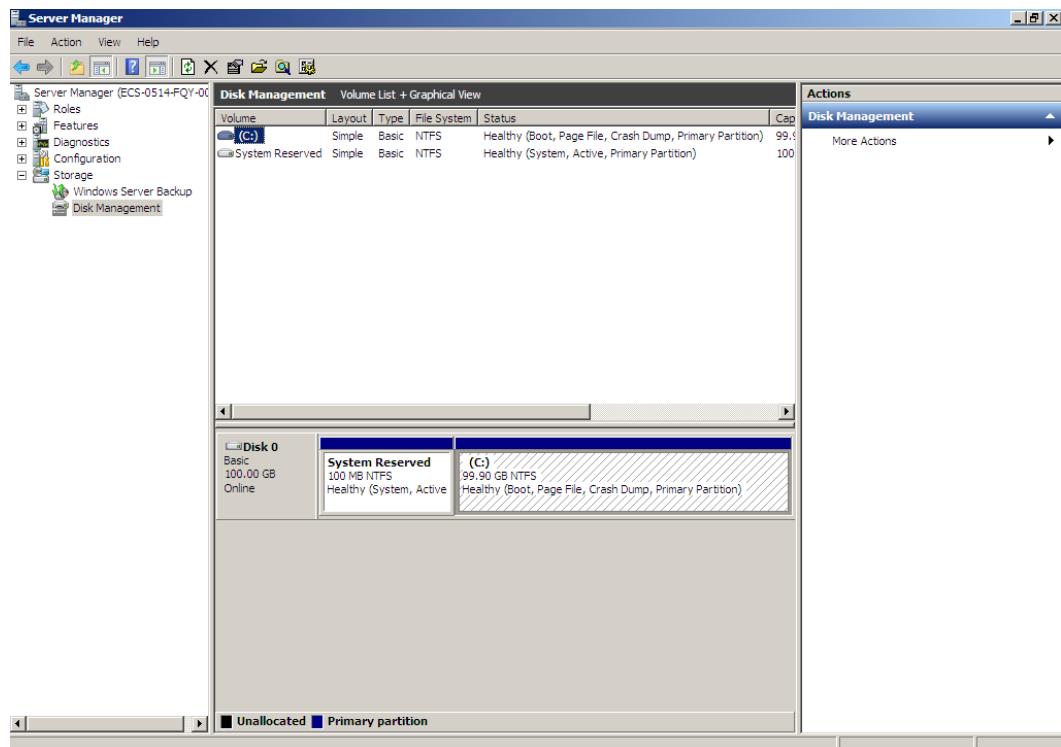
- Step 1** On the desktop of the server, right-click **Computer** and choose **Manage** from the shortcut menu.

The **Server Manager** window is displayed.

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


The **Disk Management** window is displayed.

Figure 17-27 Refresh (shrink volume)



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

Figure 17-28 Shrink Volume




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




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

**Figure 17-29 Shrink (shrink volume)**




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

**Figure 17-30 Unallocated (shrink volume)**




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

**Figure 17-31** New Simple Volume (shrink volume)




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

**Figure 17-32** New Simple Volume Wizard (shrink volume)



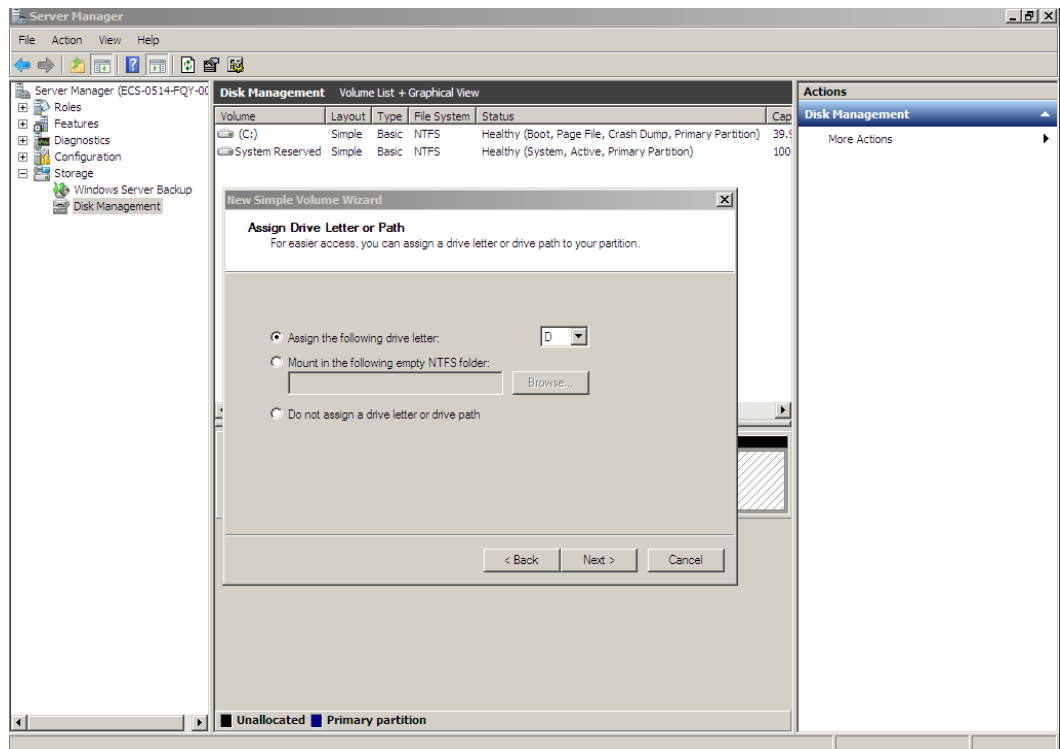
**Step 7** On the displayed **Specify Volume Size** page, set **Simple volume size in MB** and click **Next**. In this example, the default size is used.

**Figure 17-33** Specify Volume Size (shrink volume)



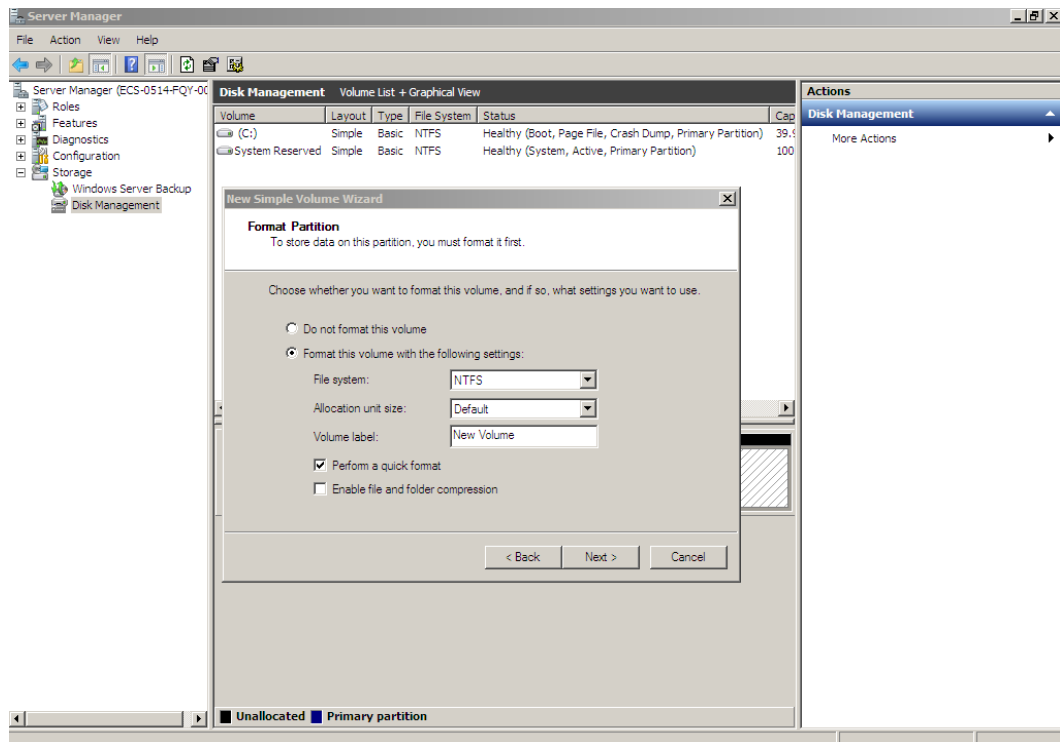
**Step 8** On the displayed **Assign Drive Letter and Path** page, click **Assign the following drive letter**, select a drive letter, and click **Next**. In this example, drive letter **D** is selected.

**Figure 17-34** Assign Drive 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**.

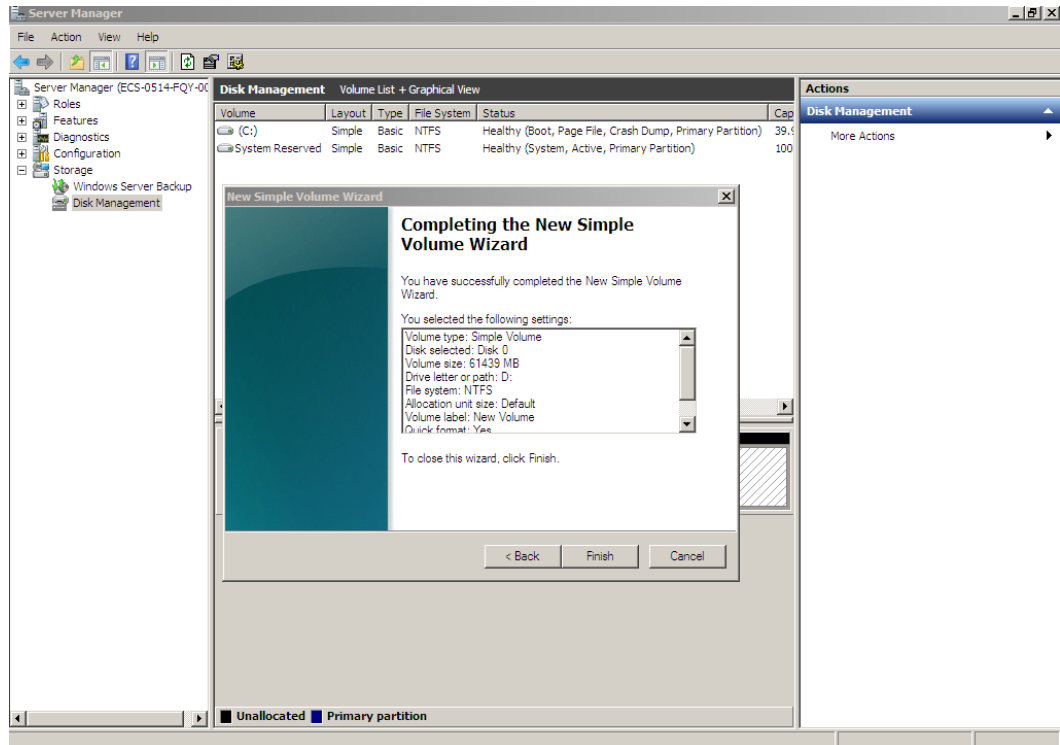
**Figure 17-35** Format Partition (shrink volume)




**Step 10** Click **Finish**.

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

**Figure 17-36** Completing the New Simple Volume Wizard (new volume D:)



**Figure 17-37** New Volume (D:)



----End

## Data Disk: Add the Additional Space to the Original Volume

In this example, the data disk has 100 GiB originally, and 50 GiB is added on the 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 17-38** Disk Management (data disk)




**NOTE**

If you cannot see 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.

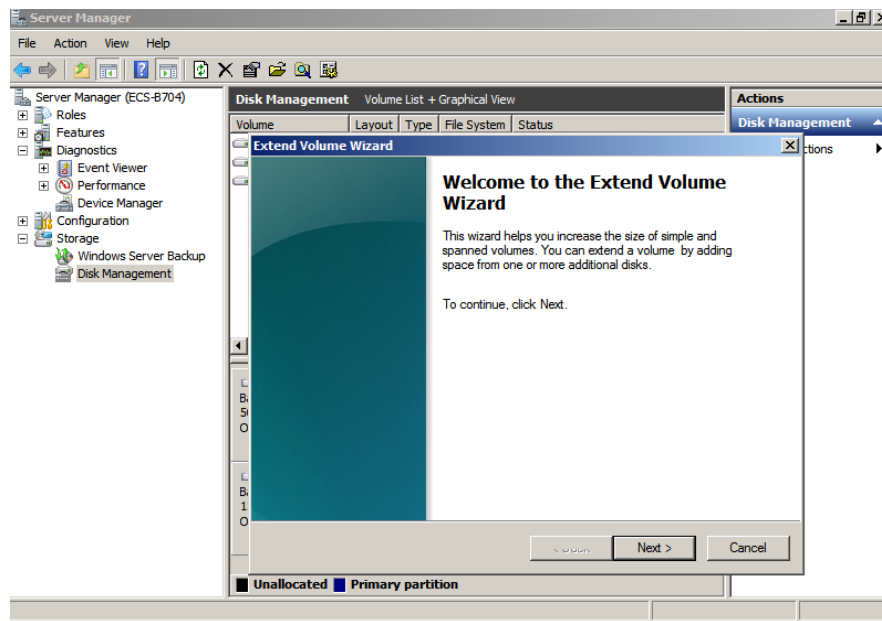
**Step 4** Right-click the target volume and choose **Extend Volume**.

**Figure 17-39** Choosing Extend Volume (Windows Server 2008)



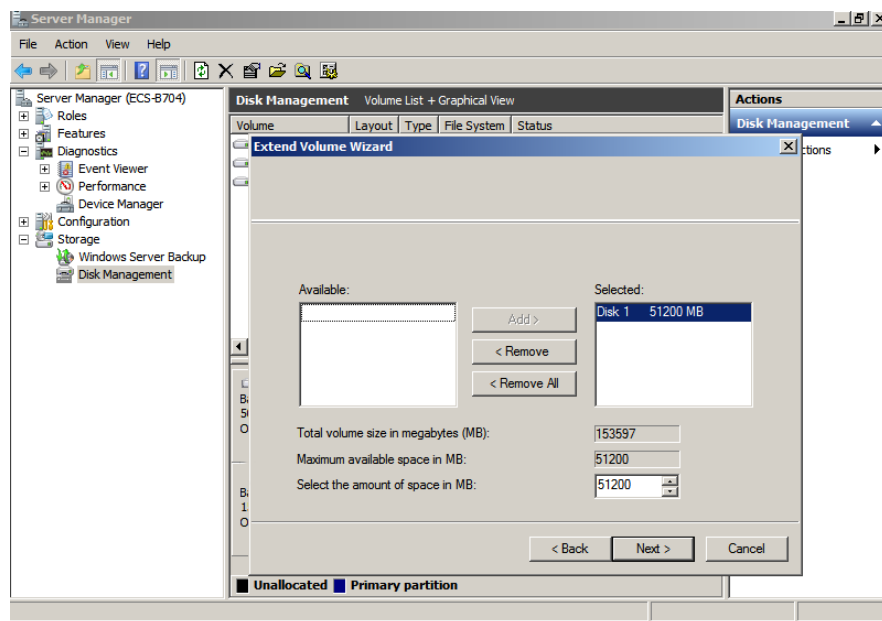
**Step 5** On the displayed **Extend Volume Wizard** window, click **Next**.

**Figure 17-40** 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 17-41** Selecting space (Windows Server 2008)

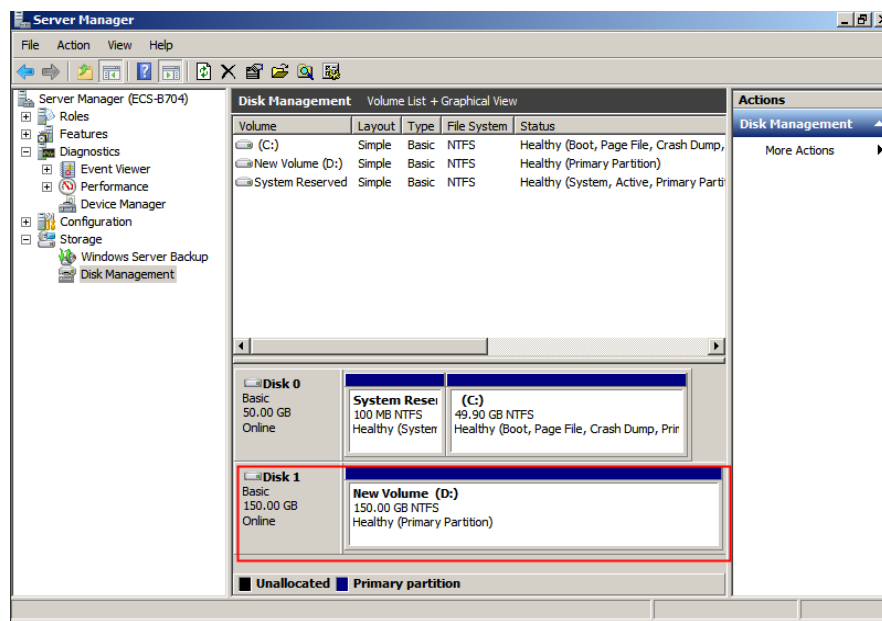


**Step 7** Click **Finish**.

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



**Figure 17-42** Capacity expansion succeeded (Windows Server 2008)



----End

## Data Disk: Create a New Volume with the Additional Space

In this example, the data disk has 40 GiB originally, and 60 GiB is added on the 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.

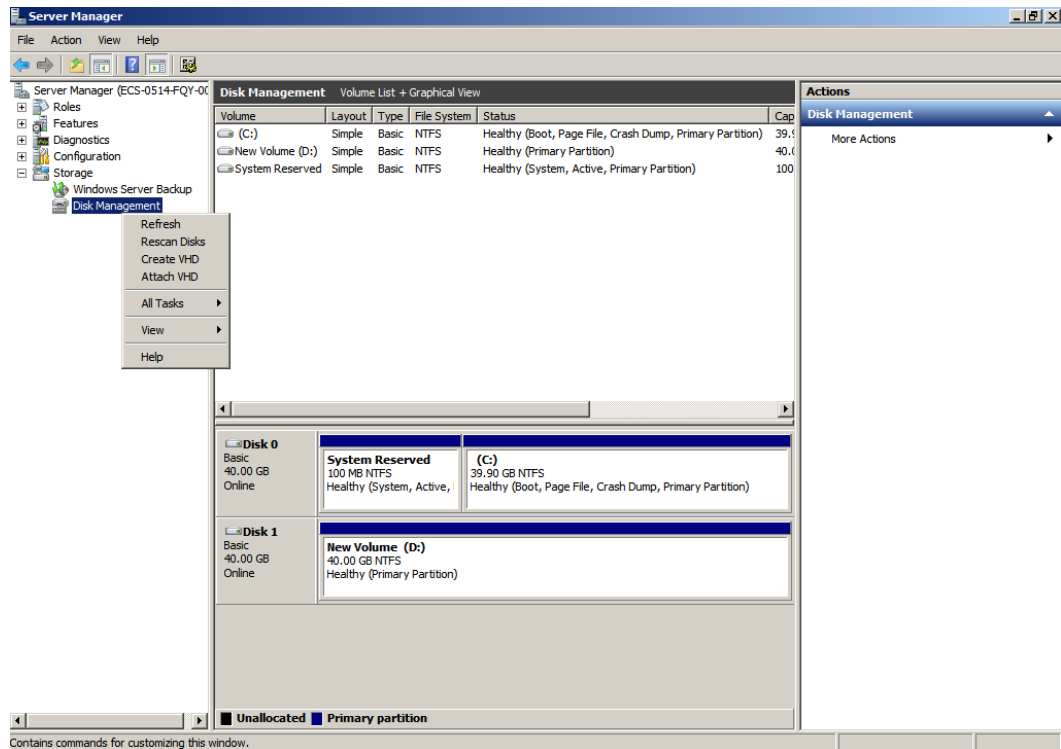
**Step 1** On the desktop of the server, right-click **Computer** and choose **Manage** from the shortcut menu.

The **Server Manager** window is displayed.

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

The **Disk Management** window is displayed.

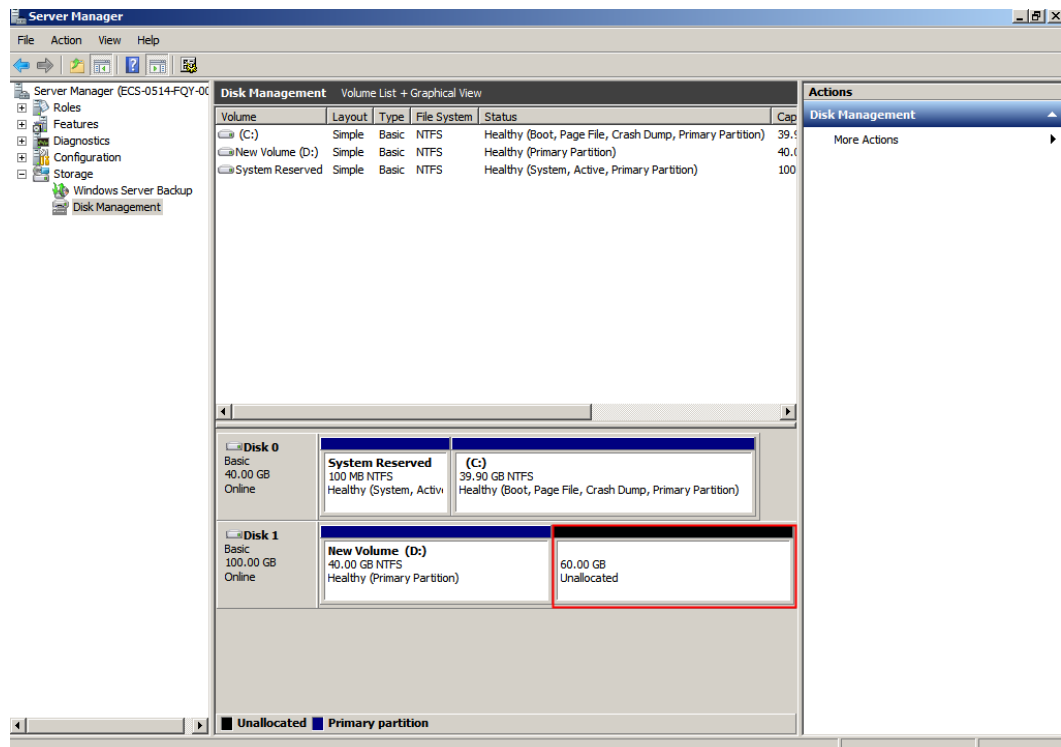
Figure 17-43 Refresh (data disk)



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

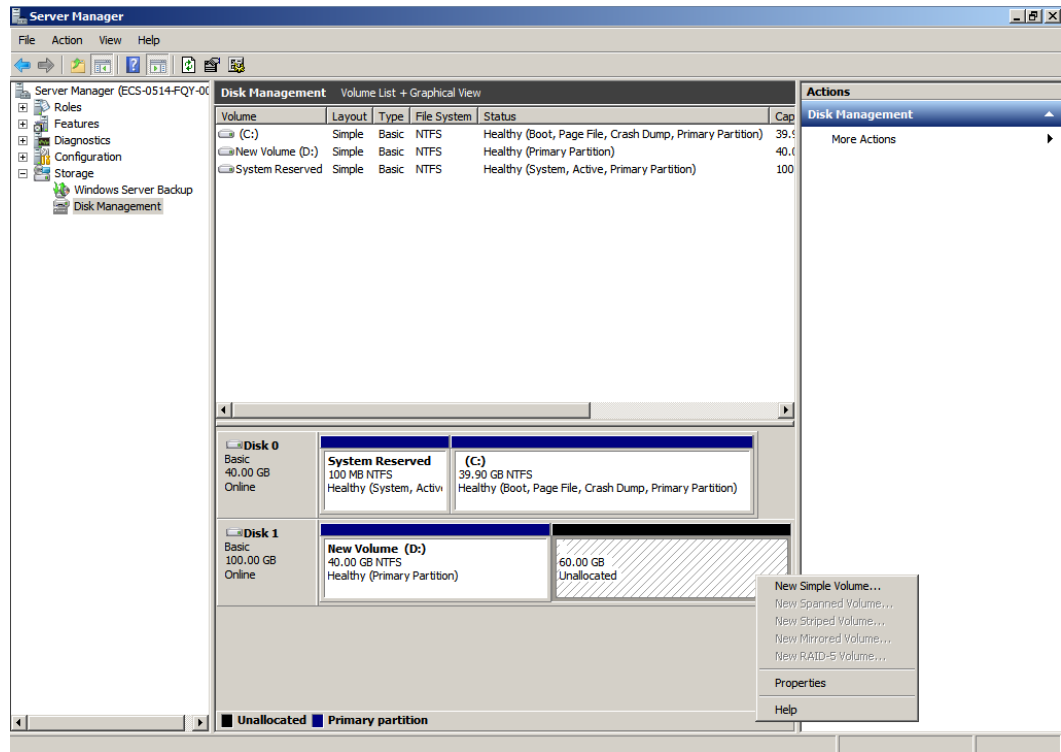
After the refresh, the additional space is displayed in the right area and is unallocated.

Figure 17-44 Unallocated disk space (data disk)



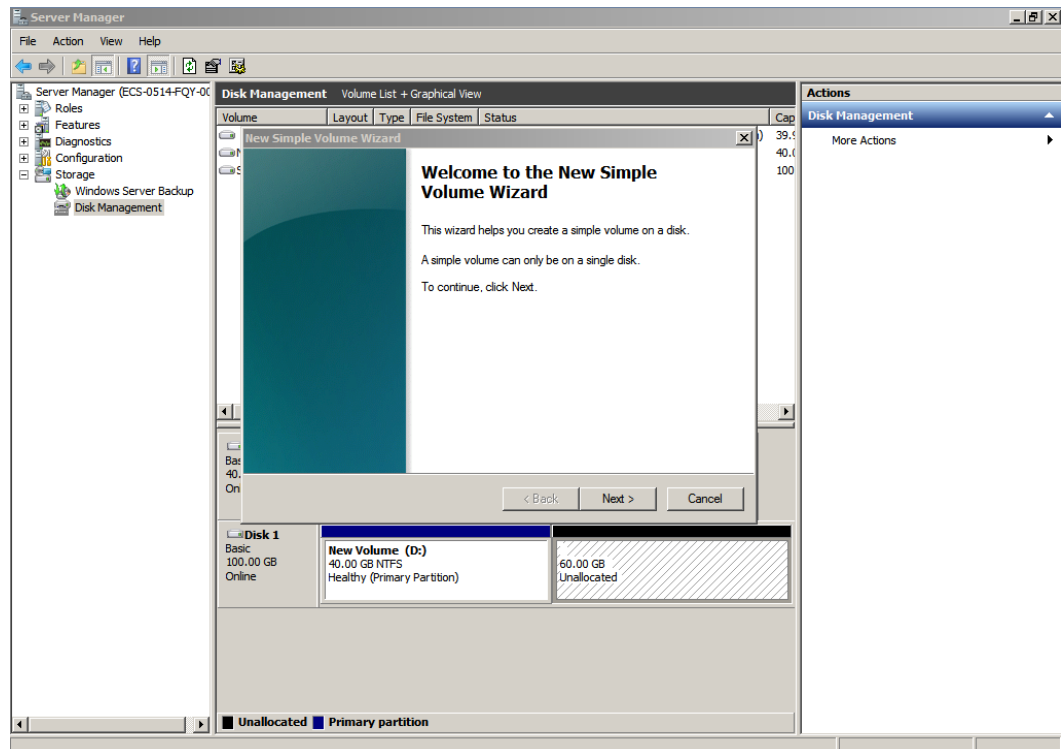
**Step 4** In the **Unallocated** area of **Disk 1**, right-click the blank area and choose **New Simple Volume**.

**Figure 17-45** New Simple Volume (data disk)




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

**Figure 17-46** New Simple Volume Wizard (data disk)



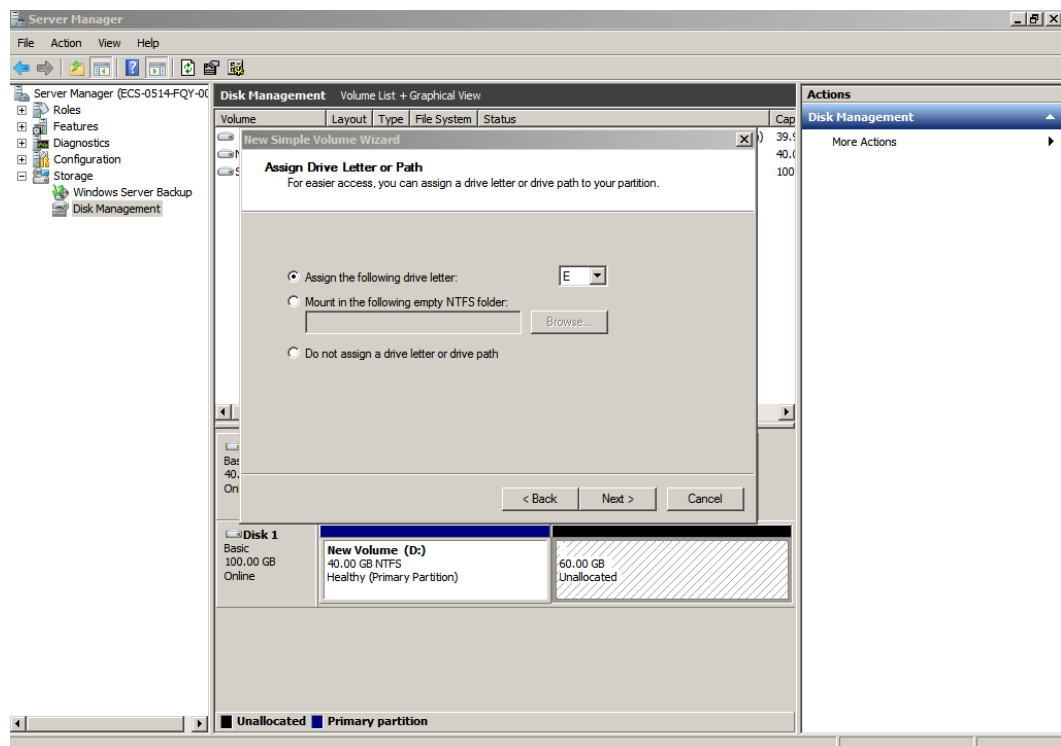
**Step 6** On the displayed **Specify Volume Size** page, set **Simple volume size in MB** and click **Next**. In this example, the default size is used.

**Figure 17-47** Specify Volume Size (data disk)



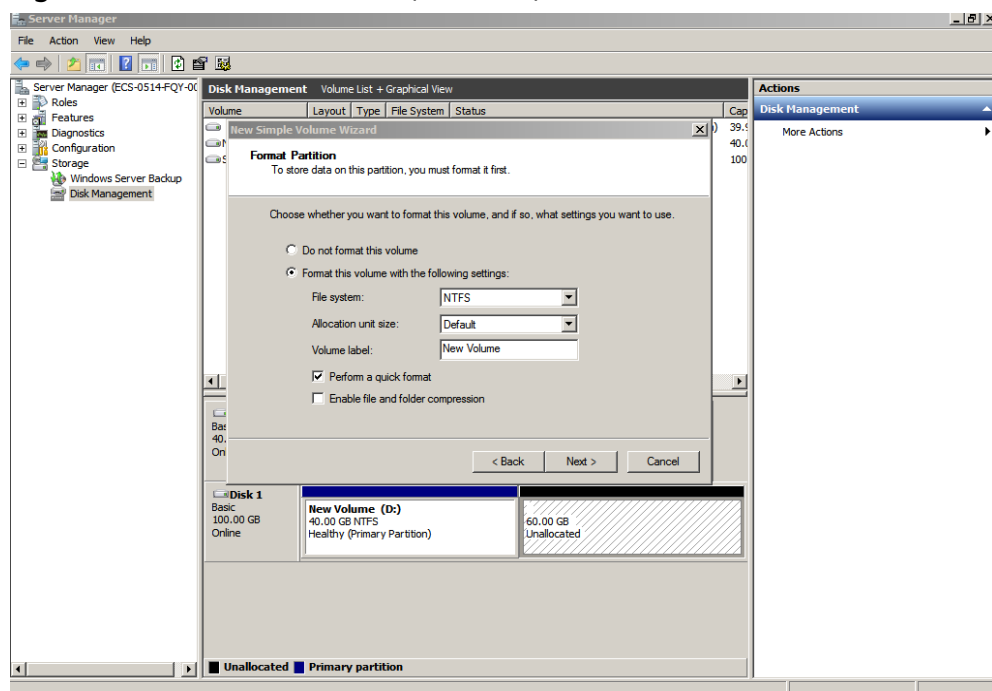
**Step 7** On the displayed **Assign Drive Letter and Path** page, click **Assign the following drive letter**, select a drive letter, and click **Next**. In this example, drive letter **E** is selected.

**Figure 17-48 Assign Drive 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**.

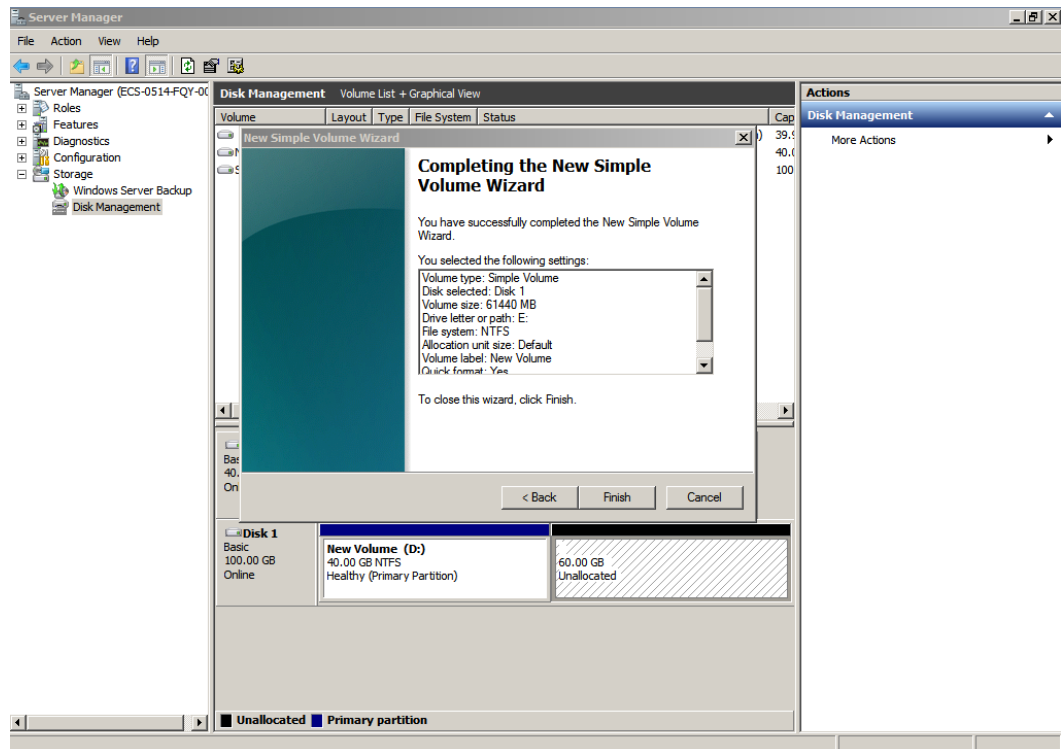
**Figure 17-49 Format Partition (data disk)**




**Step 9** Click **Finish**.

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

**Figure 17-50** Completing the New Simple Volume Wizard (new volume E:)



**Figure 17-51** New Volume (E:)



----End

## 17.4.16 How Do I Extend the Disk Partition and File System in Linux?

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

1. For how to use `fdisk` and `parted` to check the disk partition style, see [How Do I Check the Disk Partition Style in Linux?](#)
2. For how to select an appropriate partition style, see [Table 17-8](#).

**Table 17-8** Operation instructions of extending partitions and file systems in Linux

Scenario	Partition Style	Disk Function	OS Requirements	File System Format	Capacity Expansion Tool	Example Configuration
	GPT or MBR	System disk Data disk	<ul style="list-style-type: none"> <li>To extend partitions and file systems of a system disk, the kernel version must be later than 3.6.0.</li> <li>To extend partitions and file systems of a data disk, there is no requirements on the OS version.</li> </ul>	ext* (such as ext2, ext3, and ext4), xfs, and btrfs	growpart	<ul style="list-style-type: none"> <li>Device name: /dev/vdb</li> <li>Existing partition: /dev/vdb1</li> <li>Space added: 50 GiB</li> </ul>



Scenario	Partition Style	Disk Function	OS Requirements	File System Format	Capacity Expansion Tool	Example Configuration
	MBR	System disk	The kernel version is earlier than 3.6.0.	ext* (such as ext2, ext3, and ext4), xfs, and btrfs	dracut-modules-growroot	<ul style="list-style-type: none"> <li>• Device name: /dev/vda</li> <li>• File system format: ext4</li> <li>• Mount point: /mnt/sda</li> <li>• Partition name: /dev/vda1</li> <li>• Space added: 60 GiB</li> <li>• Partition style: MBR</li> </ul>

Scenario	Partition Style	Disk Function	OS Requirements	File System Format	Capacity Expansion Tool	Example Configuration
	MBR	System disk Data disk	None	ext* (such as ext2, ext3, and ext4), xfs, and btrfs	<ul style="list-style-type: none"> <li>• fdisk</li> <li>• parted</li> </ul>	<ul style="list-style-type: none"> <li>• Device name: /dev/vdb</li> <li>• File system format: ext4</li> <li>• Mount points: /mnt/sdc and /mnt/sdd</li> <li>• Partition 1: /dev/vdb1               <ul style="list-style-type: none"> <li>- Size: 100 GiB</li> <li>- Partition style: MBR</li> </ul> </li> <li>• Partition 2: /dev/vdb2               <ul style="list-style-type: none"> <li>- Size: 50 GiB</li> <li>- Partition style: MBR</li> </ul> </li> </ul>

Scenario	Partition Style	Disk Function	OS Requirements	File System Format	Capacity Expansion Tool	Example Configuration
	GPT	Data disk	None	ext* (such as ext2, ext3, and ext4), xfs, and btrfs	parted	<ul style="list-style-type: none"> <li>• Device name: /dev/vdb</li> <li>• File system format: ext4</li> <li>• Mount points: /mnt/sdc and /mnt/sdd</li> <li>• Partition 1: /dev/vdb1 <ul style="list-style-type: none"> <li>- Size: 2 TiB</li> <li>- Partition style: GPT</li> </ul> </li> <li>• Partition 2: /dev/vdb2 <ul style="list-style-type: none"> <li>- Size: 1 TiB</li> <li>- Partition style: GPT</li> </ul> </li> </ul>

 **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 TB, 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.

## 17.5 Detachment

## 17.5.1 If I Detach a Disk, Will I Lose the Data on My Disk?

- Encrypted:
  - The CMK is disabled or scheduled for deletion.  
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 be impossible to re-attach it later. In this case, do not detach the disk without a working CMK.
  - The CMK is available.  
The disk can be detached and re-attached, and data on the disk will not be lost.  
To ensure your data safety, you are advised to follow the instructions described in [Disk Detachment Process](#).
- Non-encrypted: Data on a disk will not be lost after the disk is detached, and the disk can be re-attached later if needed. To ensure your data safety, you are advised to follow the instructions described in [Disk Detachment Process](#).
- If you detach a disk that has an ongoing backup task, the task will fail.

### Disk Detachment Process

- For disks not supporting online detachment:
  - a. Stop the server where the disk was attached.
  - b. 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*.

## 17.5.2 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.

### NOTE

- In Linux, a system disk is typically mounted on `/dev/vda`. In Windows, a system disk is normally **Volume (C:)**.
- 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:)**.
- 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.
- Data disks: A data disk can be detached regardless of whether it is offline or online.
  - 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*.

## 17.6 Deletion

### 17.6.1 How Can I Recover Data from a Disk That Was Accidentally Deleted?

Disk data was deleted: Check whether the disk has any snapshots or backups created.

- If there are, use a snapshot or backup to restore the disk data to the state when the snapshot or backup was created. For details, see [Rolling Back Disk Data from a Snapshot](#) or "Restoring from a Cloud Disk Backup" in the *Cloud Backup and Recovery User Guide*.

---

#### NOTICE

If the disk was deleted after the last snapshot backup was created, that incremental data cannot be restored.

- 
- If there are not, the disk data cannot be restored.

## 17.7 Capacity

### 17.7.1 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.

### 17.7.2 What Should I Do If My Disk Starts to Run Out of Space?

If your disk space starts to fill up, you can:

1. Create a new disk and attach it to the server. For details, see [Creating an EVS Disk](#).
2. Expand the capacity of the existing disk. Both system disks and data disks can be expanded. An expansion operation includes two steps:
  - a. Expand the disk capacity on the console.
  - b. Log in to the server and extend the partition and file system.For details, see [Expanding the EVS Disk Capacity](#).
3. Free up the space on the disk. For details, see [How Do I Clean Up My Disk Space on a Windows Server?](#)

## 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.

## 17.7.3 How Do I Clean Up My Disk Space on a Windows Server?

### Scenarios

When the EVS disk space on a server is insufficient, the server running speed will be affected, which will further affect user experience. You can clean up the disk space using either of the following methods:

- [Cleaning Up Disk Space Using the System Built-in Cleanup Tool](#)

The following example uses a Windows Server 2016 server to illustrate how to clean up disk space on a server. In addition, you are advised to do the following to save disk space:

- Periodically compress and save the files that are not frequently used.
- Periodically use the disk cleanup tool to clean up disk space, delete unnecessary files, and clean up the recycle bin.
- Uninstall unnecessary programs to release disk space.

### Cleaning Up Disk Space Using the System Built-in Cleanup Tool

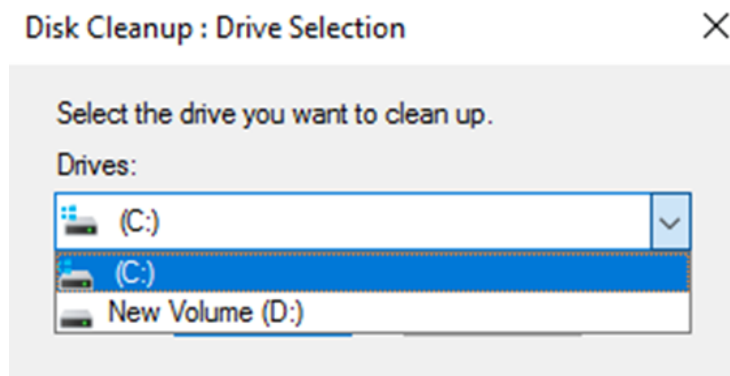
**Step 1** On the server desktop, click the start icon in the lower left corner.

The start menu is displayed.

**Step 2** In the navigation pane on the left, choose **Windows Administrative Tools > Disk Cleanup**.

The **Disk Cleanup: Drive Selection** window is displayed.

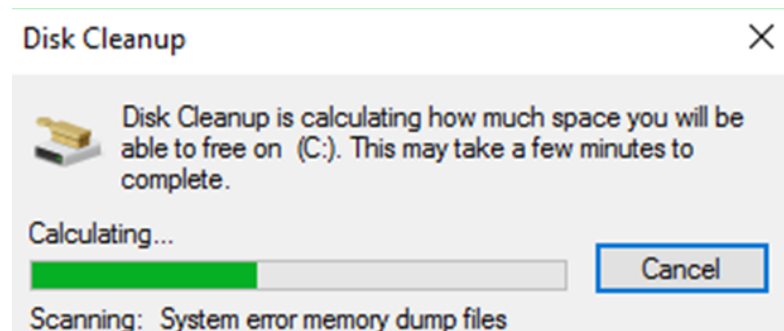
**Figure 17-52** Disk Cleanup: Drive Selection



**Step 3** Select the target disk from the drop-down list. In this example, disk (C:) is selected.

The **Disk Cleanup** window is displayed, and the system automatically calculates the space that can be freed on disk (C:).

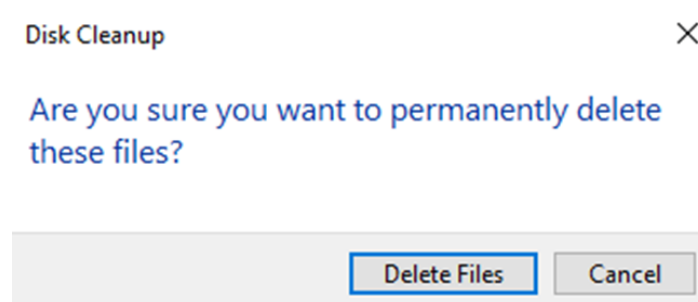
**Figure 17-53** Disk Cleanup



**Step 4** After the automatic calculation is complete, select the files to be deleted on the displayed window and click **OK**.

A confirmation dialog box is displayed.

**Figure 17-54** Deletion confirmation



**Step 5** Click **Delete Files** to clean up the disk space.

----End

## Uninstalling Unnecessary Programs on Control Panel

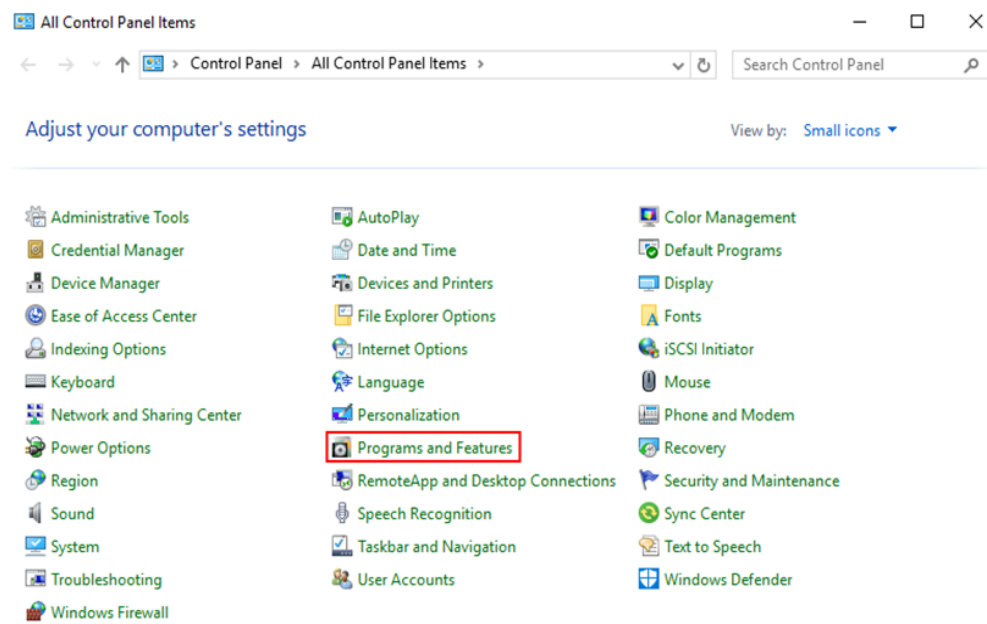
**Step 1** On the server desktop, click the start icon in the lower left corner.

The start menu is displayed.

**Step 2** In the navigation pane on the left, choose **Windows System > Control Panel**.

The **All Control Panel Items** window is displayed.

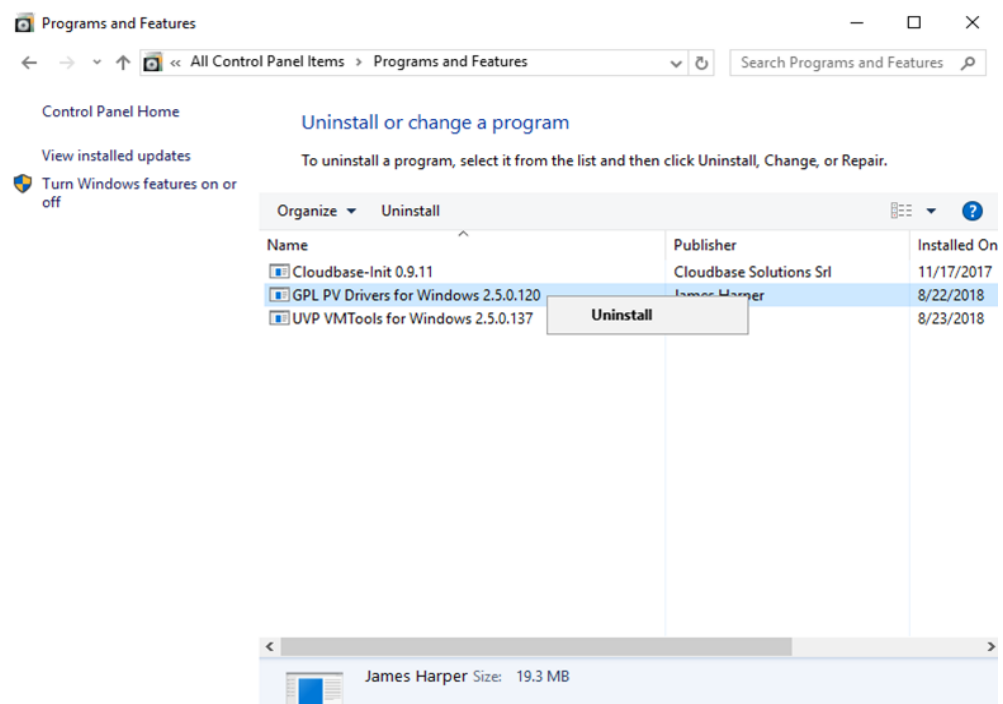
**Figure 17-55** All Control Panel Items



**Step 3** In the list, select **Programs and Features**.

The **Programs and Features** window is displayed.

**Figure 17-56** Programs and Features

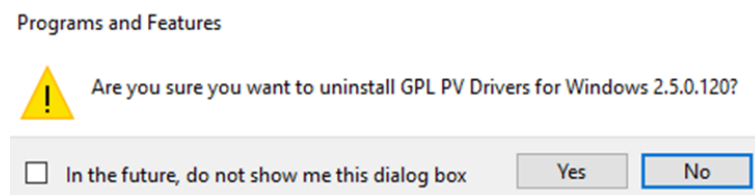


**Step 4** In the program list, right-click the program to be uninstalled and choose **Uninstall** from the shortcut menu.



A confirmation dialog box is displayed.

**Figure 17-57** Uninstallation confirmation



**Step 5** Click **Yes** to uninstall the program.

----End

## 17.7.4 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.

## 17.7.5 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 [Initializing EVS Data Disks](#).

## 17.7.6 How Can I View the Available Capacity of My Disk?

You can view your disk usage in either of the following ways:

- View the disk usage manually.  
The details depend on the OS. This FAQ uses Windows Server 2008, Windows Server 2016, and Linux as examples to describe how to view the disk usage.
  - [Viewing the Disk Usage in Linux](#)
  - [Viewing the Disk Usage in Windows Server 2008](#)
  - [Viewing the Disk Usage in Windows Server 2016](#)
- [Installing the Agent to View the Disk Usage](#)

### Viewing the 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 query 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
Filesystem Type Size Used Avail Use% Mounted on
/dev/vda1 ext4 43G 2.0G 39G 5% /
devtmpfs devtmpfs 509M 0 509M 0% /dev
tmpfs tmpfs 520M 0 520M 0% /dev/shm
tmpfs tmpfs 520M 7.2M 513M 2% /run
tmpfs tmpfs 520M 0 520M 0% /sys/fs/cgroup
tmpfs tmpfs 104M 0 104M 0% /run/user/0
/dev/vdb1 ext4 43G 51M 40G 1% /mnt/sdc
```

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 the 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, choose **Storage > Disk Management**.

The sizes and available spaces of the volumes on the current disk are displayed in the middle pane.

Figure 17-58 Disk Management page



----End

## Viewing the Disk Usage in Windows Server 2016

In this section, Windows Server 2016 Standard 64bit is used as an example.

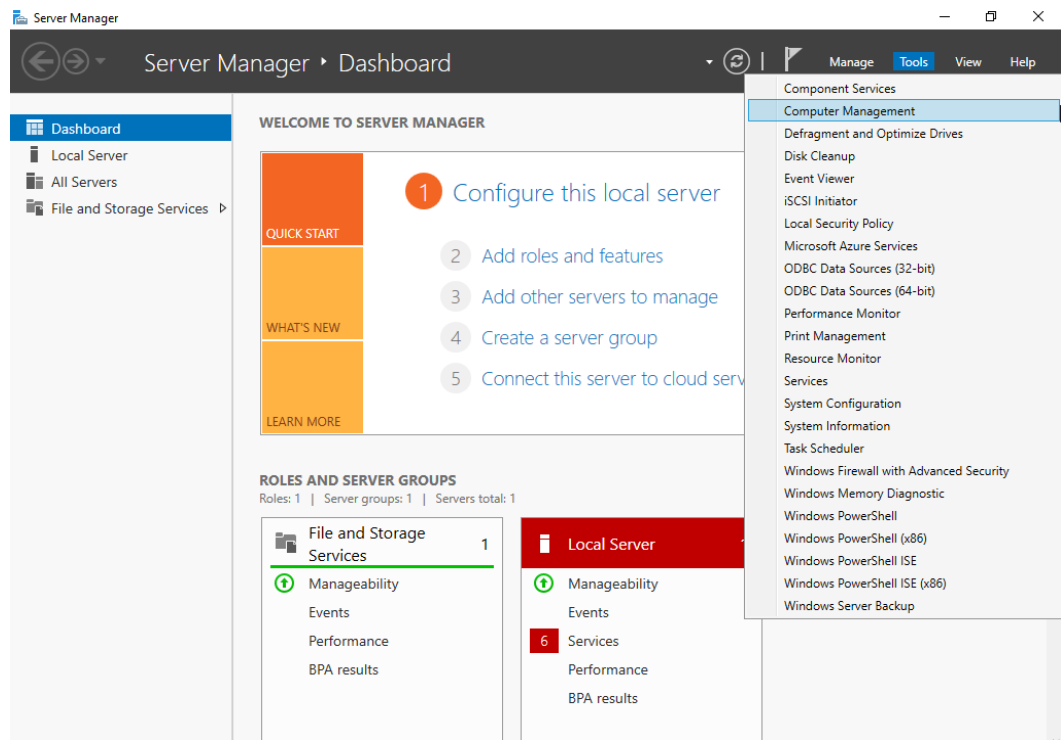
**Step 1** On the desktop of the server, click the start icon in the lower left corner.

The **Windows Server** window is displayed.

**Step 2** Click **Server Manager**.

The **Server Manager** window is displayed.

**Figure 17-59** 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 17-60** Disk list page



----End

## Installing the Agent to View the Disk Usage

Some disk monitoring metrics require that the agent to be installed.

For details about how to install the Agent on an ECS, see section "Installing and Configuring the Agent" in the *Cloud Eye User Guide*.

**Table 17-9** Disk metrics

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_free	(Agent) Available Disk Space	Free space on the disks Unit: GB <ul style="list-style-type: none"> <li>• Linux: Run the <b>df -h</b> command to check the value in the <b>Avail</b> column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</li> <li>• 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 (-), periods (.), and swung dashes (~).</li> </ul>	≥ 0	ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_total	(Agent) Disk Storage Capacity	<p>Total space on the disks, including used and free Unit: GB</p> <ul style="list-style-type: none"> <li>Linux: Run the <b>df -h</b> command to check the value in the <b>Size</b> column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</li> <li>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 (-), periods (.), and swung dashes (~).</li> </ul>	≥ 0	ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_used	(Agent) Used Disk Space	<p>Used space on the disks Unit: GB</p> <ul style="list-style-type: none"> <li>Linux: Run the <b>df -h</b> command to check the value in the <b>Used</b> column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</li> <li>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 (-), periods (.), and swung dashes (~).</li> </ul>	≥ 0	ECS - Mount point	1 minute



Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_usedPercent	(Agent) Disk Usage	<p>Percentage of total disk space that is used, which is calculated as follows:  <b>Disk Usage = Used Disk Space/Disk Storage Capacity</b></p> <p>Unit: percent</p> <ul style="list-style-type: none"> <li>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 (-), periods (.), and swung dashes (~).</li> <li>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 (-), periods (.), and swung dashes (~).</li> </ul>	0-100	ECS - Mount point	1 minute

### 17.7.7 How Can I Monitor My Disk Usages?

Some disk monitoring metrics require that the agent to be installed.

For details about how to install the Agent on an ECS, see section "Installing and Configuring the Agent" in the *Cloud Eye User Guide*.

**Table 17-10** Disk metrics

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_free	(Agent) Available Disk Space	<p>Free space on the disks Unit: GB</p> <ul style="list-style-type: none"> <li>Linux: Run the <b>df -h</b> command to check the value in the <b>Avail</b> column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</li> <li>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 (-), periods (.), and swung dashes (~).</li> </ul>	≥ 0	ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_total	(Agent) Disk Storage Capacity	<p>Total space on the disks, including used and free Unit: GB</p> <ul style="list-style-type: none"> <li>Linux: Run the <b>df -h</b> command to check the value in the <b>Size</b> column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</li> <li>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 (-), periods (.), and swung dashes (~).</li> </ul>	≥ 0	ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_used	(Agent) Used Disk Space	<p>Used space on the disks Unit: GB</p> <ul style="list-style-type: none"> <li>Linux: Run the <b>df -h</b> command to check the value in the <b>Used</b> column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</li> <li>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 (-), periods (.), and swung dashes (~).</li> </ul>	≥ 0	ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_usage_percent	(Agent) Disk Usage	<p>Percentage of total disk space that is used, which is calculated as follows:  <b>Disk Usage = Used Disk Space/Disk Storage Capacity</b></p> <p>Unit: percent</p> <ul style="list-style-type: none"> <li>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 (-), periods (.), and swung dashes (~).</li> <li>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 (-), periods (.), and swung dashes (~).</li> </ul>	0-100	ECS - Mount point	1 minute

### 17.7.8 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 an EVS 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 [Expanding the EVS Disk Capacity](#).

## 17.7.9 Why the Space of My New Disk Is Full After I Uploaded Only 500 MB of Files to the Disk?

Troubleshoot this issue by performing the following steps:

1. Check whether the disk partition usage is 100% or almost 100%.  
df -h

Figure 17-61 Checking the partition usage

```
[root@ecs-a058 /]# df -h
Filesystem      Size  Used Avail Use% Mounted on
devtmpfs        7.8G   0  7.8G   0% /dev
tmpfs           7.8G  20K  7.8G   1% /dev/shm
tmpfs           7.8G  8.7M  7.8G   1% /run
tmpfs           7.8G   0  7.8G   0% /sys/fs/cgroup
/dev/vda1       99G   99G   0 100% /
tmpfs           1.6G   0  1.6G   0% /run/user/0
```

In this example, the `/dev/vda1` partition usage is 100%.

2. Check the disk space usage.  
df -i

Figure 17-62 Checking the disk space usage

```
[root@ecs-a058 /]# df -i
Filesystem      Inodes   IUsed   IFree  IUse% Mounted on
devtmpfs        2030569   369 2030200    1% /dev
tmpfs           2033158    2 2033156    1% /dev/shm
tmpfs           2033158   507 2032651    1% /run
tmpfs           2033158   16 2033142    1% /sys/fs/cgroup
/dev/vda1       6553600 139249 6414351    3% /
tmpfs           2033158    1 2033157    1% /run/user/0
```

In this example, the disk space usage is low.

3. Check the deleted process files in the system.

```
lsuf | grep deleted
```

**Figure 17-63** Checking the deleted process files in the system

```
[root@ecs-a058 /]# lsuf | grep delete
agetty      2687      root      txt       REG       253,1     49640    1319348   /usr/sbin/agetty;5ea3f3
mysqld      8019     mysql     5u       REG       253,1     0        528810   /tmp/ibTwxfsx (deleted)
mysqld      8019     mysql     6u       REG       253,1     0        528811   /tmp/ib0KEqPb (deleted)
mysqld      8019     mysql     7u       REG       253,1     0        528812   /tmp/ibX0VBc0 (deleted)
mysqld      8019     mysql     8u       REG       253,1     0        528813   /tmp/ib73Mfc9 (deleted)
mysqld      8019     mysql     14u      REG       253,1     0        528814   /tmp/ibQ0UESN (deleted)
mysqld      8019     8020     mysql     5u       REG       253,1     0        528810   /tmp/ibTwxfsx (deleted)
mysqld      8019     8020     mysql     6u       REG       253,1     0        528811   /tmp/ib0KEqPb (deleted)
mysqld      8019     8020     mysql     7u       REG       253,1     0        528812   /tmp/ibX0VBc0 (deleted)
mysqld      8019     8020     mysql     8u       REG       253,1     0        528813   /tmp/ib73Mfc9 (deleted)
mysqld      8019     8020     mysql     14u     REG       253,1     0        528814   /tmp/ibQ0UESN (deleted)
mysqld      8019     8021     mysql     5u       REG       253,1     0        528810   /tmp/ibTwxfsx (deleted)
mysqld      8019     8021     mysql     6u       REG       253,1     0        528811   /tmp/ib0KEqPb (deleted)
mysqld      8019     8021     mysql     7u       REG       253,1     0        528812   /tmp/ibX0VBc0 (deleted)
mysqld      8019     8021     mysql     8u       REG       253,1     0        528813   /tmp/ib73Mfc9 (deleted)
mysqld      8019     8021     mysql     14u     REG       253,1     0        528814   /tmp/ibQ0UESN (deleted)
mysqld      8019     8022     mysql     5u       REG       253,1     0        528810   /tmp/ibTwxfsx (deleted)
mysqld      8019     8022     mysql     6u       REG       253,1     0        528811   /tmp/ib0KEqPb (deleted)
mysqld      8019     8022     mysql     7u       REG       253,1     0        528812   /tmp/ibX0VBc0 (deleted)
mysqld      8019     8022     mysql     8u       REG       253,1     0        528813   /tmp/ib73Mfc9 (deleted)
mysqld      8019     8022     mysql     14u     REG       253,1     0        528814   /tmp/ibQ0UESN (deleted)
mysqld      8019     8023     mysql     5u       REG       253,1     0        528810   /tmp/ibTwxfsx (deleted)
mysqld      8019     8023     mysql     6u       REG       253,1     0        528811   /tmp/ib0KEqPb (deleted)
mysqld      8019     8023     mysql     7u       REG       253,1     0        528812   /tmp/ibX0VBc0 (deleted)
mysqld      8019     8023     mysql     8u       REG       253,1     0        528813   /tmp/ib73Mfc9 (deleted)
mysqld      8019     8023     mysql     14u     REG       253,1     0        528814   /tmp/ibQ0UESN (deleted)
mysqld      8019     8024     mysql     5u       REG       253,1     0        528810   /tmp/ibTwxfsx (deleted)
mysqld      8019     8024     mysql     6u       REG       253,1     0        528811   /tmp/ib0KEqPb (deleted)
mysqld      8019     8024     mysql     7u       REG       253,1     0        528812   /tmp/ibX0VBc0 (deleted)
mysqld      8019     8024     mysql     8u       REG       253,1     0        528813   /tmp/ib73Mfc9 (deleted)
mysqld      8019     8024     mysql     14u     REG       253,1     0        528814   /tmp/ibQ0UESN (deleted)
mysqld      8019     8025     mysql     5u       REG       253,1     0        528810   /tmp/ibTwxfsx (deleted)
mysqld      8019     8025     mysql     6u       REG       253,1     0        528811   /tmp/ib0KEqPb (deleted)
mysqld      8019     8025     mysql     7u       REG       253,1     0        528812   /tmp/ibX0VBc0 (deleted)
mysqld      8019     8025     mysql     8u       REG       253,1     0        528813   /tmp/ib73Mfc9 (deleted)
mysqld      8019     8025     mysql     14u     REG       253,1     0        528814   /tmp/ibQ0UESN (deleted)
mysqld      8019     8026     mysql     5u       REG       253,1     0        528810   /tmp/ibTwxfsx (deleted)
mysqld      8019     8026     mysql     6u       REG       253,1     0        528811   /tmp/ib0KEqPb (deleted)
mysqld      8019     8026     mysql     7u       REG       253,1     0        528812   /tmp/ibX0VBc0 (deleted)
mysqld      8019     8026     mysql     8u       REG       253,1     0        528813   /tmp/ib73Mfc9 (deleted)
mysqld      8019     8026     mysql     14u     REG       253,1     0        528814   /tmp/ibQ0UESN (deleted)
mysqld      8019     8027     mysql     5u       REG       253,1     0        528810   /tmp/ibTwxfsx (deleted)
```

Roughly calculate the total size of the deleted files based on the returned command output. If it is almost the same as the used space of the disk, the disk space may be used up by the deleted processes that have not been released.

4. Go to the location of a deleted file to check whether the file is still there.

```
ll /tmp/
```

Note that variable `/tmp/` in the command indicates the path of the deleted file.

5. If the file is not there, run the following command to terminate the process, or restart the server to release the used space.

```
kill -9 PID
```

Note that variable `PID` in the command indicates the process ID.

6. Check that the process is terminated.

```
lsuf | grep deleted
```

7. Check that the disk partition usage is no longer 100%.

```
df -h
```

## 17.8 Snapshot

## 17.8.1 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.

## 17.8.2 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*.

## 17.8.3 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.

## 17.8.4 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 an EVS 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.

## 17.8.5 Why Can't I Roll Back My Disk Data from a Snapshot?

Possible causes are as follows:

- Snapshot data can only be rolled back when the status of the snapshot's source disk is **Available** or **Rollback failed**. If the snapshot's source disk is **In-use**, detach the disk and then use the snapshot to roll back disk data. After the rollback succeeds, re-attach the disk.
- Snapshot data can only be rolled back to source EVS disks. Rollback to a different disk is not possible.
- If the snapshot is being created, it cannot be used to roll back disk data.
- A snapshot whose name starts with **autobk\_snapshot\_vbs\_**, **manualbk\_snapshot\_vbs\_**, **autobk\_snapshot\_csbs\_**, or **manualbk\_snapshot\_csbs\_** is automatically generated during backup. Such a snapshot can only be viewed. It cannot be used to roll back the disk data.



## 17.8.6 Can I Roll Back Data from a Snapshot After Reinstalling the OS or Formatting a 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.

## 17.8.7 How Is the Snapshot Size Calculated?

The first snapshot is a full snapshot, which backs up all data (data blocks) on the EVS disk at the time of the snapshot. Subsequent snapshots are incremental snapshots, which back up only changed data blocks since the last snapshot.

## 17.8.8 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.

## 17.8.9 Can I Perform Multiple Rollback Operations for a Snapshot?

Yes. You can use a snapshot to roll back data to the source disk multiple times as needed.

## 17.8.10 Can I Replicate Snapshots to Other Regions or Accounts?

No. Snapshots cannot be replicated to other regions or accounts.

## 17.8.11 Why Can't I Find My Snapshot?

Possible causes are:

- You have manually deleted the snapshot.
- You have deleted the snapshot's source disk. Then, all the snapshots created for this disk will be deleted automatically.
- You have reinstalled or changed the server OS. System disk snapshots will be automatically deleted after a server OS has been reinstalled or changed.

## 17.8.12 Can I Use a Snapshot to Create a Disk and Roll Back Disk Data at Almost the Same Time?

The following two scenarios are supported:

- You can first use a snapshot to create a disk and then use the snapshot to roll back data of its source disk.
- You can first use a snapshot to roll back data of its source disk and then use the snapshot to create a disk.

## 17.8.13 Can I Modify Data in My EVS Snapshot?

No. EVS disk snapshots are read-only. After snapshots are created, data in the snapshots cannot be modified.

## 17.9 Performance

### 17.9.1 How Do 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 Server 2019 Standard 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**

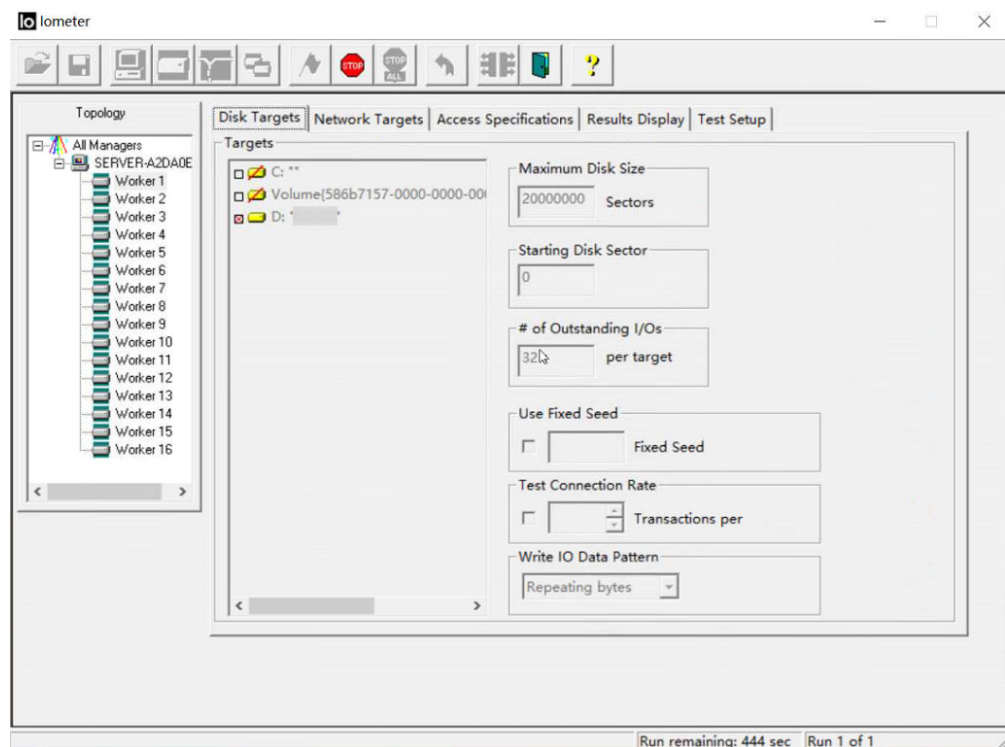
If you delete the partition and select another start sector number for 4-KiB alignment, you will lose all the data on that partition.

**Step 4** Use Iometer to test the disk performance. For details, see the Iometer product documentation.

When the disk IOPS and throughput are tested, the parameter settings for Iometer and fio are the same. For details, see [Table 17-11](#).


The following example uses Iometer to test the disk performance.

1. Set the workflow.

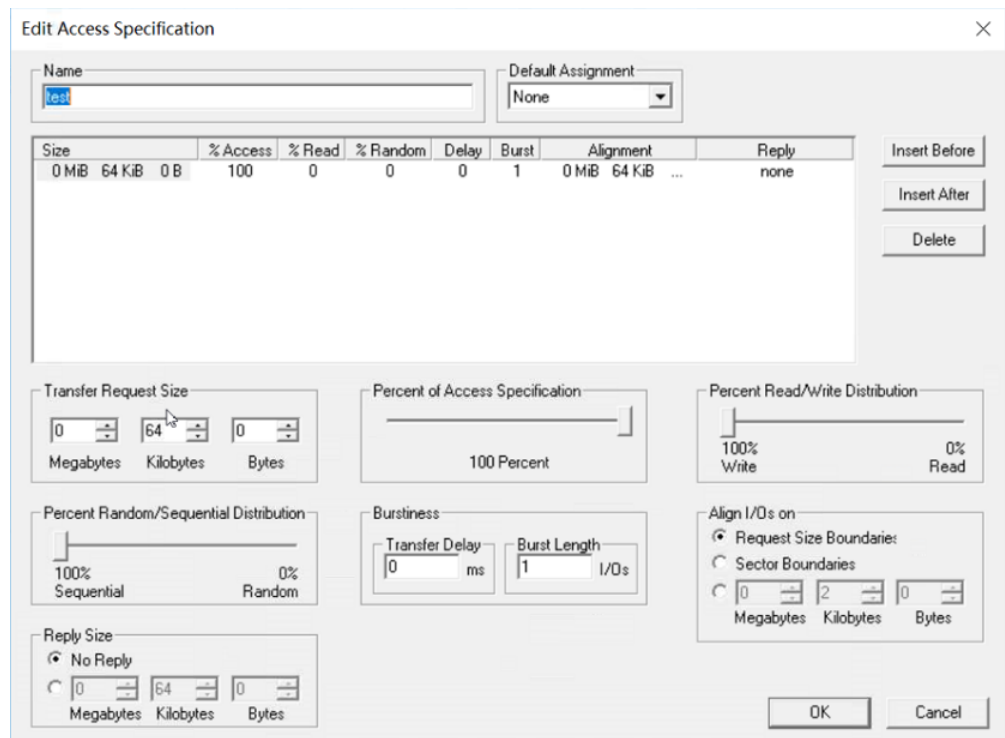


2. Set the test run time.

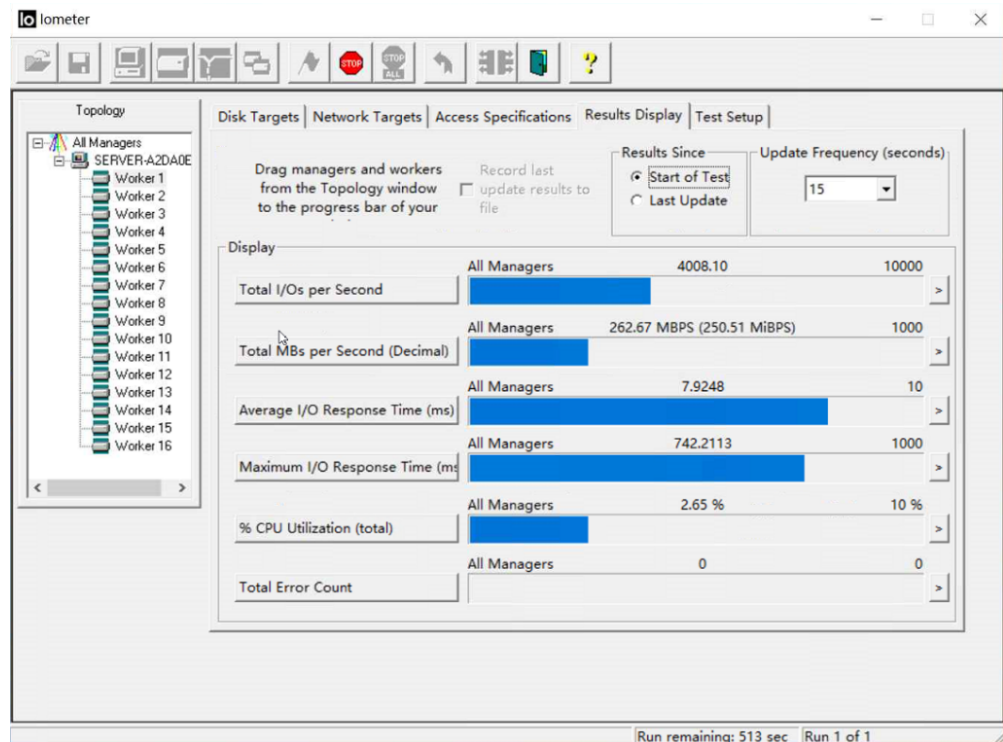
In this example, the test run time is set to 10 minutes, with 60 seconds ramp up time. Disk performance is tested after the writes are stable.



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.

**Step 1** Log in to the server and switch to user **root**.

**Step 2** Install the performance test tool **fiio**.

```
yum install fiio
```

**Step 3** 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 4](#).
- 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.

---

**Step 4** 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 single-queue, 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 17-11** lists the fio test parameters.

**Table 17-11** Parameter description

Parameter	Description
direct	Defines whether direct I/O is used. <ul style="list-style-type: none"><li>- Set to <b>0</b>: buffered I/O is used.</li><li>- Set to <b>1</b>: direct I/O is used.</li></ul>
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 = <code>iodepth x numjobs</code> Examples: <ul style="list-style-type: none"><li>- If there is a single thread and <b>-iodepth=32</b>, the I/O queue depth of this thread is 32 and the total concurrent I/Os of fio is 32 (32 x 1).</li><li>- If there are three threads and <b>-iodepth=32</b>, the I/O queue depth of each thread is 32 and the total concurrent I/Os of fio is 96 (32 x 3).</li></ul>
rw	Defines the test read/write policy. <ul style="list-style-type: none"><li>- <b>randread</b>: random read</li><li>- <b>randwrite</b>: random write</li><li>- <b>read</b>: sequential read</li><li>- <b>write</b>: sequential write</li><li>- <b>randrw</b>: mixed random read/write</li></ul>

Parameter	Description
ioengine	<p>Defines how fio delivers the I/O request (synchronously or asynchronously).</p> <ul style="list-style-type: none"> <li>- 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.</li> <li>- 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.</li> </ul>
bs	<p>Defines the I/O block size. The unit can be KiB, Kb, MiB, and Mb, and the default value is 4 KiB.</p>
size	<p>Defines the amount of data processed by the test I/Os. If parameters, such as <b>runtime</b>, are not specified, the test ends when fio has processed all the specified data amount.</p> <p>The value can be a number with a unit or percentage. A number with a unit indicates the read/write data amount, for example <b>size=10G</b>, 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 <b>size=20%</b>, indicating the read/write data amount takes 20% of the total file space.</p>
numjobs	<p>Defines the number of concurrent threads.</p>
runtime	<p>Defines the test time.</p> <p>If this parameter is not specified, the test ends until the specified data amount is processed by the block size defined using parameter <b>size</b>.</p>
group_reporting	<p>Defines the test result display mode. The parameter value displays the statistics on a single thread instead of that on all jobs.</p>
filename	<p>Defines the name of the test file or device.</p> <ul style="list-style-type: none"> <li>- If a file is specified, the performance of the file system is tested. Example: <b>-filename=/opt/fiotest/fiotest.txt</b></li> <li>- If a device name is specified, the performance of the raw disk is tested. Example: <b>-filename=/dev/vdb</b></li> </ul> <p><b>NOTICE</b> If the test is performed on a disk already has partitions and file systems created as well as data on it, user parameter <b>filename</b> to specify a file so that the original file system is not damaged and the data is not overwritten.</p>
name	<p>Defines the test task name.</p>



----End

## 17.9.2 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.

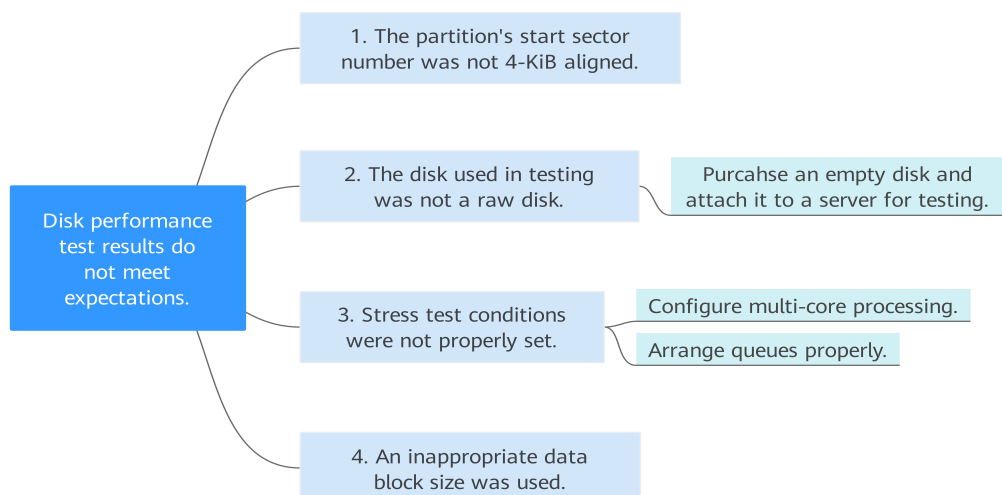
---

#### NOTICE

Some operations may result in data loss. It is recommended that you use raw disks for performance test.

---

Figure 17-64 Troubleshooting



**Table 17-12** Troubleshooting

Possible Cause	Solution
The partition's start sector number is not 4-KiB aligned.	Go to <a href="#">Check Whether Partition's Start Sector Number Is 4-KiB Aligned</a> . 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. <ul style="list-style-type: none"> <li>When testing the disk IOPS, set the data block size to a small value, for example, 4 KiB.</li> <li>When testing the disk throughput, set the data block size to a large value, for example, 1024 KiB.</li> </ul>

## 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.

----End

## 17.9.3 How Can I Handle a Slowdown in Disk Read/Write Speed or Increased I/Os?

### 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 **iotstat -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 to encounter a performance bottleneck.

### Solution

If your disk contains important data, create a new disk from the disk backup, so you do not lose any of the disk data. The procedure is as follows:

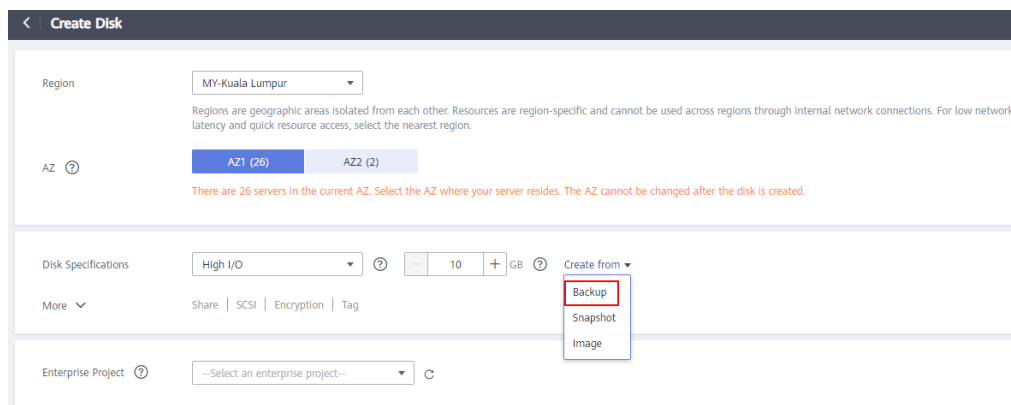
1. Create a backup for the disk.

**Figure 17-65** 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 17-66 Create from backup



## 17.9.4 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:  $\text{Disk IOPS} = \text{Min. (Max. IOPS, Min. IOPS} + \text{IOPS per GiB} \times \text{Disk size)}$ . For a 500 GiB disk, its IOPS is calculated as follows:  $\text{Disk IOPS} = \text{Min. (33,000, 1,500} + 50 \times 500) = 26,500$

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 (26,500), the queue depth should reach 26.

## 17.10 Sharing

### 17.10.1 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.

### 17.10.2 How Many Servers Can I Attach a Shared Disk to?

A shared disk can be attached to up to 16 servers.

### 17.10.3 How Can I Attach a Shared Disk to Multiple Servers?

A shared disk can be attached to multiple servers on the console. You can choose to attach it to servers one by one or in a batch.

For details, see [Attaching a Shared Disk](#).

### 17.10.4 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.

### 17.10.5 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.

## 17.11 Backup

### 17.11.1 Do I Need to Stop the Server Before Performing a Disk Backup?

No. You can back up disks when the server is 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. For applications that require strong consistency, such as databases and email systems, you are advised to enable application-consistent backup.

## 17.11.2 Can I Back Up and Restore My EVS Disk to a Different Region?

EVS disks already attached to ECS support cross-region backup and restoration. Such disks can be backed up together with their ECS using CBR cloud server backup, which supports cross-region replication. You can replicate server backups to your desired region and use the replicated backups to create images and provision servers.

EVS disks that have not been attached to ECS do not support cross-region backup and restoration.

## 17.11.3 How Do I Check the Backup Data of My Disk?

1. Create a new disk from a disk backup. For details, see section "Creating a Disk from a Cloud Disk Backup" 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 check the disk data.

# A Change History

Released On	Description
2024-03-20	This issue is the third official release. Updated the following content: Supplemented contents related to EVS encryption.
2023-03-13	This issue is the second official release. Updated the following content: Supplemented related contents.
2022-08-08	This issue is the first official release.

Released On	Description
2024-06-20	This issue is the sixth official release, which incorporates the following change: <ul style="list-style-type: none"> <li>Optimized the procedure in section "Create an EVS Disk."</li> <li>Added support for organization tag policies and modified the description in section "Tag Overview."</li> </ul>
2023-06-12	This issue is the fifth official release, which incorporates the following change: Added screenshots in disk purchase and attachment sections.
2023-03-23	This issue is the fourth official release, which incorporates the following change: Changed Key Management Service (KMS) to Data Encryption Workshop (DEW) throughout the document.
2023-03-13	This issue is the third official release, which incorporates the following change: Optimized some descriptions.

Released On	Description
2022-08-08	This issue is the second official release, which incorporates the following change: <ul style="list-style-type: none"><li data-bbox="603 376 1018 416">• Optimized some descriptions.</li></ul>
2020-08-15	This issue is the first official release.