

Elastic Volume Service(EVS)

User Guide

Issue	01
Date	2023-11-14



Copyright © Huawei Technologies Co., Ltd. 2024. All rights reserved.

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

Trademarks and Permissions



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

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

Notice

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

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

Security Declaration

Vulnerability

Huawei's regulations on product vulnerability management are subject to the *Vul. Response Process*. For details about this process, visit the following web page:

<https://www.huawei.com/en/psirt/vul-response-process>

For vulnerability information, enterprise customers can visit the following web page:

<https://securitybulletin.huawei.com/enterprise/en/security-advisory>

Contents

1 Permissions Management.....	1
1.1 Creating a User and Granting EVS Permissions.....	1
1.2 EVS Custom Policies.....	2
2 Purchasing and Using an EVS Disk.....	5
3 Disk Capacity Expansion.....	8
3.1 Expansion Overview.....	8
3.2 Expanding Capacity for an In-use EVS Disk.....	11
3.3 Expanding Capacity for an Available EVS Disk.....	13
3.4 Extending Disk Partitions and File Systems (Windows Server 2008).....	15
3.5 Extending Disk Partitions and File Systems (Windows Server 2016).....	40
3.6 Extending Disk Partitions and File Systems (Linux).....	65
3.6.1 Partition and File System Extension Preparations (Linux).....	65
3.6.2 Extending Partitions and File Systems for System Disks (Linux).....	69
3.6.3 Extending Partitions and File Systems for Data Disks (Linux).....	77
4 Changing the EVS Disk Type (OBT).....	90
5 Detaching an EVS Disk.....	93
5.1 Detaching a System Disk.....	93
5.2 Detaching a Data Disk.....	94
6 Attaching an Existing Disk.....	97
6.1 Attaching an Existing System Disk.....	97
6.2 Attaching an Existing Non-Shared Disk.....	99
6.3 Attaching an Existing Shared Disk.....	100
7 Deleting EVS Disks.....	103
8 Unsubscribing from Yearly/Monthly EVS Disks.....	105
9 Managing EVS Recycle Bin.....	108
9.1 Recycle Bin Overview.....	108
9.2 Enabling the Recycle Bin.....	109
9.3 Disabling the Recycle Bin.....	110
9.4 Configuring a Recycle Bin Policy.....	111
9.5 Recovering Disks from the Recycle Bin.....	112

9.6 Permanently Deleting Disks from the Recycle Bin.....	113
10 Viewing EVS Disk Details.....	114
11 Managing Encrypted EVS Disks.....	116
12 Managing Shared EVS Disks.....	120
13 Managing EVS Backups.....	122
14 Managing EVS Snapshots (OBT).....	125
14.1 Snapshot Overview (OBT).....	125
14.2 Creating a Snapshot (OBT).....	126
14.3 Deleting a Snapshot (OBT).....	129
14.4 Viewing Details of a Snapshot (OBT).....	129
14.5 Rolling Back Data from a Snapshot (OBT).....	130
14.6 Creating an EVS Disk from a Snapshot (OBT).....	131
15 Managing EVS Transfers.....	133
16 Managing a Tag.....	136
16.1 Tag Overview.....	136
16.2 Adding a Tag.....	136
16.3 Modifying a Tag.....	137
16.4 Deleting a Tag.....	138
16.5 Searching for Disks by Tag.....	138
17 Changing EVS Disk Name.....	140
18 Viewing EVS Monitoring Data.....	141
19 Viewing EVS Monitoring Data (Agent Installed and Simplified Monitoring Metrics Used).....	145
20 Auditing.....	149
21 Managing Quotas.....	150
21.1 Querying EVS Resource Quotas.....	150
21.2 Increasing EVS Resource Quotas.....	151
A Appendix.....	153
A.1 EVS Disk Status.....	153
A.2 EVS Snapshot Status.....	155
B Change History.....	157

1 Permissions Management

1.1 Creating a User and Granting EVS Permissions

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

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

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

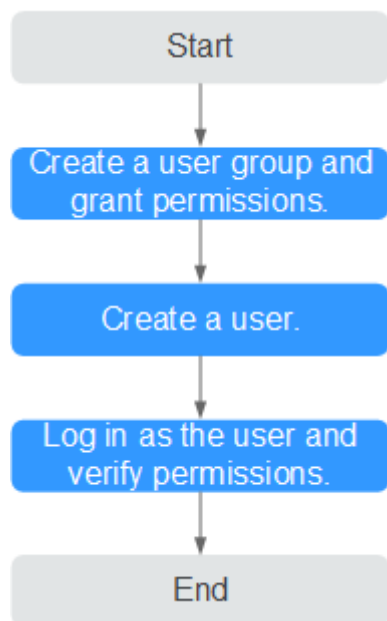
This section describes the procedure for granting permissions (see [Figure 1-1](#)).

Prerequisites

Learn about the permissions (see [EVS Permissions](#)) supported by EVS and choose policies or roles according to your requirements.

Process Flow

Figure 1-1 Process for granting EVS permissions



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

1.2 EVS Custom Policies

Custom policies can be created to supplement the system-defined policies of EVS. For the actions supported for custom policies, see [Permissions Policies and Supported Actions](#).

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 [Creating a Custom Policy](#). 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"
      ]
    }
  ]
}
```

- Example 2: Create a custom policy containing multiple actions.

A custom policy can contain the actions of one or multiple services.

Example policy containing multiple actions:

```
{
  "Version": "5.0",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "evs:volumes:create",

```

```
        "evs:volumes:list"
      ],
    },
    {
      "Effect": "Allow",
      "Action": [
        "evs:volumes:create",
        "evs:volumes:list",
        "evs:volumes:get",
        "evs:types:get",
        "evs:quotas:get",
        "evs:volumes:use"
      ]
    },
    {
      "Effect": "Allow",
      "Action": [
        "ecs:cloudServerFlavors:get",
        "ecs:cloudServers:listServersDetails",
        "ecs:cloudServers:showServer",
        "ecs:cloudServers:attach",
        "ecs:cloudServers:detachVolume"
      ]
    }
  ]
}
```

2 Purchasing and Using an EVS Disk


EVS disks can be attached to servers and be used as system disks and data disks. For details, see [Table 2-1](#).

Table 2-1 Method of purchase

Disk	Description	Method
System disk	System disks are purchased along with servers and cannot be purchased separately.	<ul style="list-style-type: none">• Purchasing an ECS• Creating a BMS
Data disk	Data disks can be purchased along with servers or separately.	<ul style="list-style-type: none">• Purchasing an ECS• Creating a BMS• Purchase an EVS Disk

[Figure 2-1](#) shows how to purchase a data disk separately.

Figure 2-1 Process overview



1. **Preparations:** Register with Huawei Cloud, and top up your account. For details, see [Preparations](#).
2. **Purchase a disk.** Configure the disk parameters, including the disk type, size, name, and other information. For details, see [Purchase an EVS Disk](#).
3. **Attach the data disk.** Attach the separately purchased disk to a server. For details, see the following sections:
 - [Attaching a Non-Shared Disk](#)
 - [Attaching a Shared Disk](#)
4. **Initialize the data disk.** Log in to the server and initialize the data disk before using it. For details about how to initialize the disk, see the following sections:
 - [Introduction to Data Disk Initialization Scenarios and Partition Styles](#)
 - Windows
 - [Initializing a Windows Data Disk \(Windows Server 2008\)](#)
 - [Initializing a Windows Data Disk \(Windows Server 2019\)](#)
 - [Initializing a Windows Data Disk Larger Than 2 TiB \(Windows Server 2008\)](#)
 - [Initializing a Windows Data Disk Larger Than 2 TiB \(Windows Server 2012\)](#)
 - Linux
 - [Initializing a Linux Data Disk \(fdisk\)](#)

- [Initializing a Linux Data Disk \(parted\)](#)
- [Initializing a Linux Data Disk Larger Than 2 TiB \(parted\)](#)

3 Disk Capacity Expansion

3.1 Expansion Overview

What Is Capacity Expansion?

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

Both system disks and data disks can be expanded. A system disk can be expanded to up to 1 TiB, and a data disk to 32 TiB. Disk capacity can only be expanded, not reduced.

NOTE

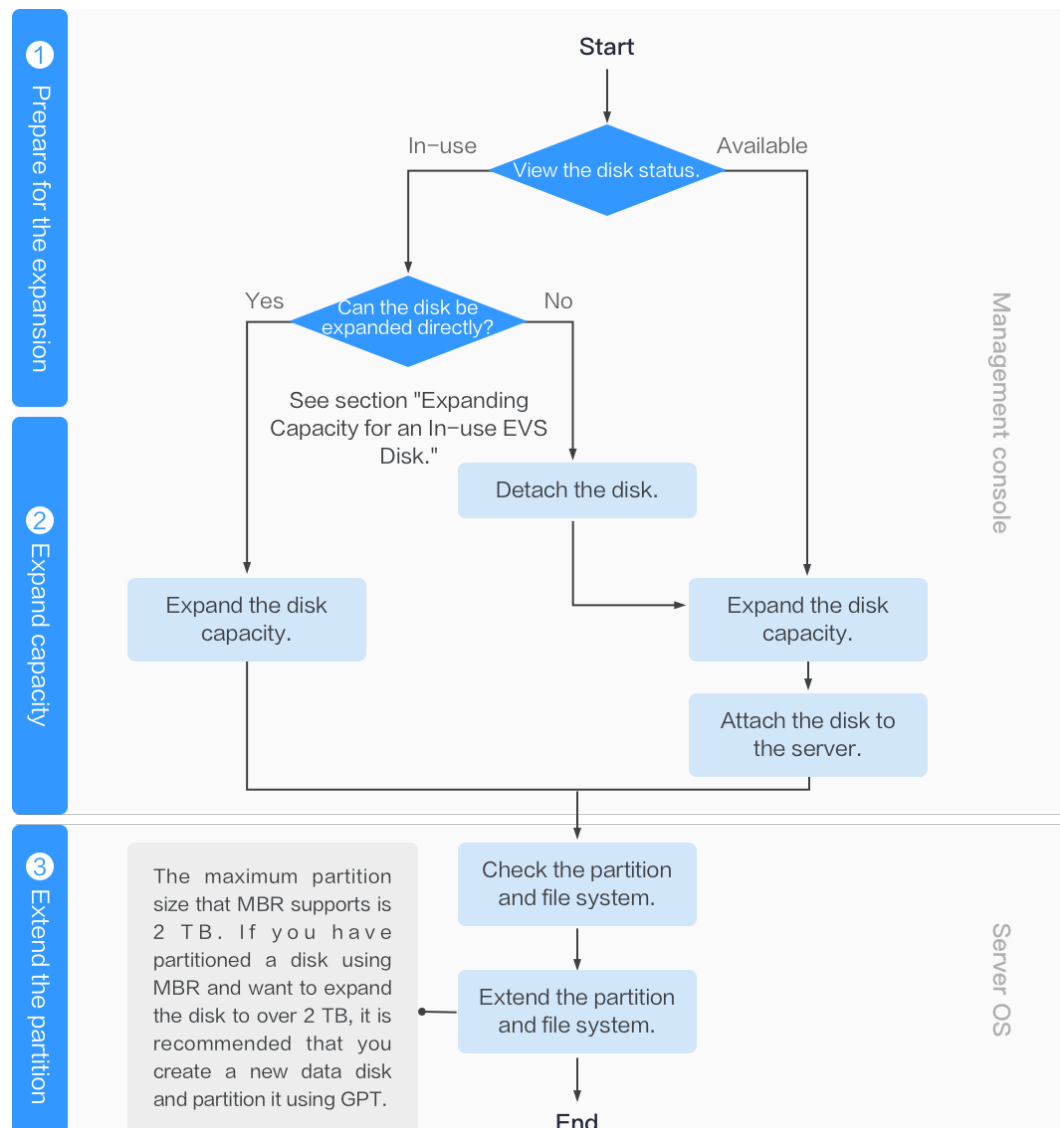
If you detach a system disk and then attach it to an ECS as a data disk, the maximum capacity of this disk is still 1 TiB.

How to Expand the Disk Capacity?

A capacity expansion operation includes the following steps:

1. [Expand the disk capacity on the management console.](#)
2. [Log in to the server and extend the disk partition and file system.](#)

Figure 3-1 Capacity expansion procedure



Expand the Disk Capacity on the Management Console

Choose a proper expansion method based on the disk status. For how to view the disk status, see [Viewing EVS Disk Details](#).

- For an In-use disk:
 - The disk has been attached to a server. Check whether the disk can be expanded in the In-use state by referring to [Constraints](#).
 - If yes, expand the disk capacity according to [Expanding Capacity for an In-use EVS Disk](#).
 - If no, detach the disk. Then, expand the disk capacity according to [Expanding Capacity for an Available EVS Disk](#).
 - For an Available disk:
 - The disk has not been attached to any server and can be directly expanded by referring to [Expanding Capacity for an Available EVS Disk](#).
- A shared disk can be expanded only when its status is **Available**.

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

After the disk has been expanded on the management console, only the disk storage capacity is enlarged, but its additional space cannot be used directly. You must log in to the server and extend the disk partition and file system. For details, see [Table 3-1](#).

Table 3-1 Extending the disk partition and file system

Capacity After Expansion	Extend Disk Partition and File System
Disk capacity \leq 2 TiB	<ul style="list-style-type: none">Windows: Extending Disk Partitions and File Systems (Windows Server 2008)Linux: Partition and File System Extension Preparations (Linux)
Disk capacity > 2 TiB	<ul style="list-style-type: none">GPT partition style: Extending Disk Partitions and File Systems (Windows Server 2008) or Partition and File System Extension Preparations (Linux)MBR partition style: Not supported The maximum disk capacity that MBR supports is 2 TiB, and the disk space exceeding 2 TiB cannot be used. If your disk uses MBR and you need to expand the disk capacity to over 2 TiB, change the partition style from MBR to GPT. Ensure that the disk data has been backed up before changing the partition style because services will be interrupted and data on the disk will be deleted during this change.

NOTE

If the server is stopped during the expansion, the additional space of a Windows system disk, Windows data disk, 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, extend the partition and file system according to the preceding steps.

Related Operations

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

3.2 Expanding Capacity for an In-use EVS Disk

Scenarios

This section describes how to expand the capacity of an In-use EVS disk on the management console. The In-use status indicates that the disk has been attached to a server. You do not need to detach the disk when expanding an In-use disk. For how to view the disk status, see [Viewing EVS Disk Details](#).


Constraints

- Disk capacity can be expanded, but cannot be reduced.
- When expanding an In-use disk, the server attached with this disk must be in the **Running** or **Stopped** state.
- A shared disk in the **In-use** state cannot be expanded. To expand such a disk, you must detach it from all its servers, wait until its status changes to **Available**, and then expand its capacity. For more information, see [Expanding Capacity for an Available EVS Disk](#).
- Only some server OSs support capacity expansion of In-use disks. For servers without such support, detach the disk and then expand its capacity. Otherwise, you may need to stop and then start the server after the expansion to make the additional space available.

Perform the following operations to check whether your server OS supports capacity expansion of In-use disks:

- a. Check your server image. Both public images and private images listed on the console support the capacity expansion of In-use disks.

How to view: Log in to the management console. In the navigation pane

on the left, click  and choose **Compute > Image Management Service**. On the **Public Images** tab, view the images of the **ECS image** type.

- b. If your server OS does not appear in the image list, check whether it is included in [Table 3-2](#).

If yes, you can expand the disk. Otherwise, you must detach the disk and then expand its capacity. For details, see [Expanding Capacity for an Available EVS Disk](#).

Table 3-2 OSs that support the capacity expansion of In-use disks

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

OS	Version
Fedora	24 64-bit or later
SUSE 12	SUSE Linux Enterprise Server 12 64-bit or later
SUSE 11	SUSE Linux Enterprise Server 11 SP4 64-bit
OpenSUSE	42.1 64-bit or later
Oracle Linux Server release 7	7.2 64-bit or later
Oracle Linux Server release 6	6.7 64-bit or later
Ubuntu Server	14.04 64-bit or later
Red Hat Enterprise Linux 7	7.3 64-bit
Red Hat Enterprise Linux 6	6.8 64-bit
EulerOS	2.2 64-bit or later
Huawei Cloud EulerOS	1.1 or later
Windows Server 2016	Windows Server 2016 R2 Enterprise 64-bit
Windows Server 2012	Windows Server 2012 R2 Standard 64-bit
Windows Server 2008	Windows Server 2008 R2 Enterprise 64-bit

Prerequisites

Disk data has been backed up using CBR or snapshots. For details about backups, see [Managing EVS Backups](#). For details about snapshots, see [Creating a Snapshot \(OBT\)](#).

Procedure

Step 1 Log in to the [management console](#).

Step 2 Under **Storage**, click **Elastic Volume Service**.

The disk list page is displayed.

Step 3 Choose a way to expand the disk by determining whether you want to check server information first.

- If yes, perform the following procedure:
 - a. In the disk list, click the name of the to-be-expanded disk.
The disk details page is displayed.
 - b. Click the **Servers** tab to view the server where the target disk has been attached.

- c. Click **Expand Capacity** in the upper right corner of the page.
The expansion page is displayed.
- If no, perform the following procedure:
 - a. In the disk list, locate the row that contains the target disk and click **Expand Capacity** in the **Operation** column.
The expansion page is displayed.

Step 4 Set the **New Capacity** parameter and click **Next**.

Step 5 On the **Details** page, check the disk details.

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

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

Step 6 In the displayed dialog box, click **OK**.

Step 7 In the disk list, view the capacity of the target disk.

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

 **NOTE**

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

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

The operations vary depending on the server OS.

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

----End

3.3 Expanding Capacity for an Available EVS Disk

Scenarios

This section describes how to expand the capacity of an Available EVS disk on the management console. The Available status indicates that the disk has not been attached to any server. For how to view the disk status, see [Viewing EVS Disk Details](#).

Constraints

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

Prerequisites

Disk data has been backed up using CBR or snapshots. For details about backups, see [Managing EVS Backups](#). For details about snapshots, see [Creating a Snapshot \(OBT\)](#).

Procedure

Step 1 Log in to the [management console](#).

Step 2 Under **Storage**, click **Elastic Volume Service**.

The disk list page is displayed.

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

The expansion page is displayed.

Step 4 Set the **New Capacity** parameter and click **Next**.

Step 5 On the **Details** page, check the disk details.

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

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

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

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

NOTE

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

Step 7 Attach the disk to the server. For details, see the following sections:

- [Attaching an Existing System Disk](#)
- [Attaching an Existing Non-Shared Disk](#)
- [Attaching an Existing Shared Disk](#)

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

The operations vary depending on the server OS.

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

----End

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

Scenarios

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

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

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

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

- For a system disk:
 - If volume (C:) already exists, you can add the additional space to volume (C:) and use it as a system volume. For details, see [System Disk: Add 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

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

Prerequisites

- You have expanded the disk capacity and attached the disk to a server on the management console. For details, see [Expanding Capacity for an In-use EVS Disk](#) or [Expanding Capacity for an Available EVS Disk](#).
- You have logged in to the server.
 - For how to log in to an ECS, see [Logging In to an ECS](#).
 - For how to log in to a BMS, see [Logging In to a BMS](#).

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

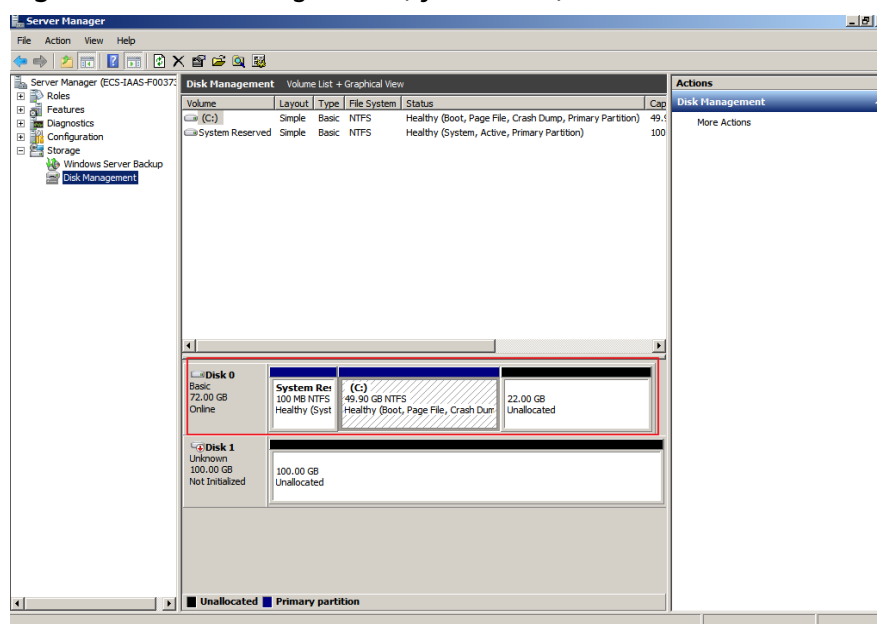
Step 1 On the desktop of the server, right-click **Computer** and choose **Manage** from the shortcut menu.

The **Server Manager** window is displayed.

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

The **Disk Management** window is displayed.

Figure 3-2 Disk Management (system disk)



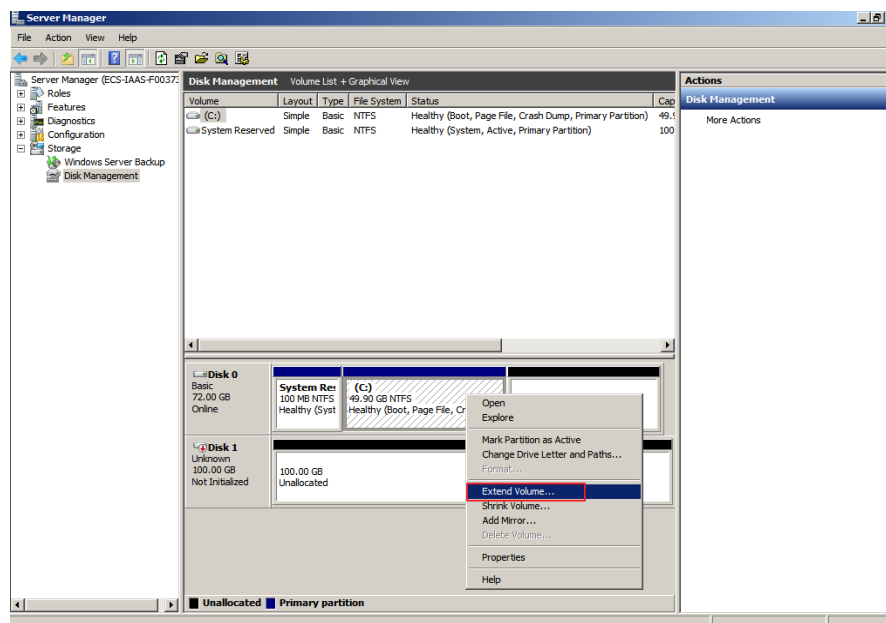
NOTE

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.

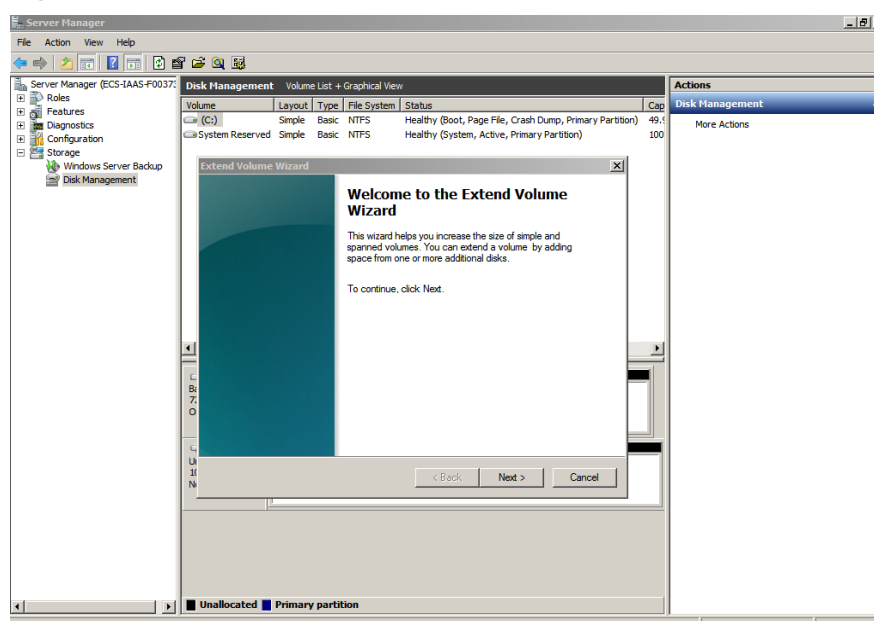
Step 4 Right-click the target volume and choose **Extend Volume**.

Figure 3-3 Choosing **Extend Volume**



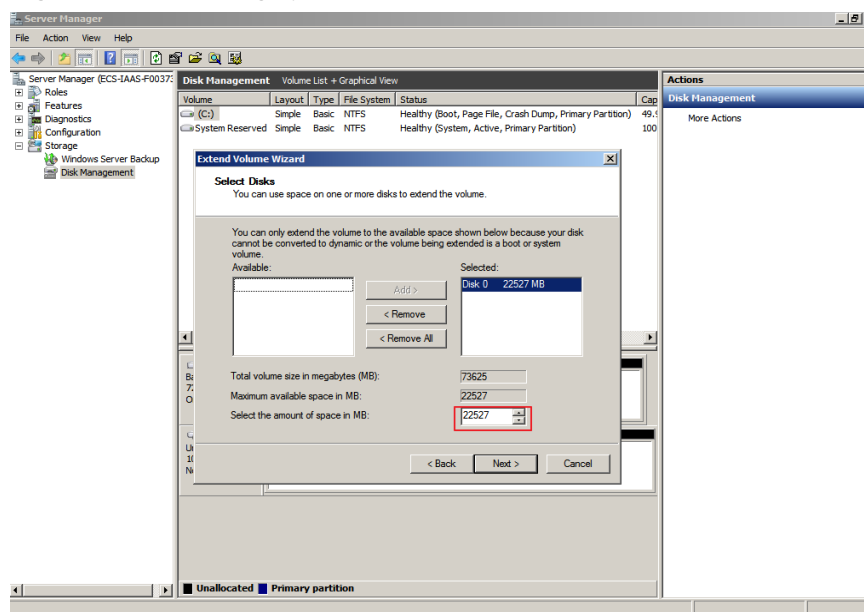
Step 5 On the displayed **Extend Volume Wizard** window, click **Next**.

Figure 3-4 Extend Volume Wizard



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

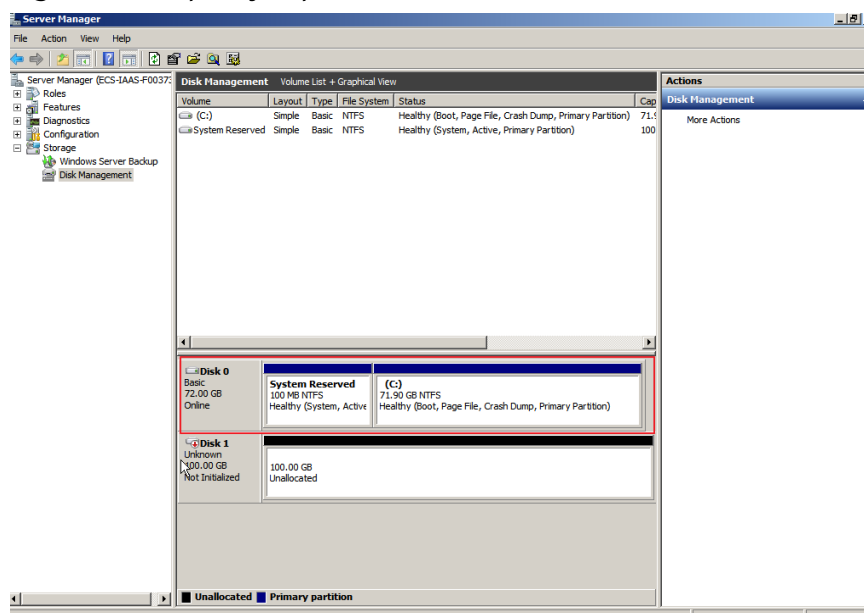
Figure 3-5 Selecting space



Step 7 Click **Finish**.

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

Figure 3-6 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 management console. The following procedure describes how to use this 60 GiB to

create a new volume, for example volume (F:), on the server. After the operation is complete, new volume (F:) has 60 GiB of capacity and can be used as a data volume.

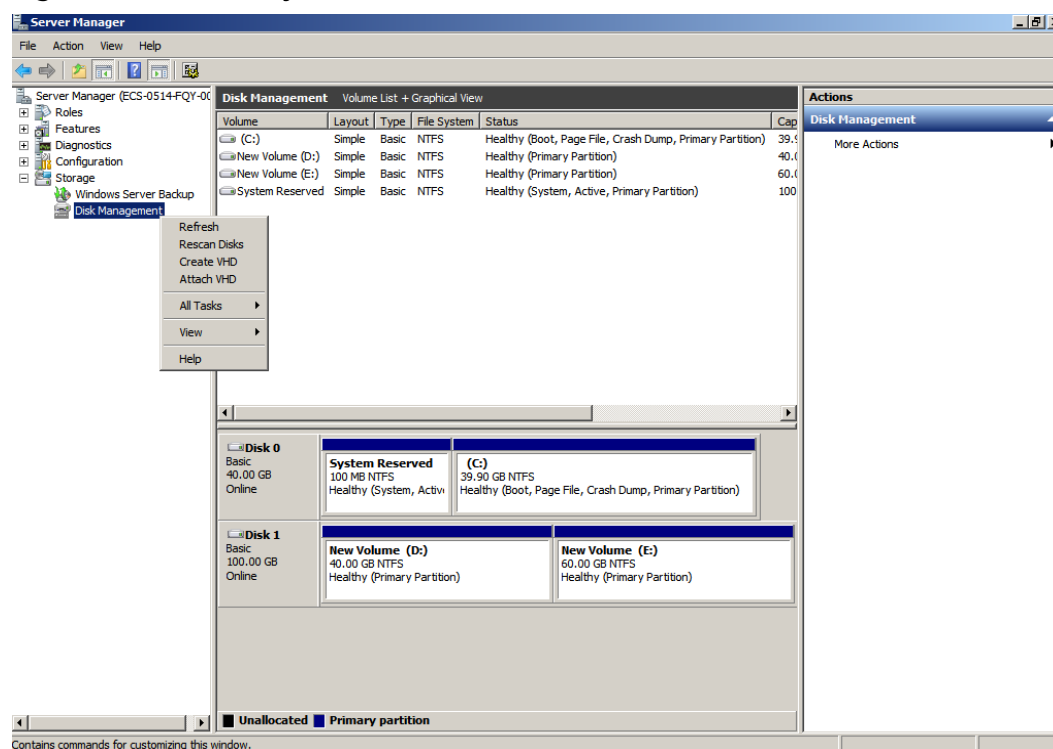
Step 1 On the desktop of the server, right-click **Computer** and choose **Manage** from the shortcut menu.

The **Server Manager** window is displayed.

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

The **Disk Management** window is displayed.

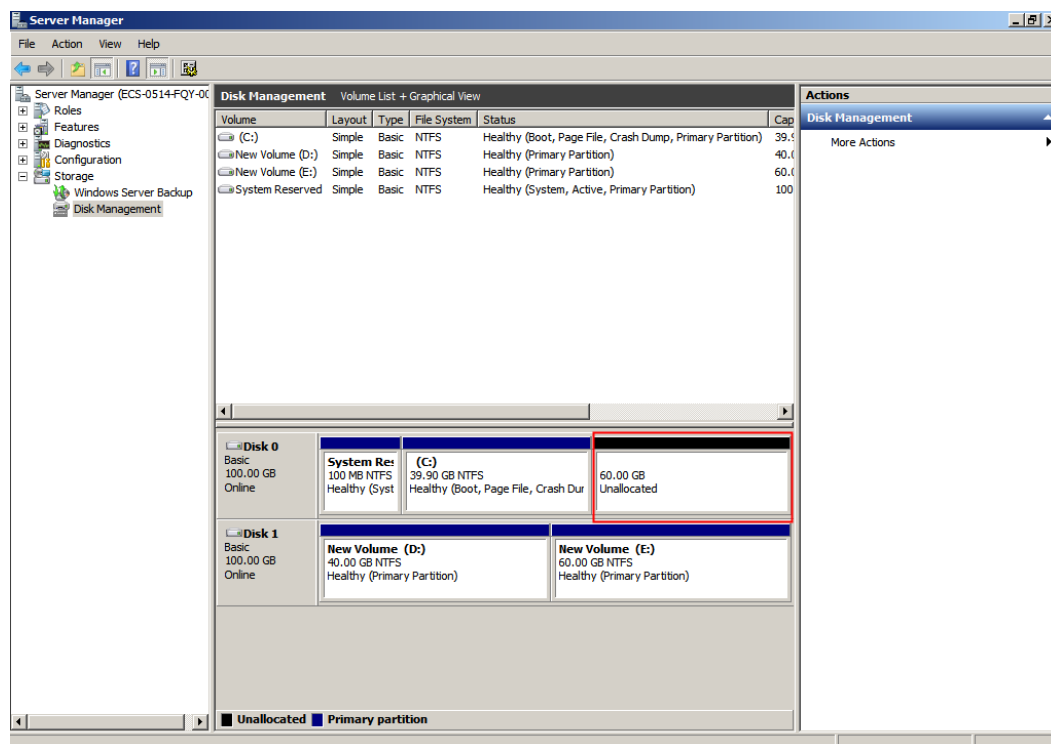
Figure 3-7 Refresh (system disk)



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

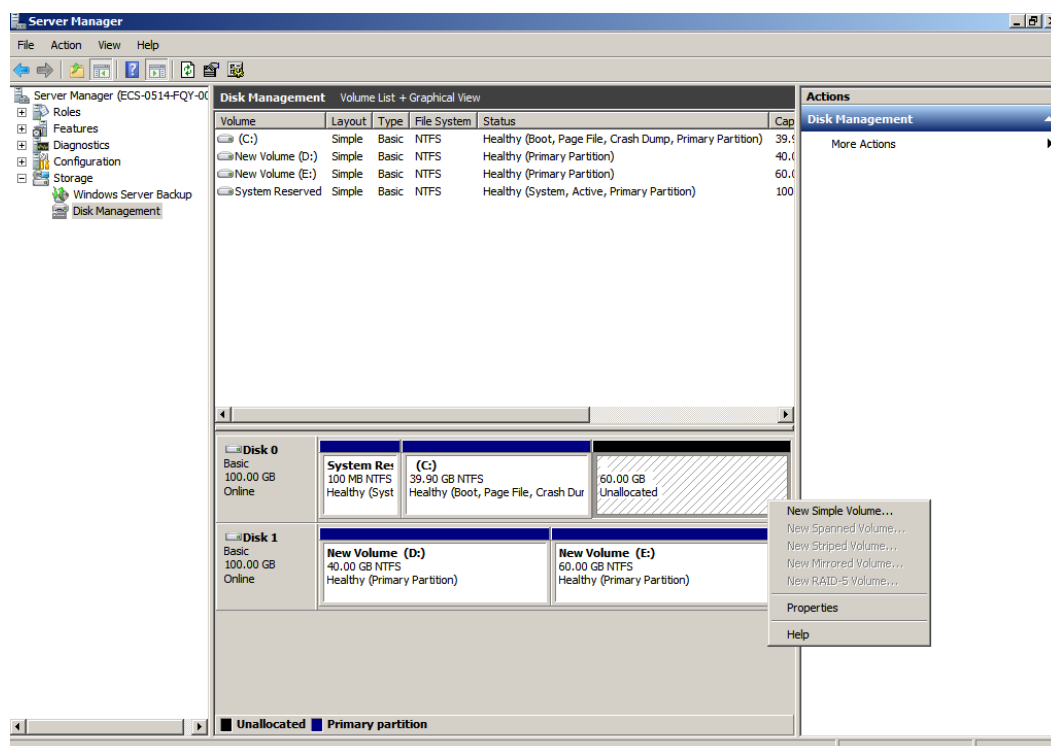
After the refresh, the additional space is displayed in the right area and is unallocated.

Figure 3-8 Unallocated disk space (system disk)



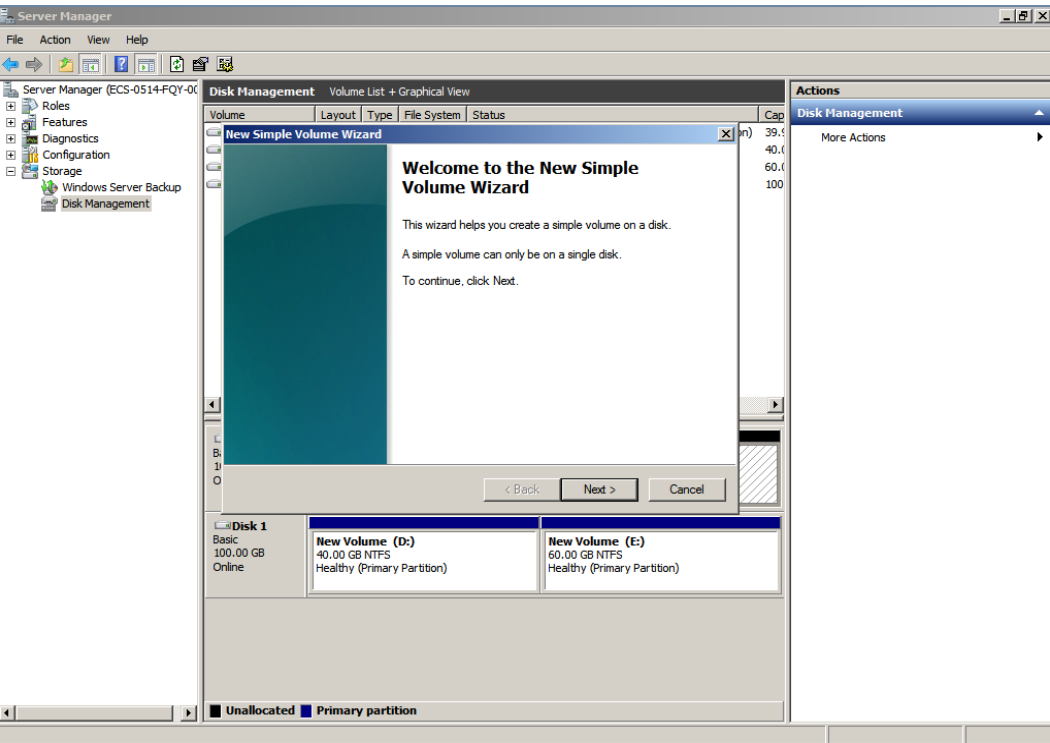
Step 4 In the **Unallocated** area of **Disk 0**, right-click the blank area and choose **New Simple Volume**.

Figure 3-9 New Simple Volume (system disk)



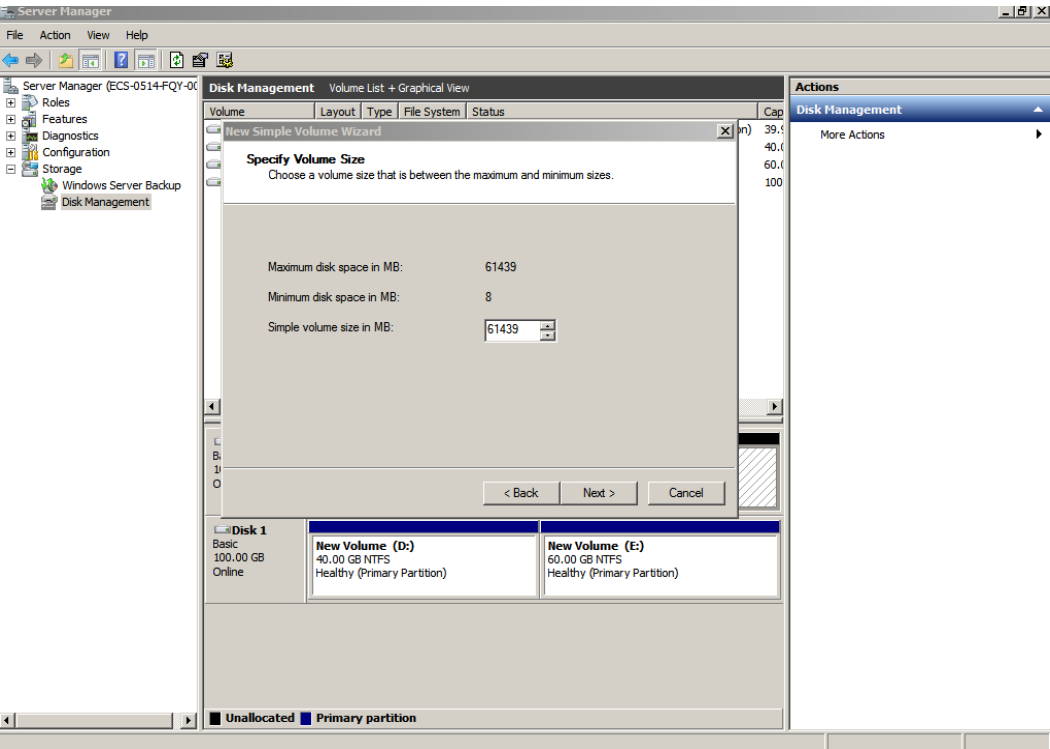
Step 5 On the displayed **New Simple Volume Wizard** window, click **Next**.

Figure 3-10 New Simple Volume Wizard (system disk)



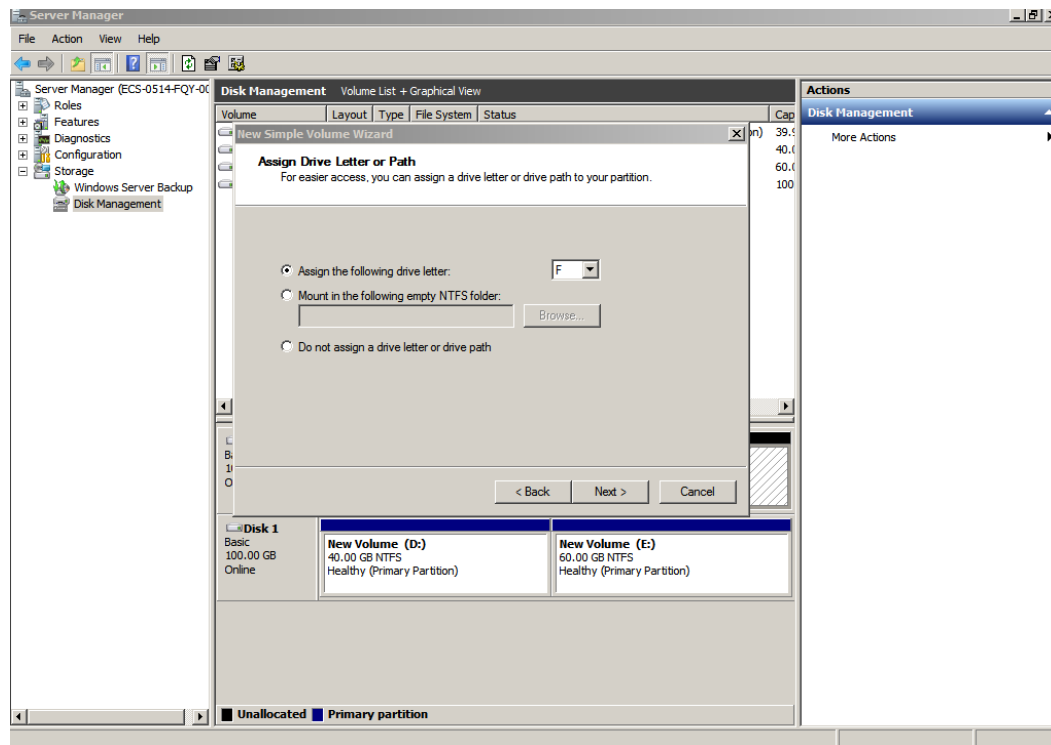
Step 6 On the displayed **Specify Volume Size** page, set **Simple volume size in MB** and click **Next**. In this example, the default size is used.

Figure 3-11 Specify Volume Size (system disk)



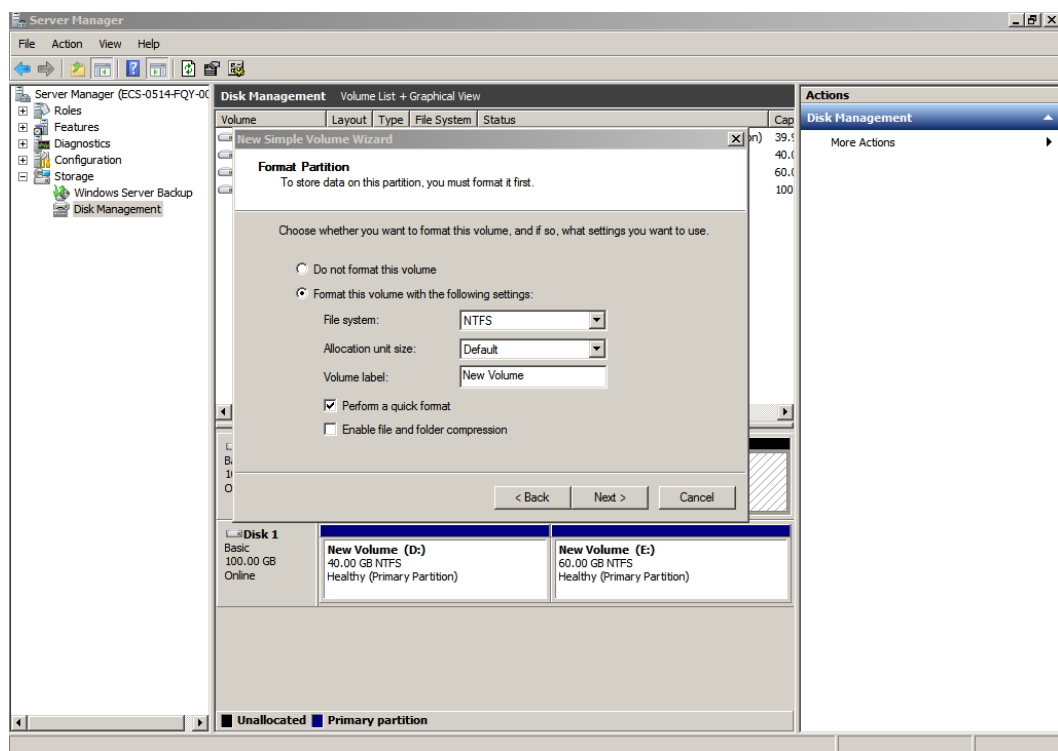
- Step 7** On the displayed **Assign Drive Letter and Path** page, click **Assign the following drive letter**, select a drive letter, and click **Next**. In this example, drive letter **F** is selected.

Figure 3-12 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 3-13 Format Partition (system disk)



Step 9 Click **Finish**.

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

Figure 3-14 Completing the New Simple Volume Wizard (new volume F:)

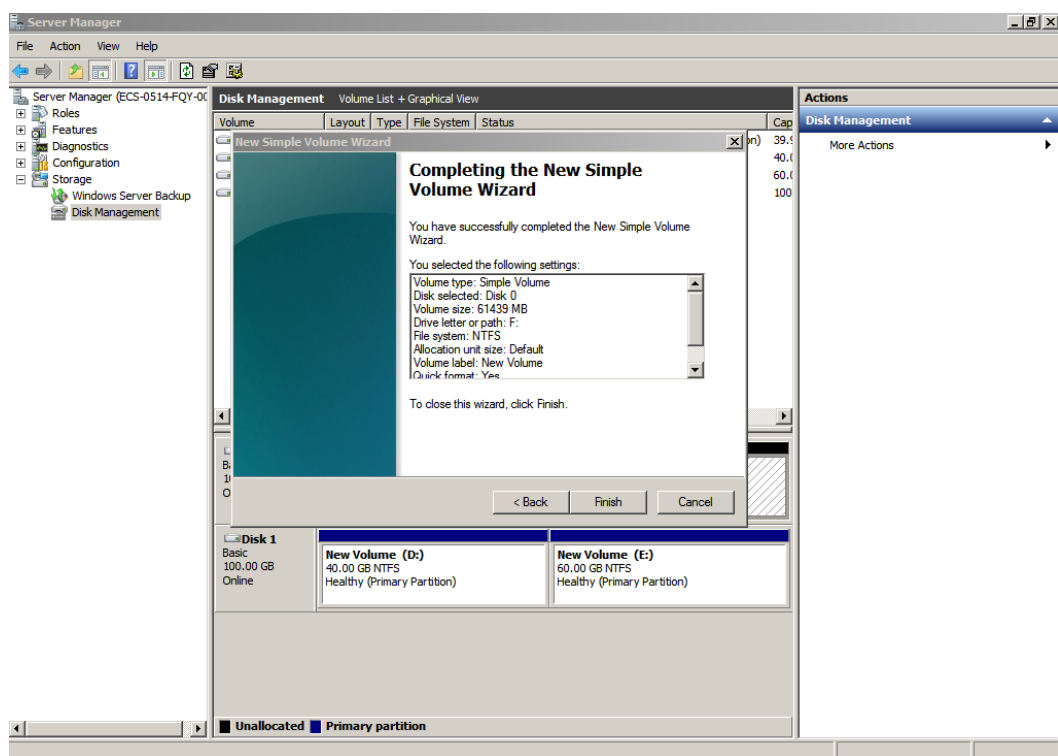
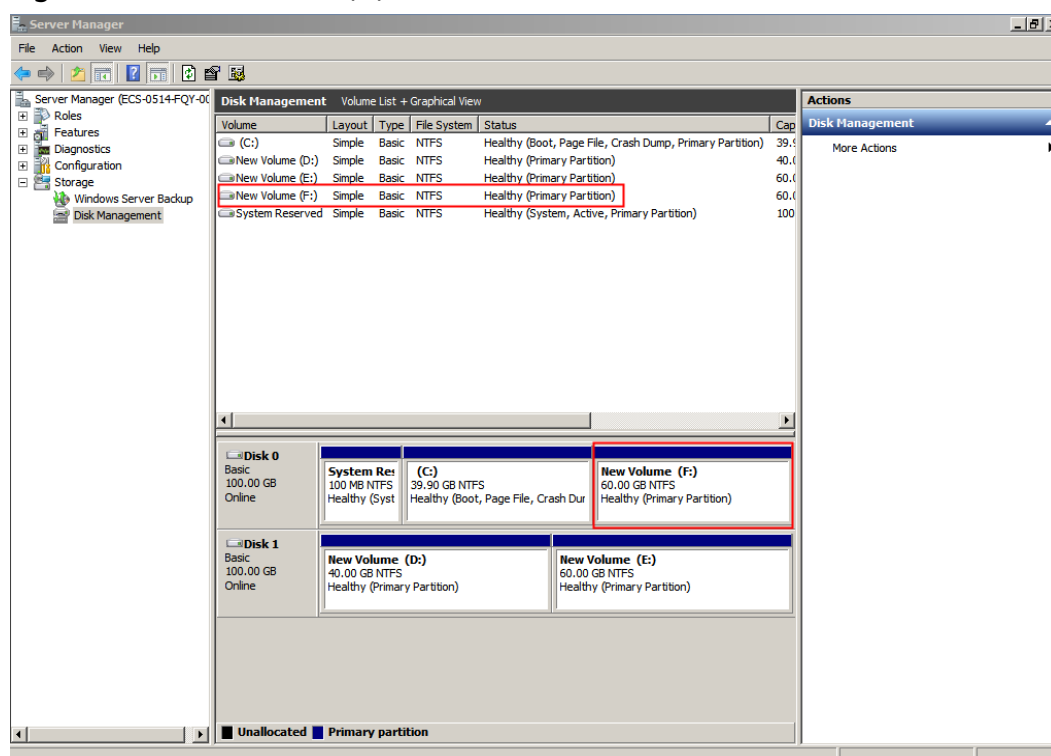


Figure 3-15 New Volume (F:)

-----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 management console and then formatted and added to volume (C:). This 60 GiB has not been used.

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

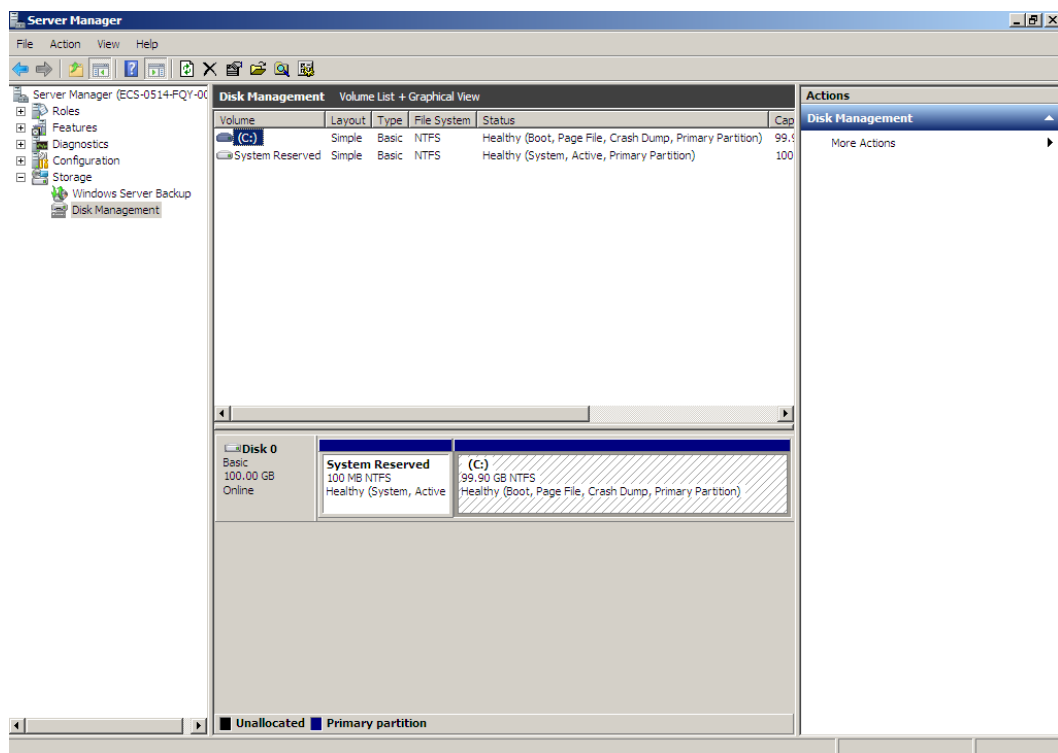
- Step 1** On the desktop of the server, right-click **Computer** and choose **Manage** from the shortcut menu.

The **Server Manager** window is displayed.

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

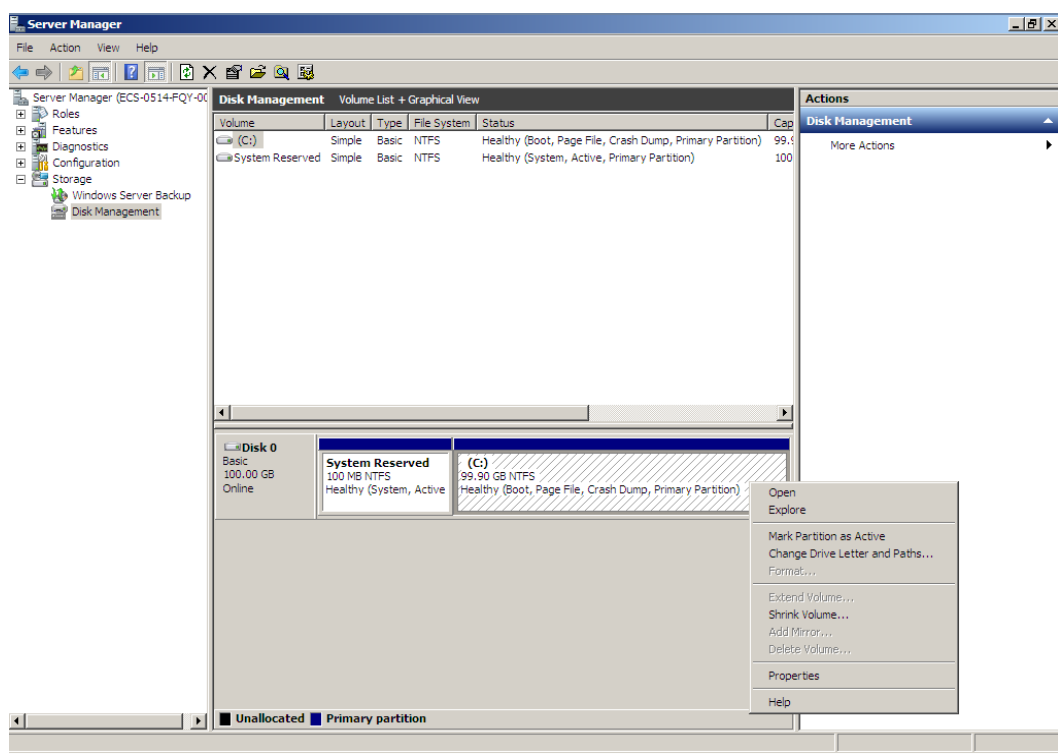
The **Disk Management** window is displayed.

Figure 3-16 Refresh (shrink volume)



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

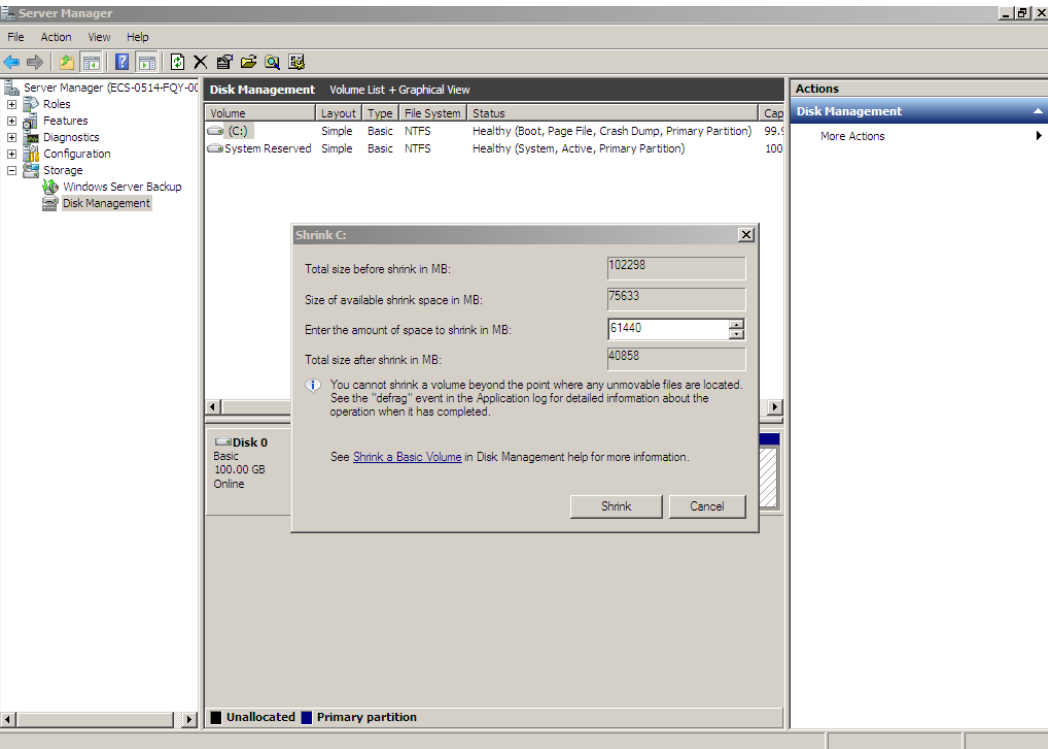
Figure 3-17 Shrink Volume



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

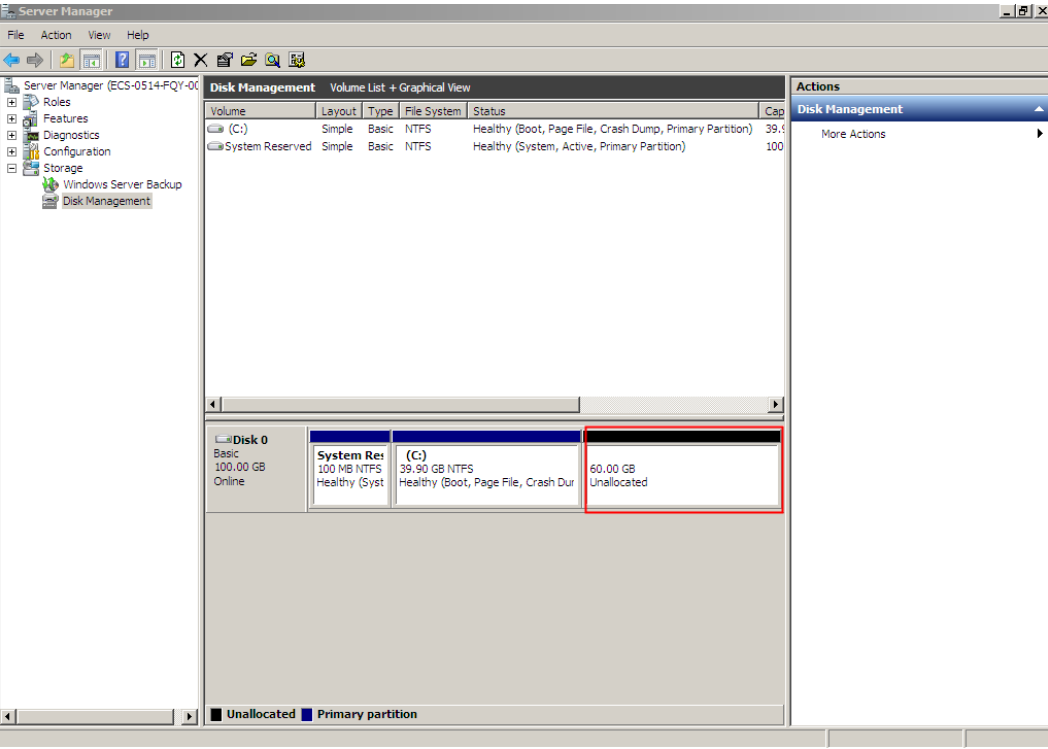
In this example, the volume available space is 60 GiB. Therefore, enter **61440** (60 × 1024 MiB).

Figure 3-18 Shrink (shrink volume)



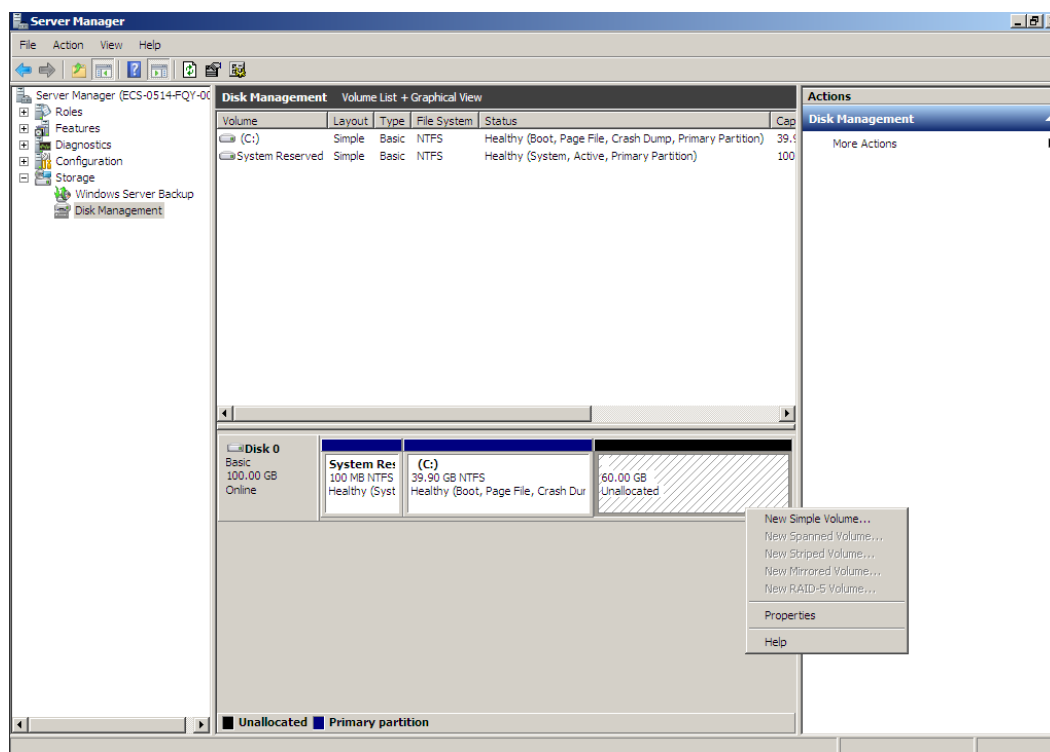
After the operation is complete, **Disk 0** has 60 GiB unallocated space.

Figure 3-19 Unallocated (shrink volume)



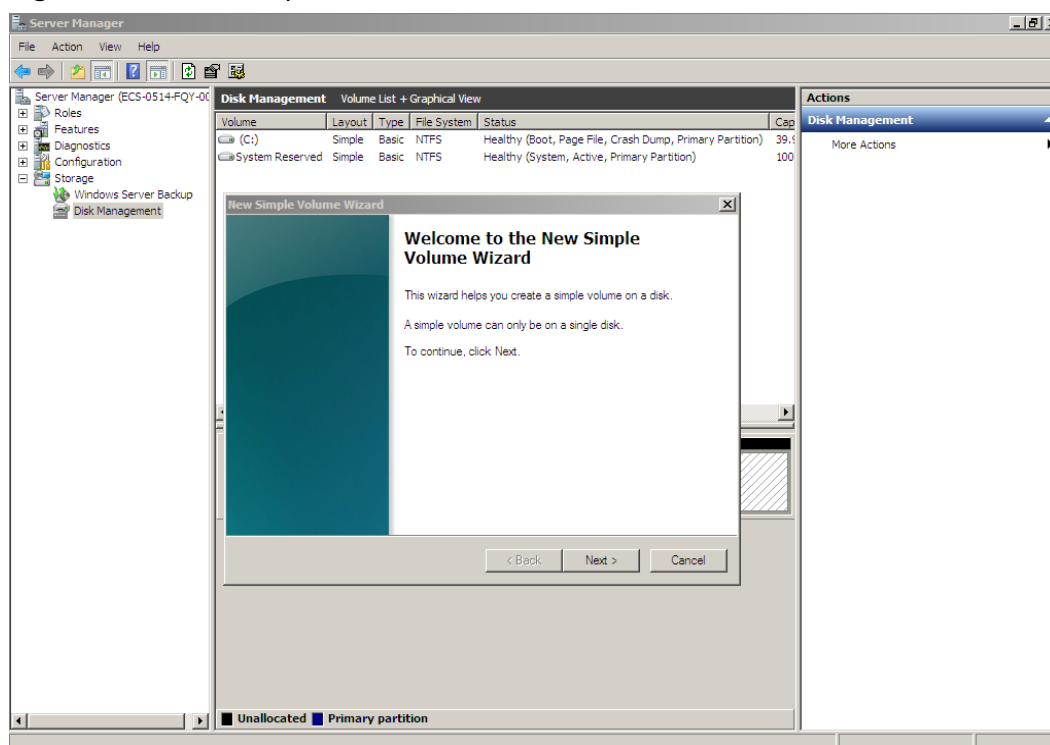
Step 5 In the **Unallocated** area of **Disk 0**, right-click the blank area and choose **New Simple Volume**.

Figure 3-20 New Simple Volume (shrink volume)



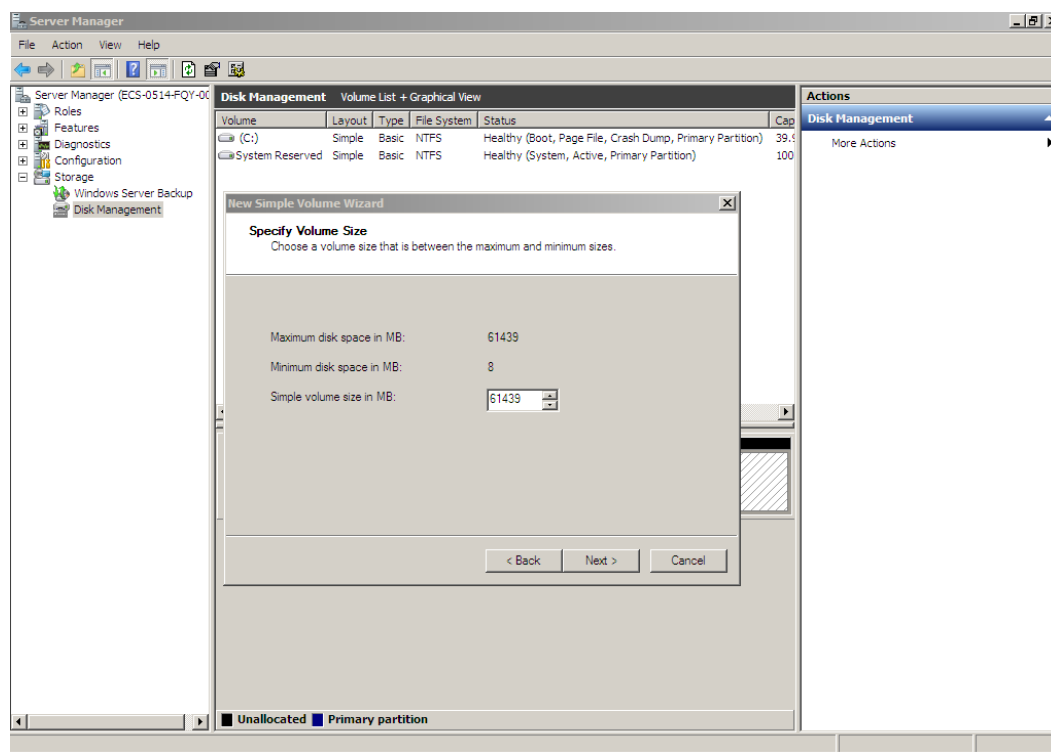
Step 6 On the displayed **New Simple Volume Wizard** window, click **Next**.

Figure 3-21 New Simple Volume Wizard (shrink volume)



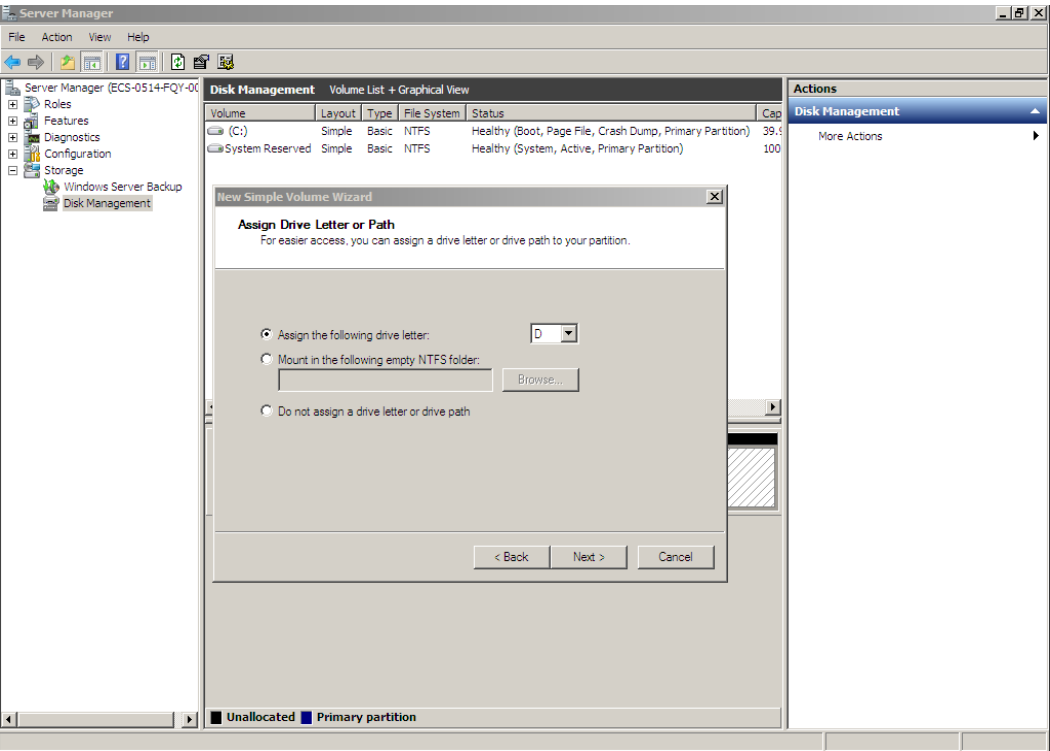
Step 7 On the displayed **Specify Volume Size** page, set **Simple volume size in MB** and click **Next**. In this example, the default size is used.

Figure 3-22 Specify Volume Size (shrink volume)



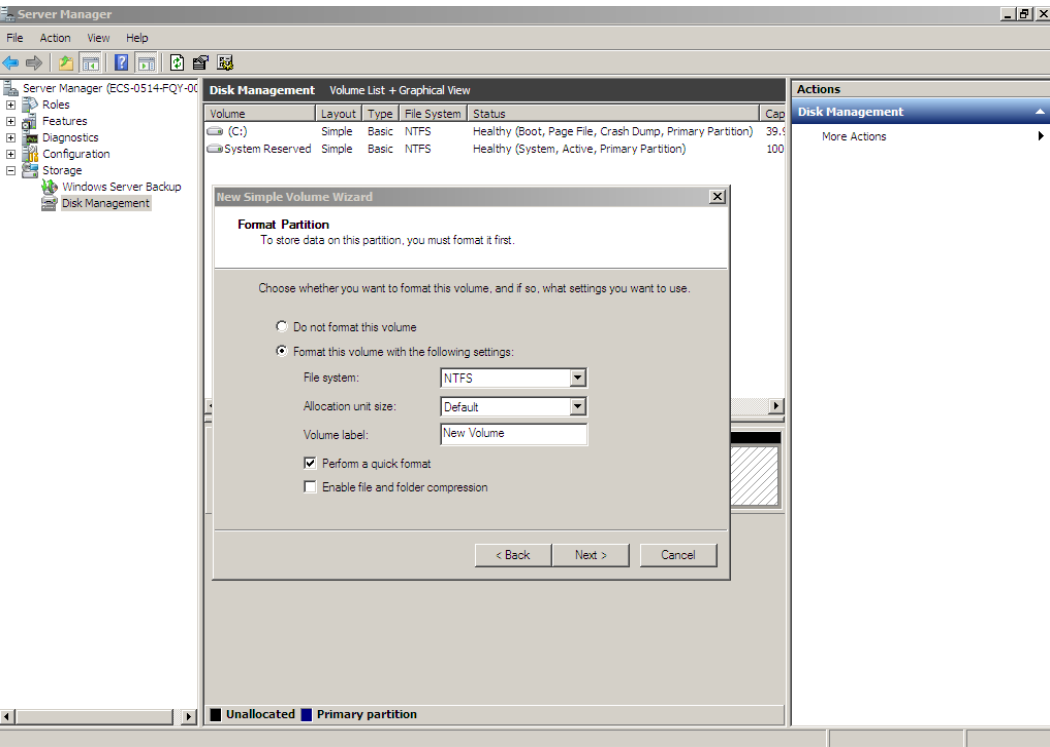
Step 8 On the displayed **Assign Drive Letter and Path** page, click **Assign the following drive letter**, select a drive letter, and click **Next**. In this example, drive letter **D** is selected.

Figure 3-23 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 3-24 Format Partition (shrink volume)



Step 10 Click **Finish**.

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

Figure 3-25 Completing the New Simple Volume Wizard (new volume D:)

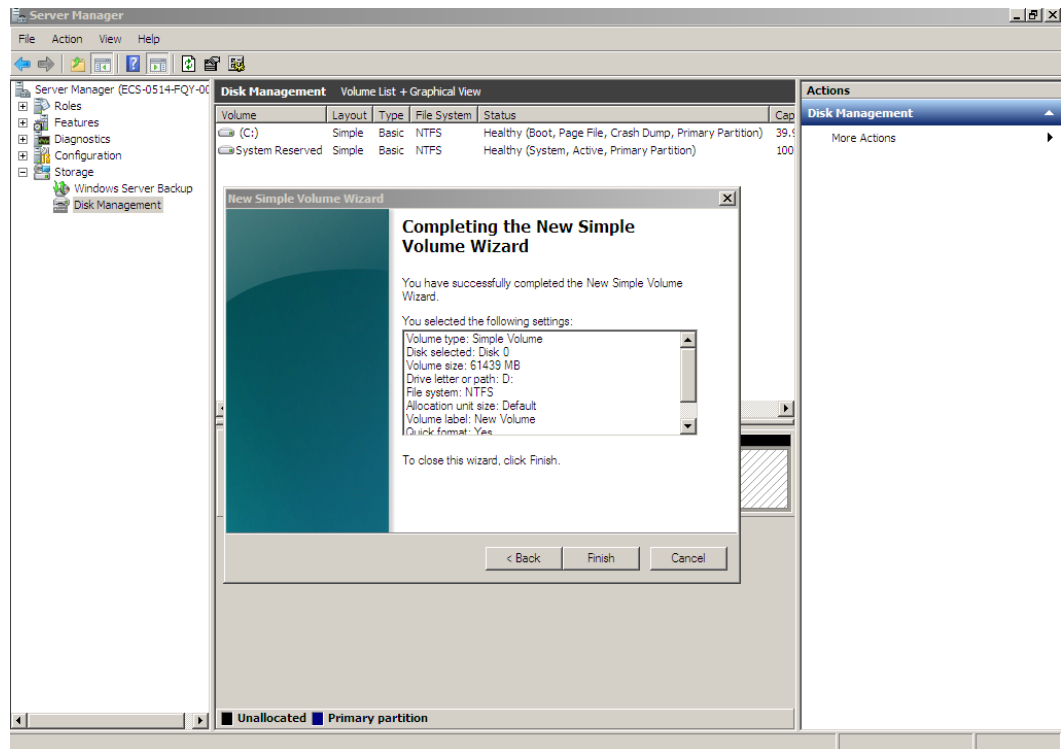
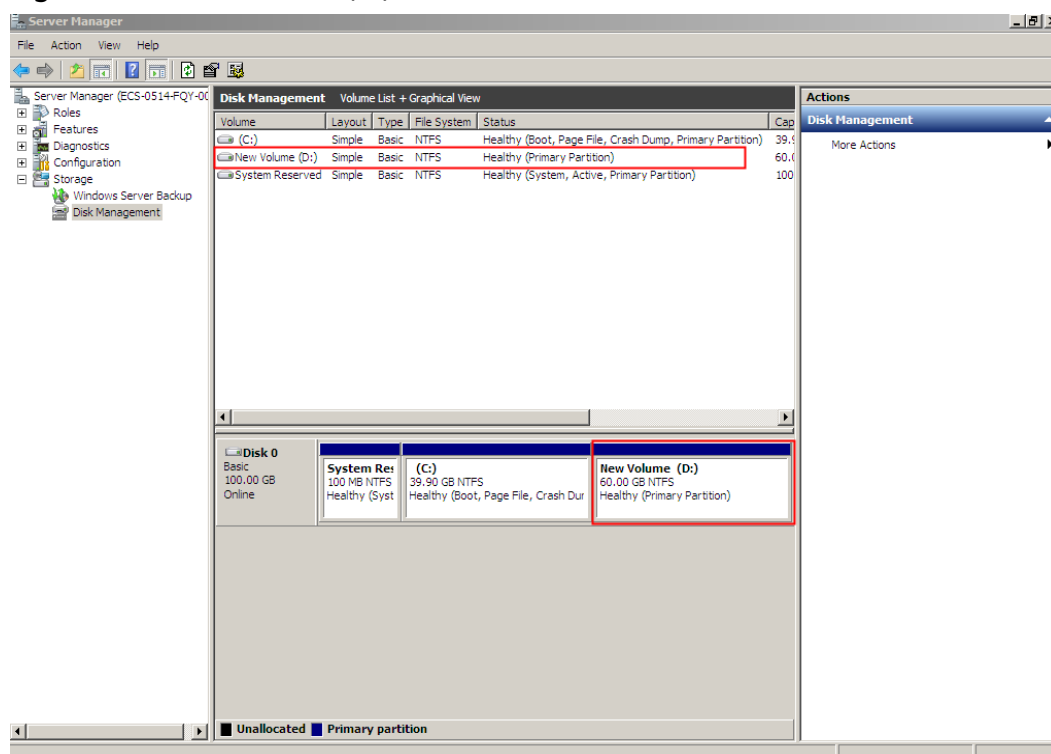


Figure 3-26 New Volume (D:)

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

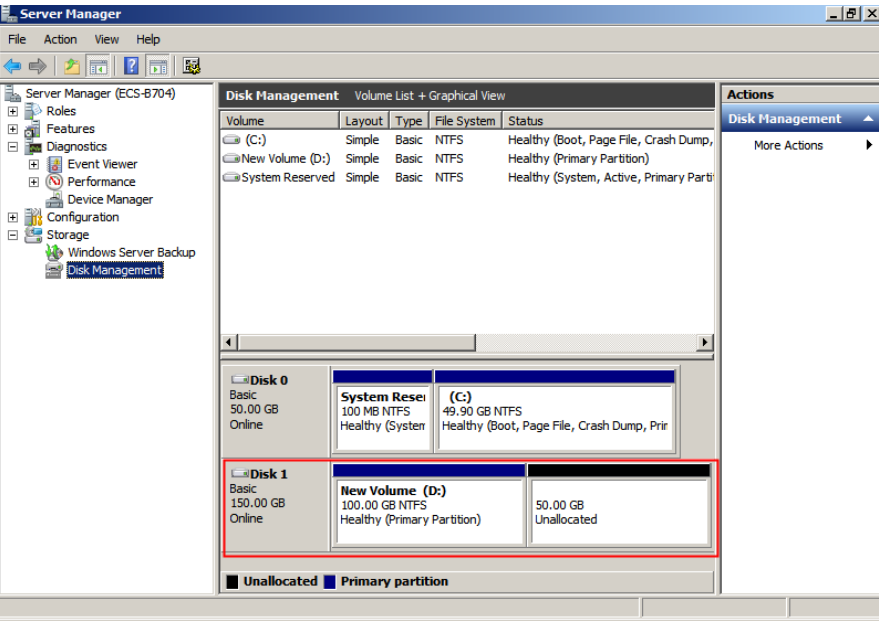
Step 1 On the desktop of the server, right-click **Computer** and choose **Manage** from the shortcut menu.

The **Server Manager** window is displayed.

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

The **Disk Management** window is displayed.

Figure 3-27 Disk Management (data disk)

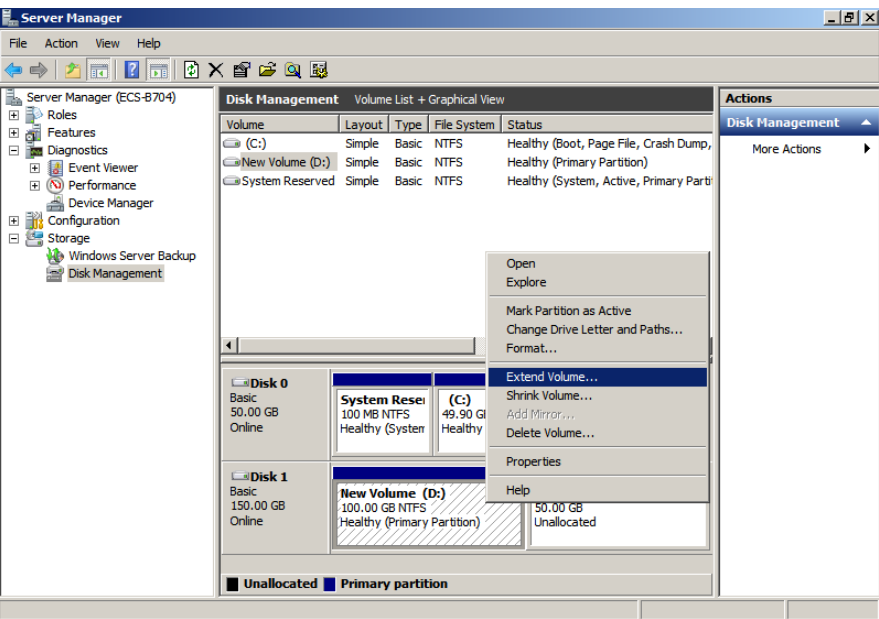


NOTE

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

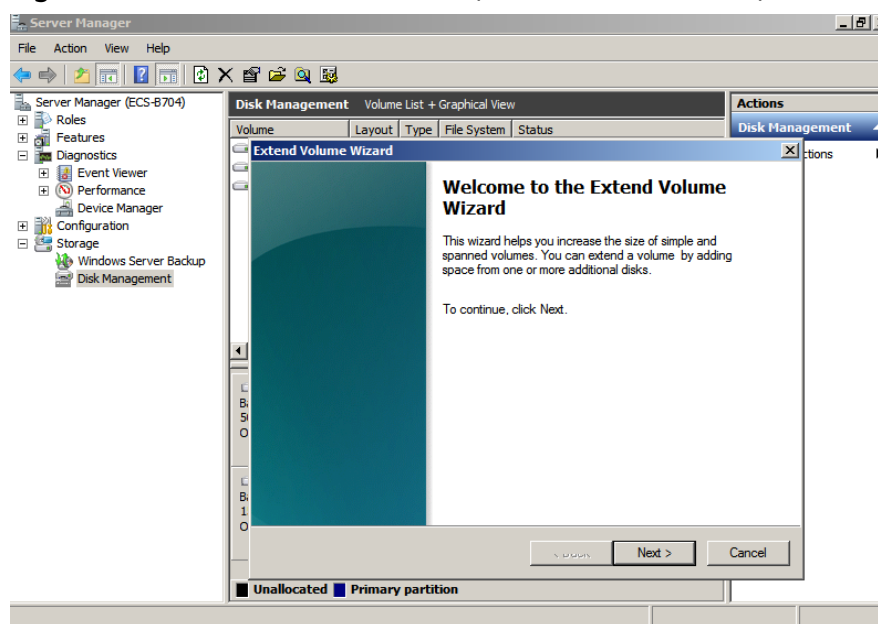
- Step 3
- On the **Disk Management** page, select the disk and volume that you want to extend. The current volume size and unallocated space are displayed.
- Step 4
- Right-click the target volume and choose **Extend Volume**.

Figure 3-28 Choosing Extend Volume (Windows Server 2008)



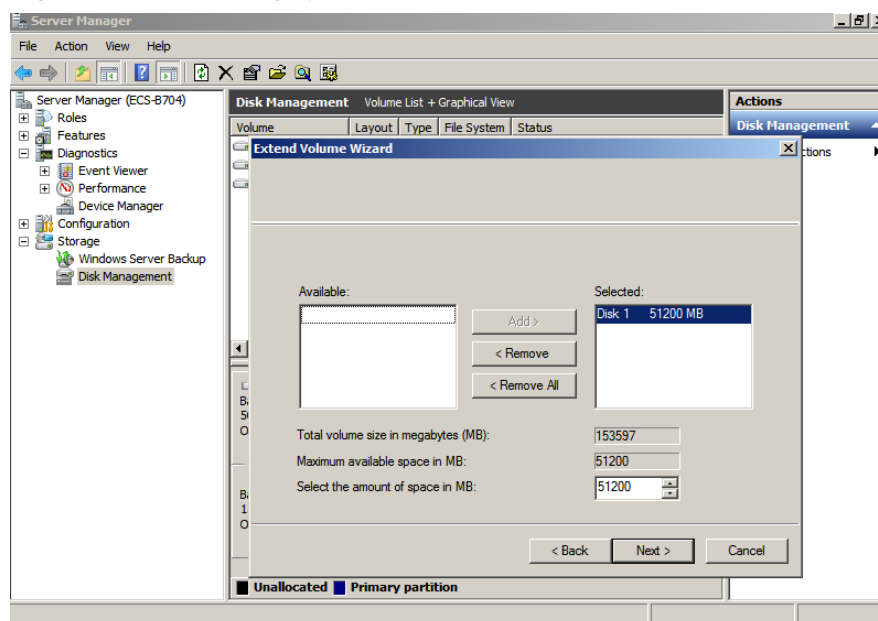
- Step 5
- On the displayed **Extend Volume Wizard** window, click **Next**.

Figure 3-29 Extend Volume Wizard (Windows Server 2008)



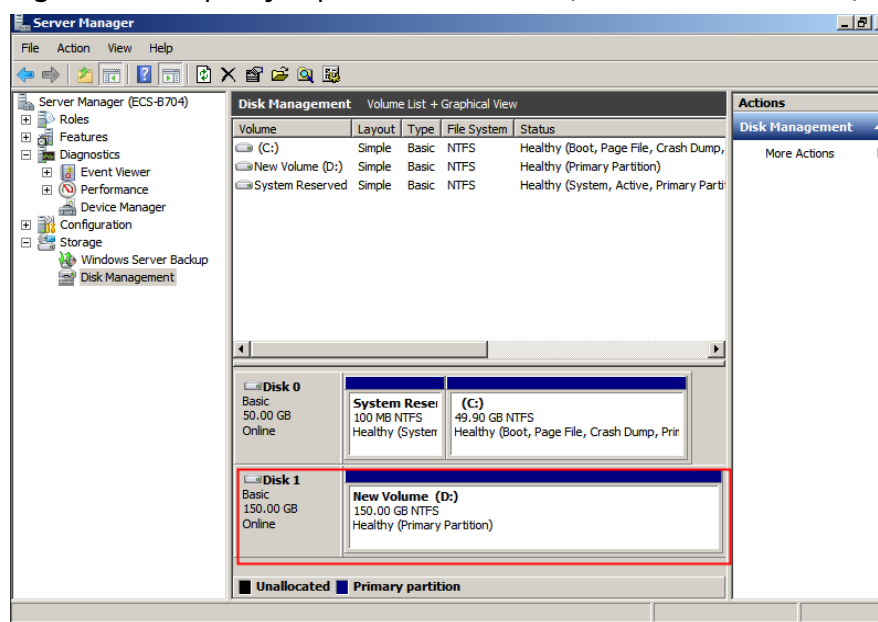
Step 6 In the text box to the right of **Select the amount of space in MB**, enter the amount of the additional space and click **Next**.

Figure 3-30 Selecting space (Windows Server 2008)



Step 7 Click **Finish**.

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

Figure 3-31 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 management console. The following procedure describes how to use this 60 GiB to create a new volume, for example volume (E:), on the server. After the operation is complete, new volume (E:) has 60 GiB of capacity and can be used as a data volume.

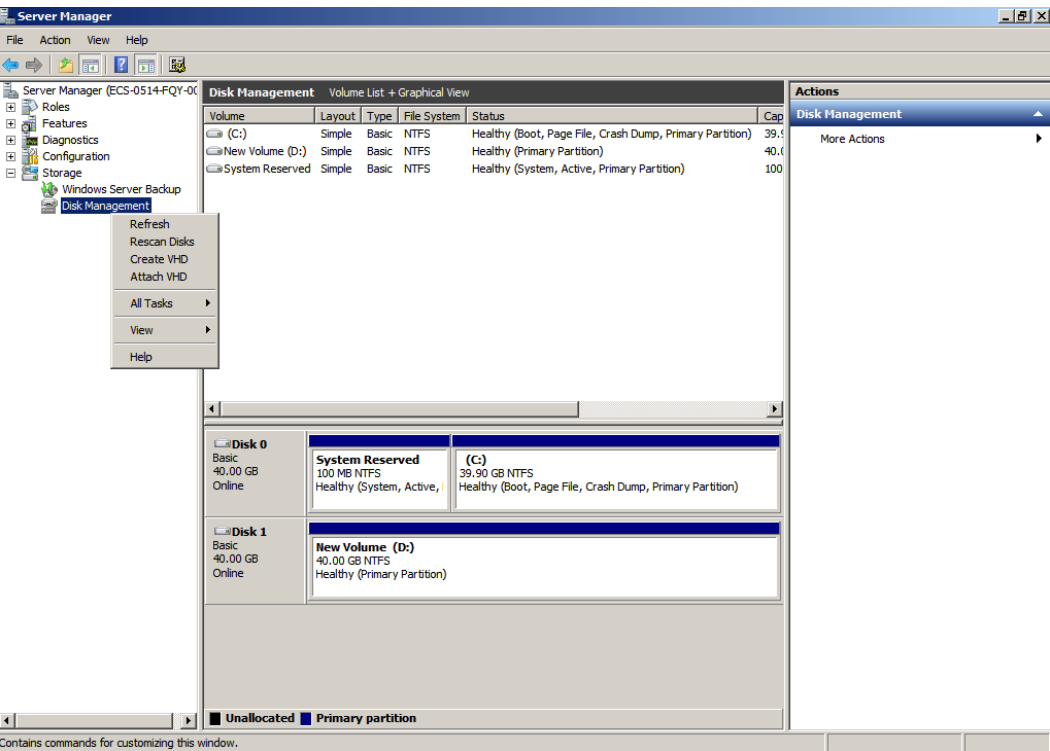
Step 1 On the desktop of the server, right-click **Computer** and choose **Manage** from the shortcut menu.

The **Server Manager** window is displayed.

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

The **Disk Management** window is displayed.

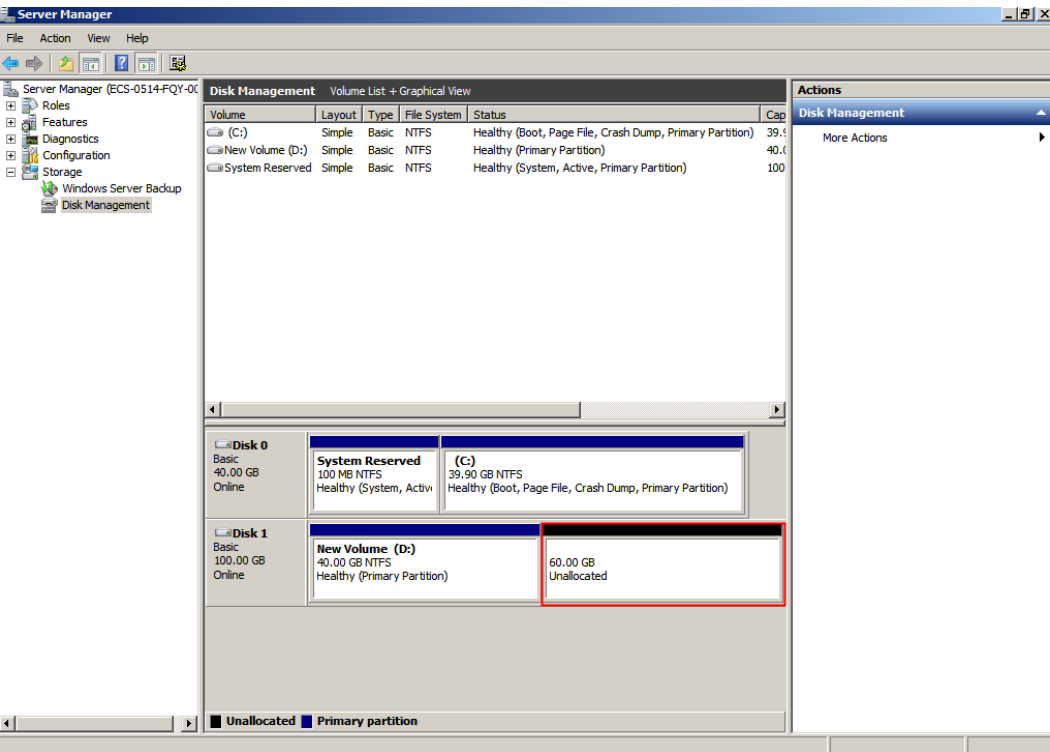
Figure 3-32 Refresh (data disk)



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

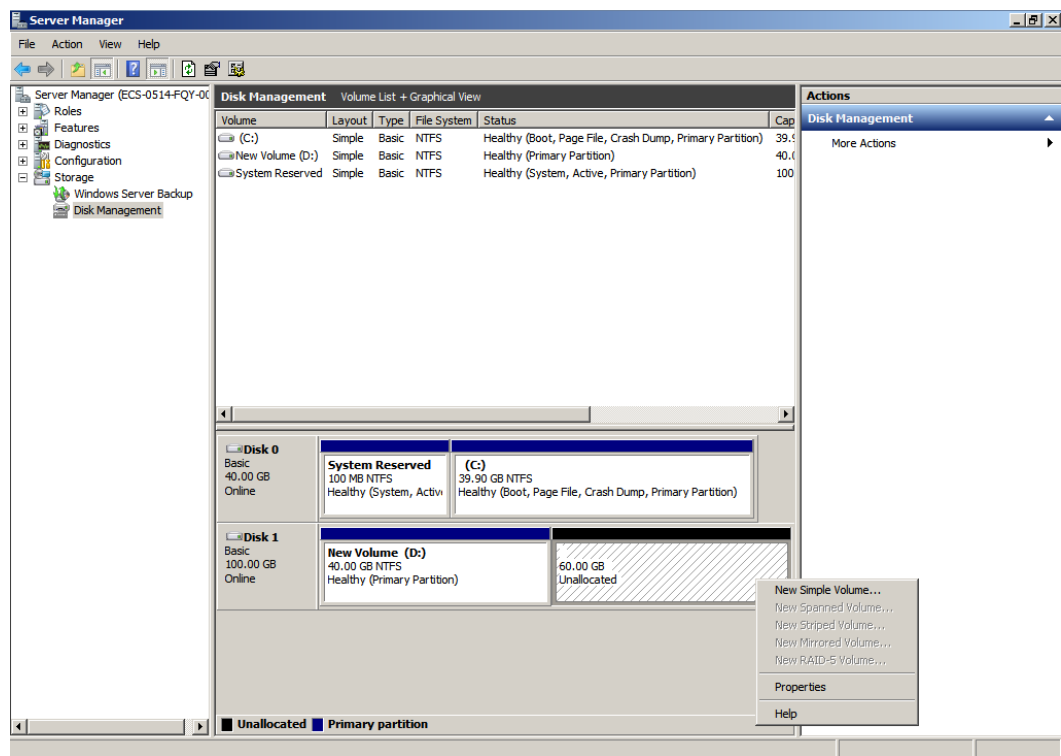
After the refresh, the additional space is displayed in the right area and is unallocated.

Figure 3-33 Unallocated disk space (data disk)



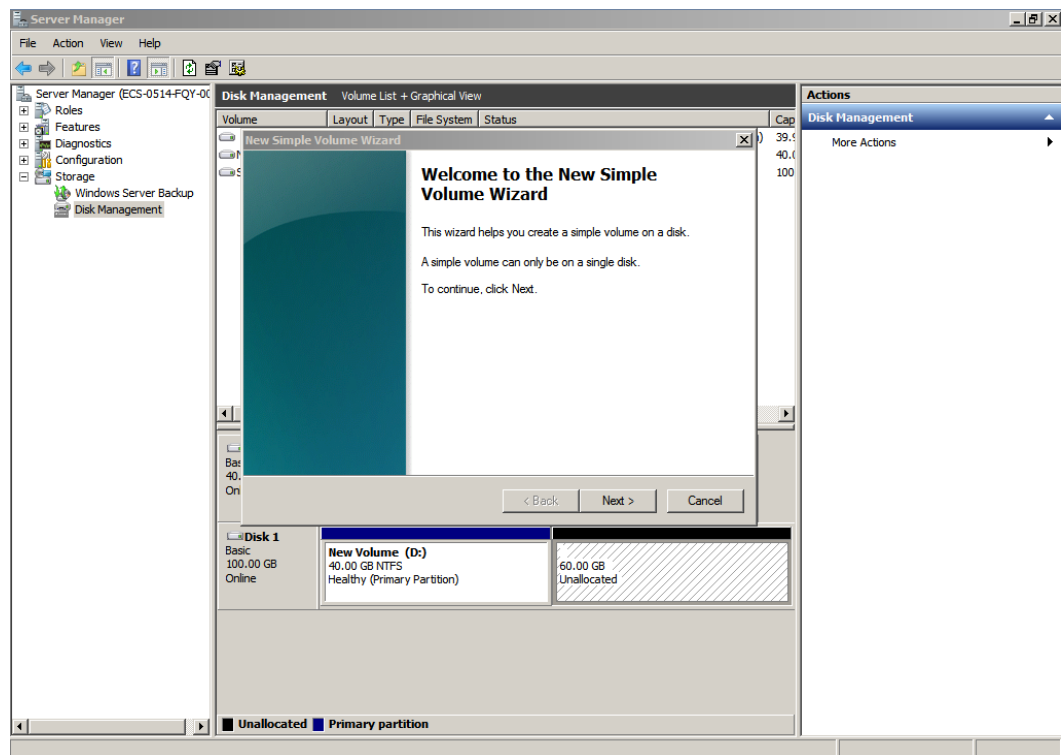
Step 4 In the **Unallocated** area of **Disk 1**, right-click the blank area and choose **New Simple Volume**.

Figure 3-34 New Simple Volume (data disk)



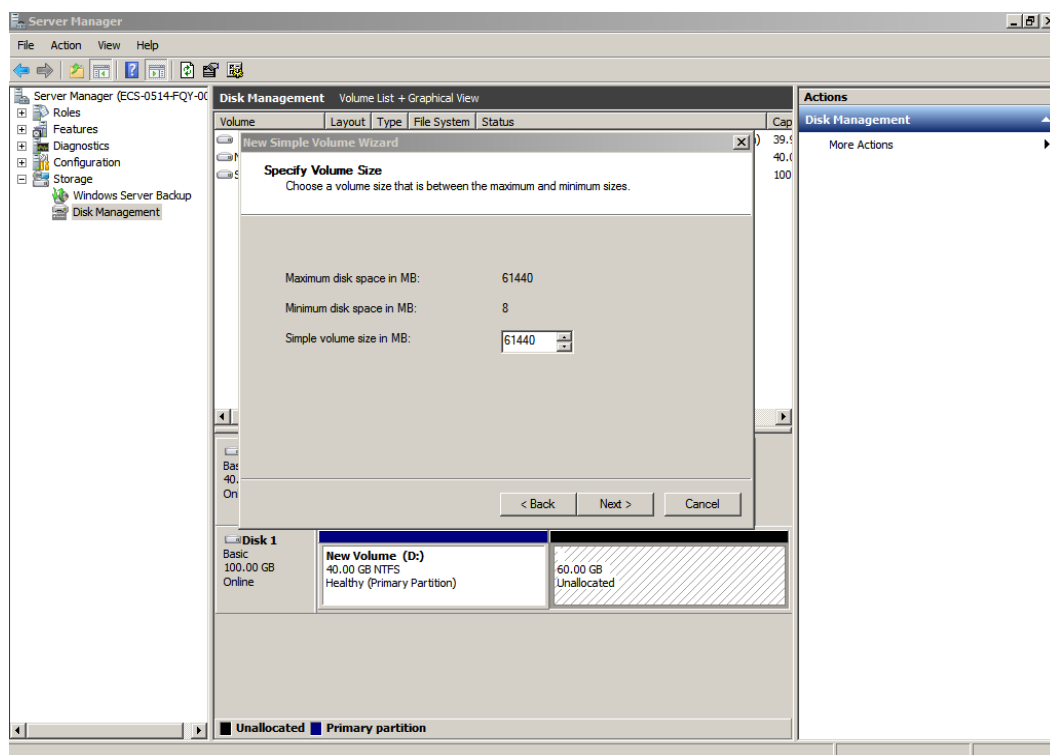
Step 5 On the displayed **New Simple Volume Wizard** window, click **Next**.

Figure 3-35 New Simple Volume Wizard (data disk)



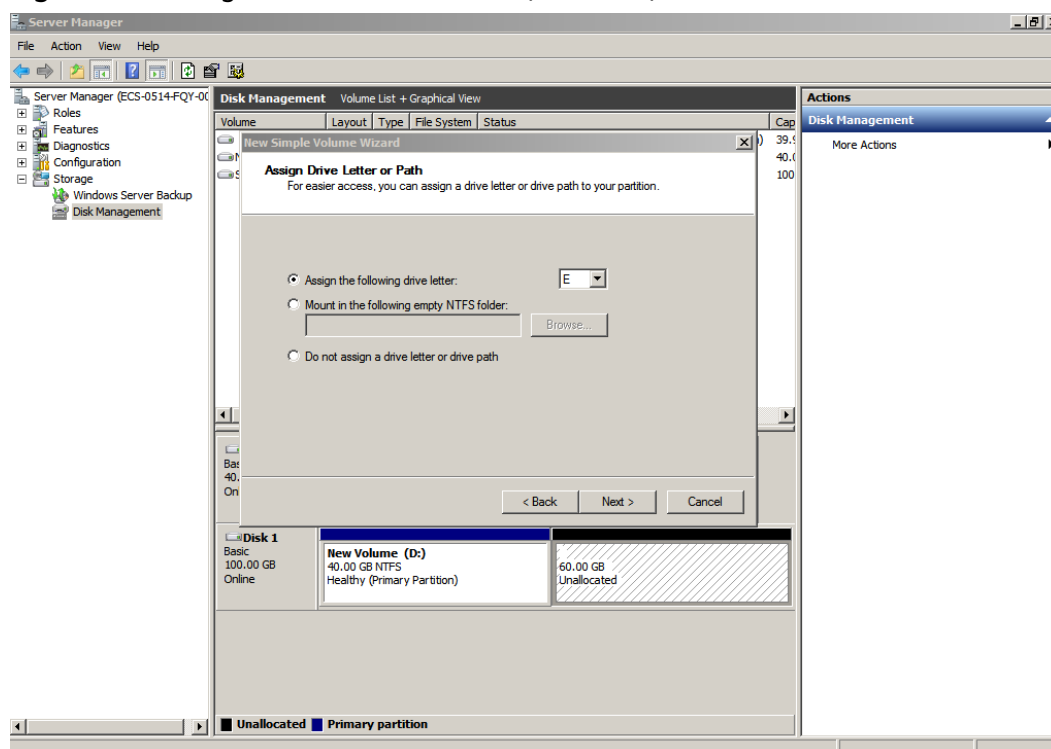
- Step 6** On the displayed **Specify Volume Size** page, set **Simple volume size in MB** and click **Next**. In this example, the default size is used.

Figure 3-36 Specify Volume Size (data disk)



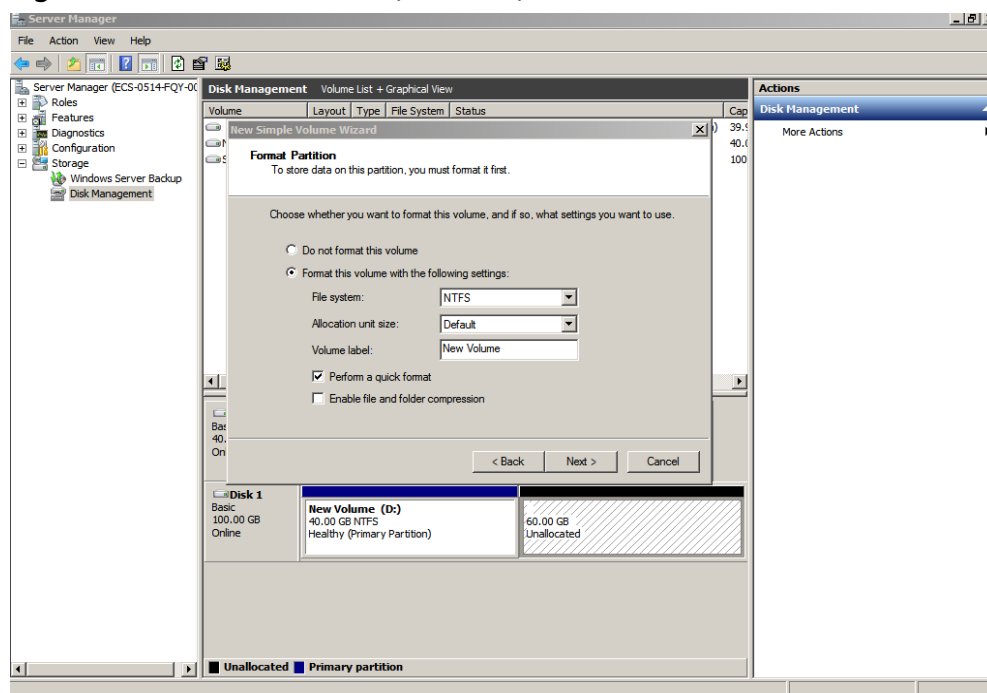
- Step 7** On the displayed **Assign Drive Letter and Path** page, click **Assign the following drive letter**, select a drive letter, and click **Next**. In this example, drive letter **E** is selected.

Figure 3-37 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 3-38 Format Partition (data disk)



Step 9 Click **Finish**.

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

Figure 3-39 Completing the New Simple Volume Wizard (new volume E:)

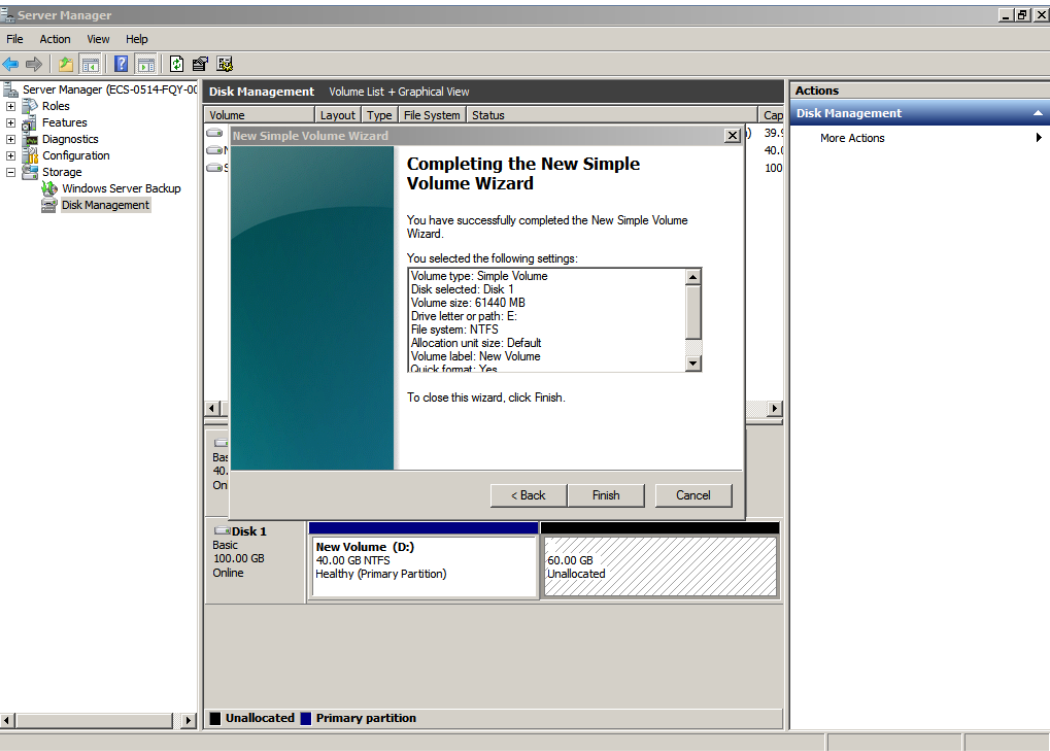
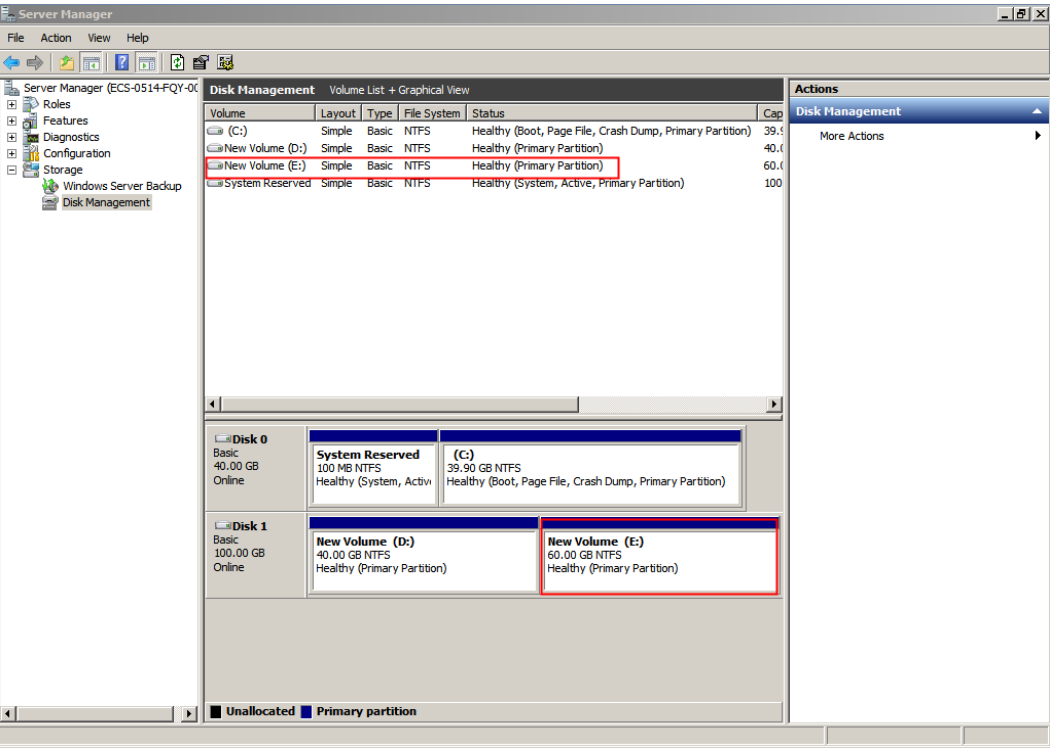


Figure 3-40 New Volume (E:)



----End

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

Scenarios

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

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

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

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

- For a system disk:
 - If volume (C:) already exists, you can add the additional space to volume (C:) and use it as a system volume. For details, see [System Disk: Add 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](#).
- 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

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

Prerequisites

- You have expanded the disk capacity and attached the disk to a server on the management console. For details, see [Expanding Capacity for an In-use EVS Disk](#) or [Expanding Capacity for an Available EVS Disk](#).
- You have logged in to the server.
 - For how to log in to an ECS, see [Logging In to an ECS](#).
 - For how to log in to a BMS, see [Logging In to a BMS](#).


System Disk: Add the Additional Space to the Original Volume

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

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

The **Disk Management** window is displayed.

Figure 3-41 Disk Management (Windows Server 2016)




NOTE

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

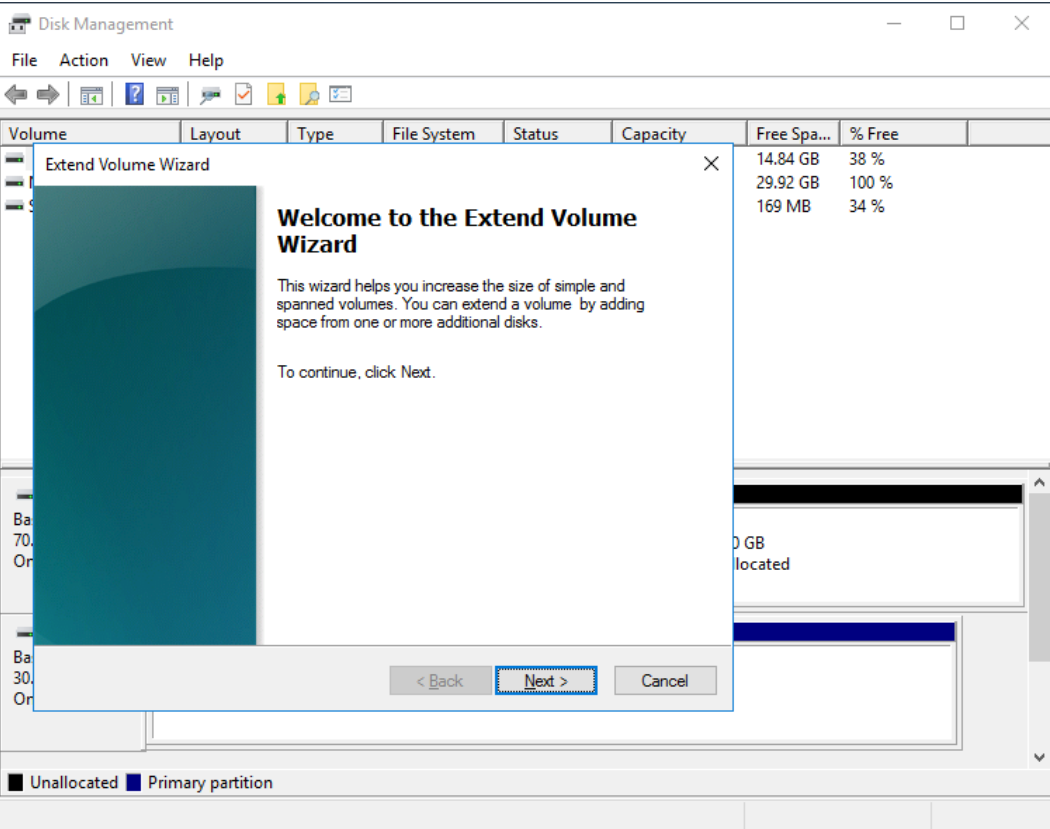
- Step 2** On the **Disk Management** page, select the disk and volume that you want to extend. The current volume size and unallocated space are displayed.
- Step 3** Right-click the target volume and choose **Extend Volume**.

Figure 3-42 Choosing Extend Volume (Windows Server 2016)



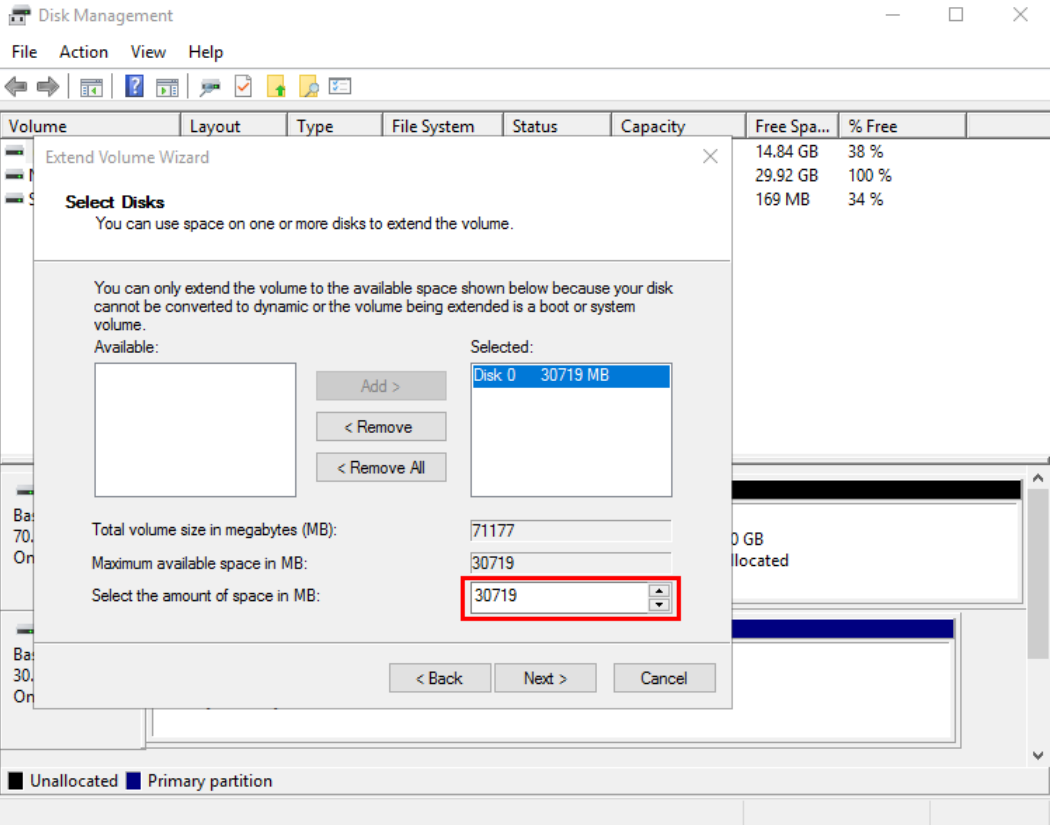
- Step 4** On the displayed **Extend Volume Wizard** window, click **Next**.

Figure 3-43 Extend Volume Wizard (Windows Server 2016)



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

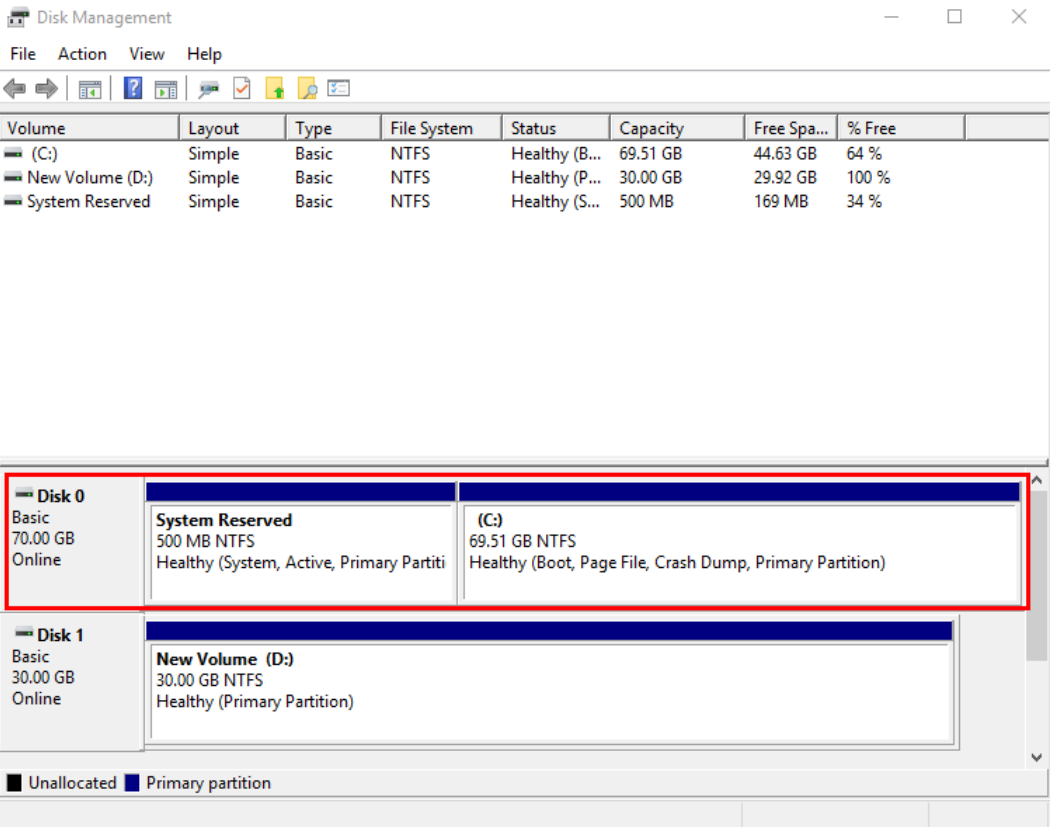
Figure 3-44 Selecting space (Windows Server 2016)



Step 6 Click **Finish**.

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

Figure 3-45 Capacity expansion succeeded (Windows Server 2016)



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

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

The **Disk Management** window is displayed.

Figure 3-46 Unallocated disk space (Windows Server 2016 system disk)

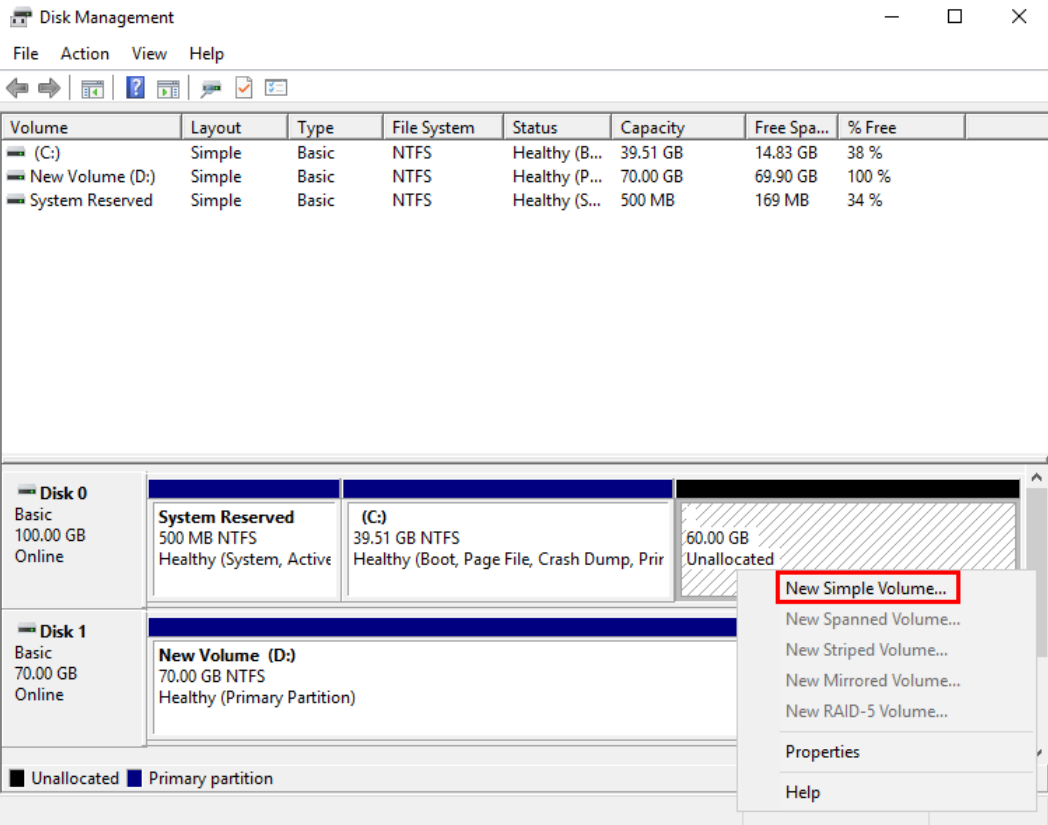


NOTE

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

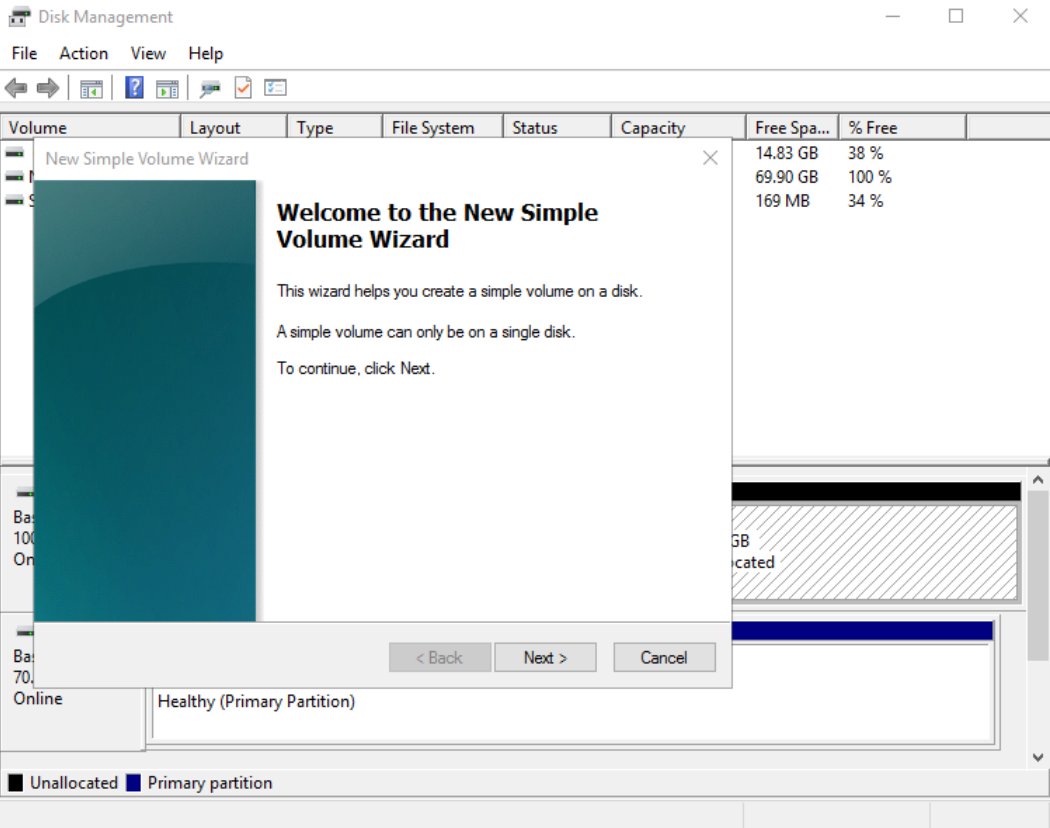
Step 2 In the **Unallocated** area of **Disk 0**, right-click the blank area and choose **New Simple Volume**.

Figure 3-47 New Simple Volume (Windows Server 2016 system disk)



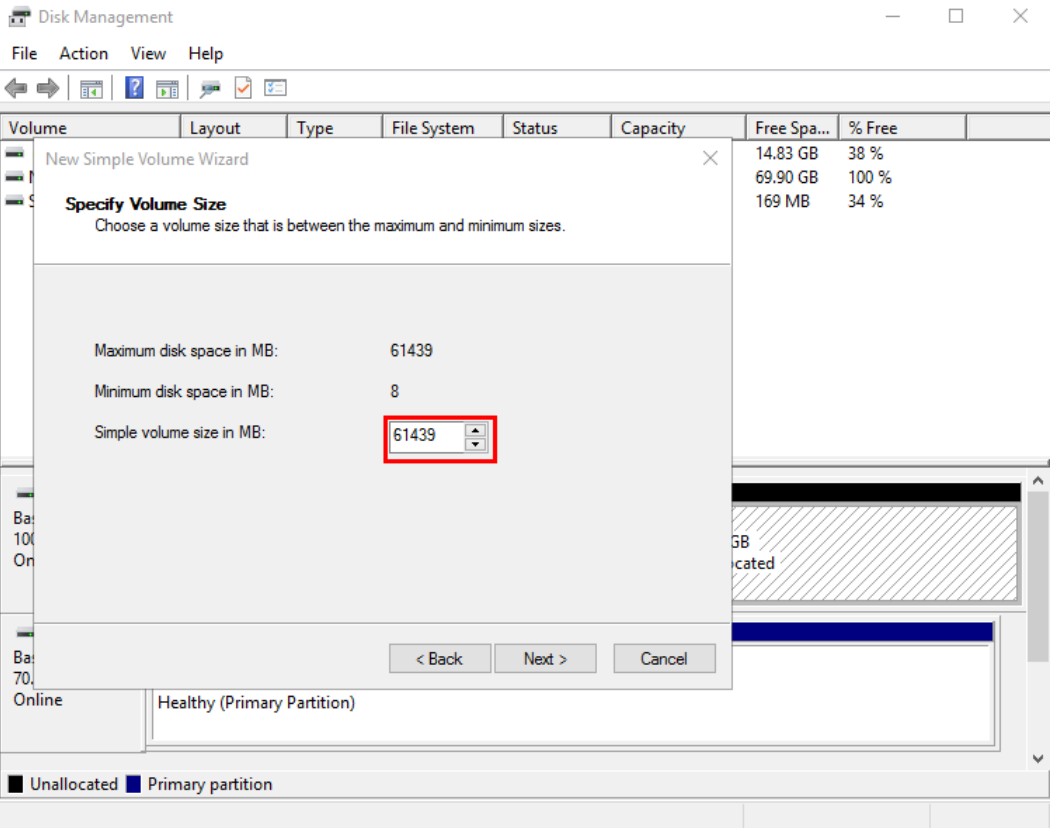
Step 3 On the displayed **New Simple Volume Wizard** window, click **Next**.

Figure 3-48 New Simple Volume Wizard (Windows Server 2016 system disk)




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

Figure 3-49 Specify Volume Size (Windows Server 2016 system disk)



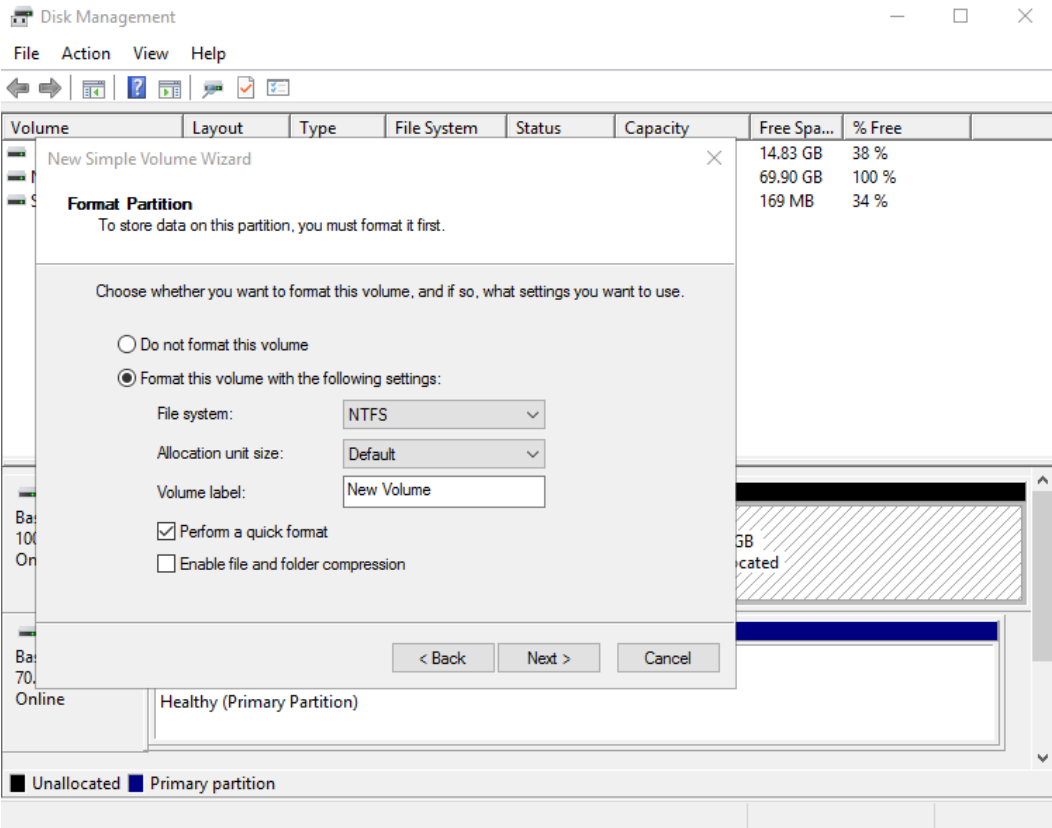
Step 5 On the displayed **Assign Drive Letter and Path** page, click **Assign the following drive letter**, select a drive letter, and click **Next**. In this example, drive letter **F** is selected.

Figure 3-50 Assign Drive Letter or Path (Windows Server 2016 system disk)



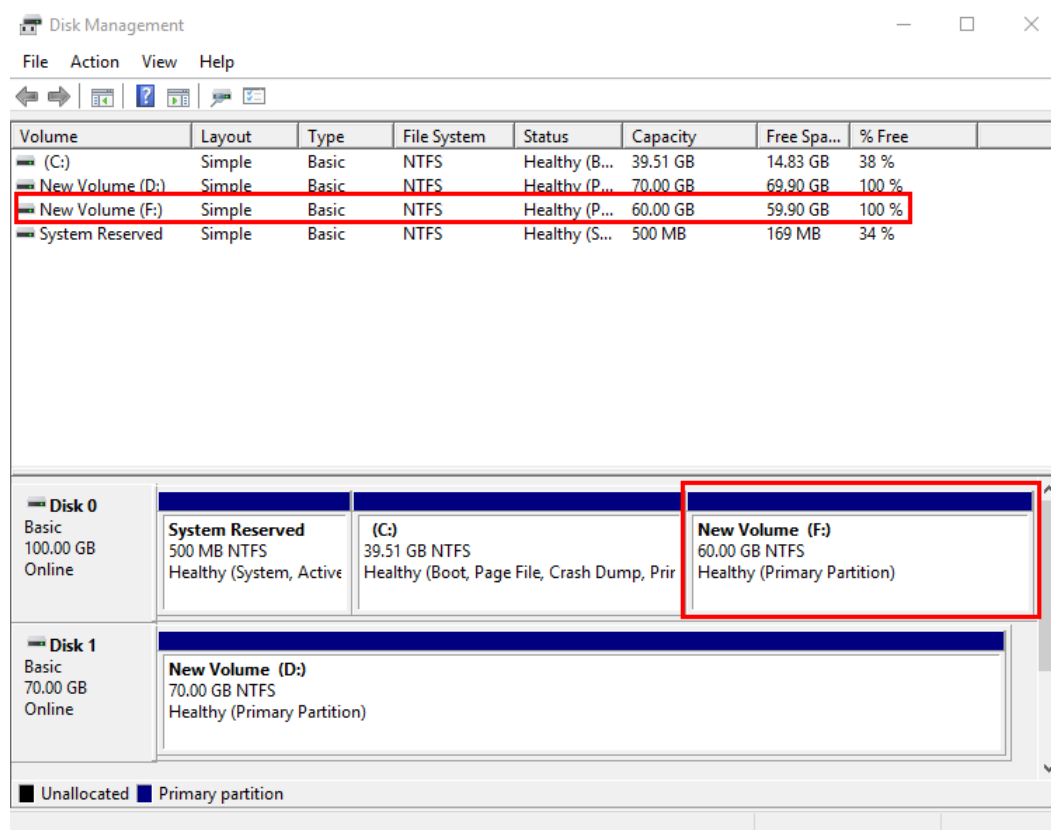
Step 6 On the displayed **Format Partition** page, click **Format this volume with the following settings**, set parameters based on the requirements, and select **Perform a quick format**. Then, click **Next**.

Figure 3-51 Format Partition (Windows Server 2016 system disk)



Step 7 Click **Finish**.

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

Figure 3-52 Volume (F:) (Windows Server 2016)

----End

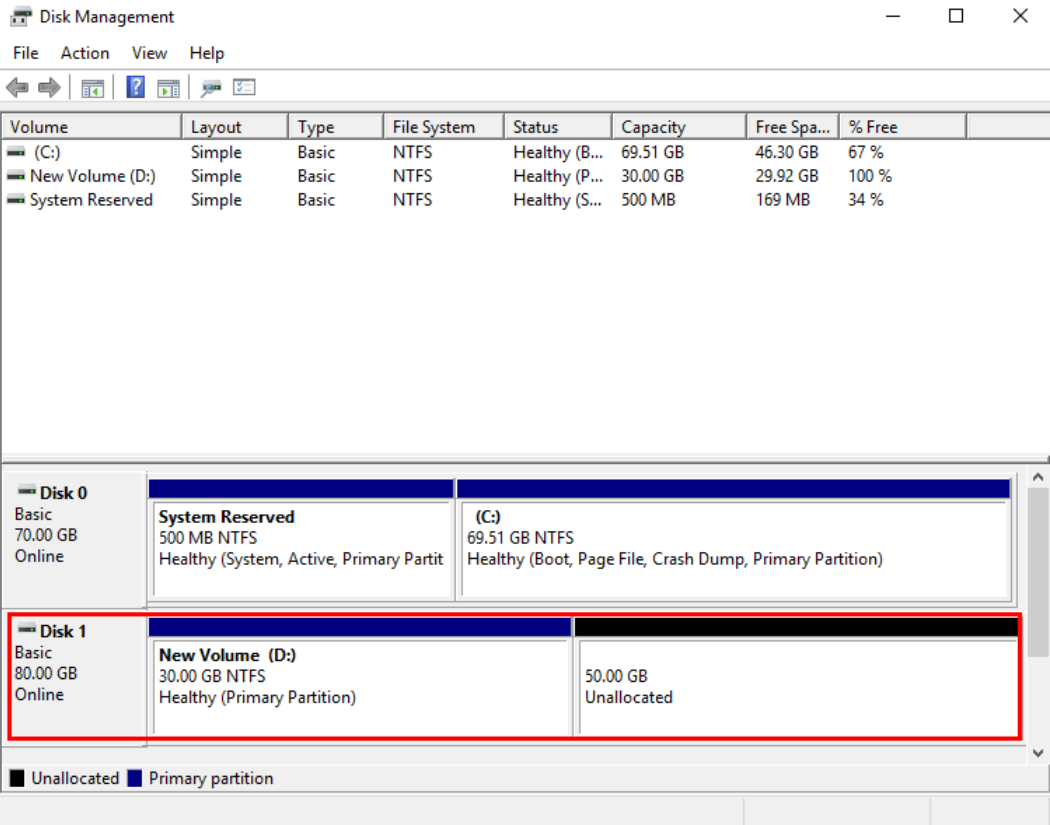
Data Disk: Add the Additional Space to the Original Volume

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

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

The **Disk Management** window is displayed.

Figure 3-53 Disk Management (Windows Server 2016 data disk)

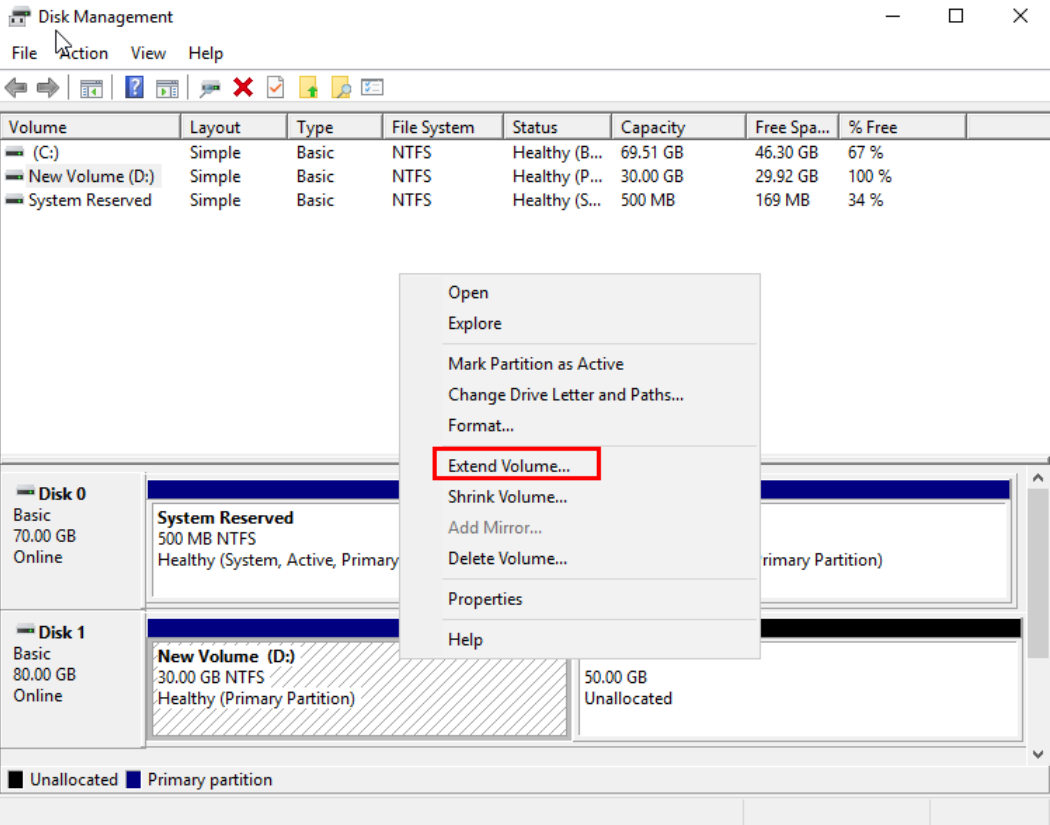


NOTE

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

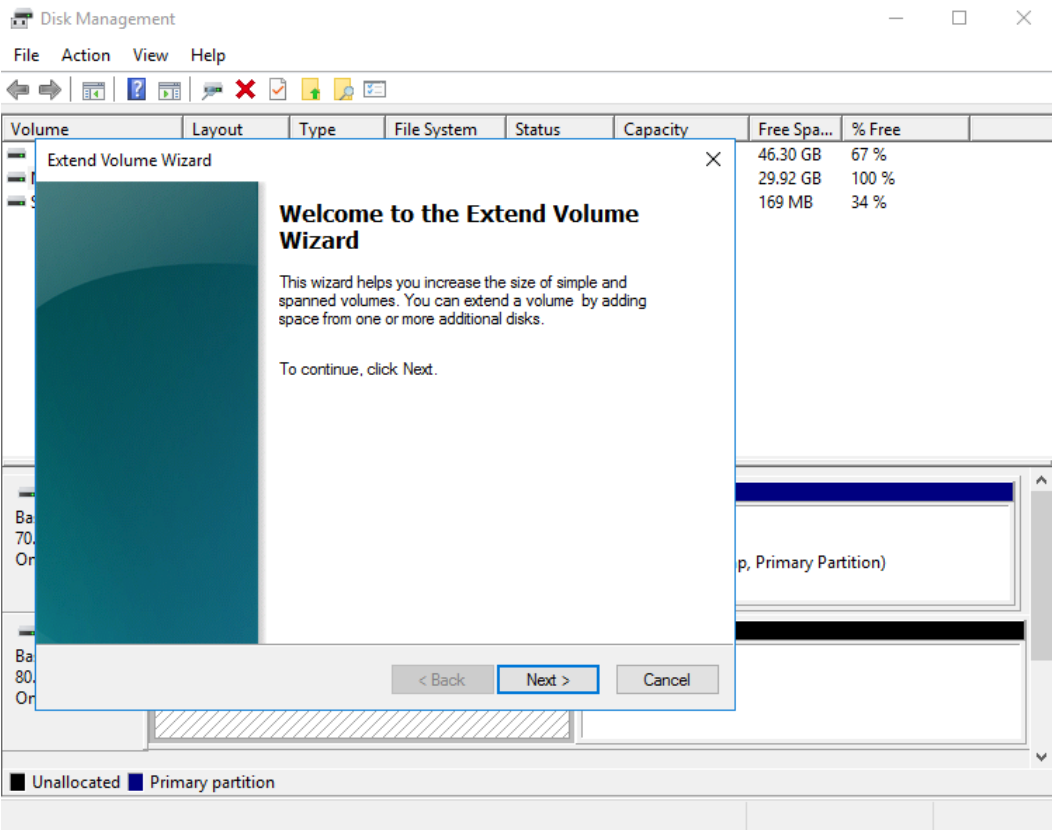
- Step 2
- On the **Disk Management** page, select the disk and volume that you want to extend. The current volume size and unallocated space are displayed.
- Step 3
- Right-click the target volume and choose **Extend Volume**.

Figure 3-54 Choosing Extend Volume (Windows Server 2016 operating system)



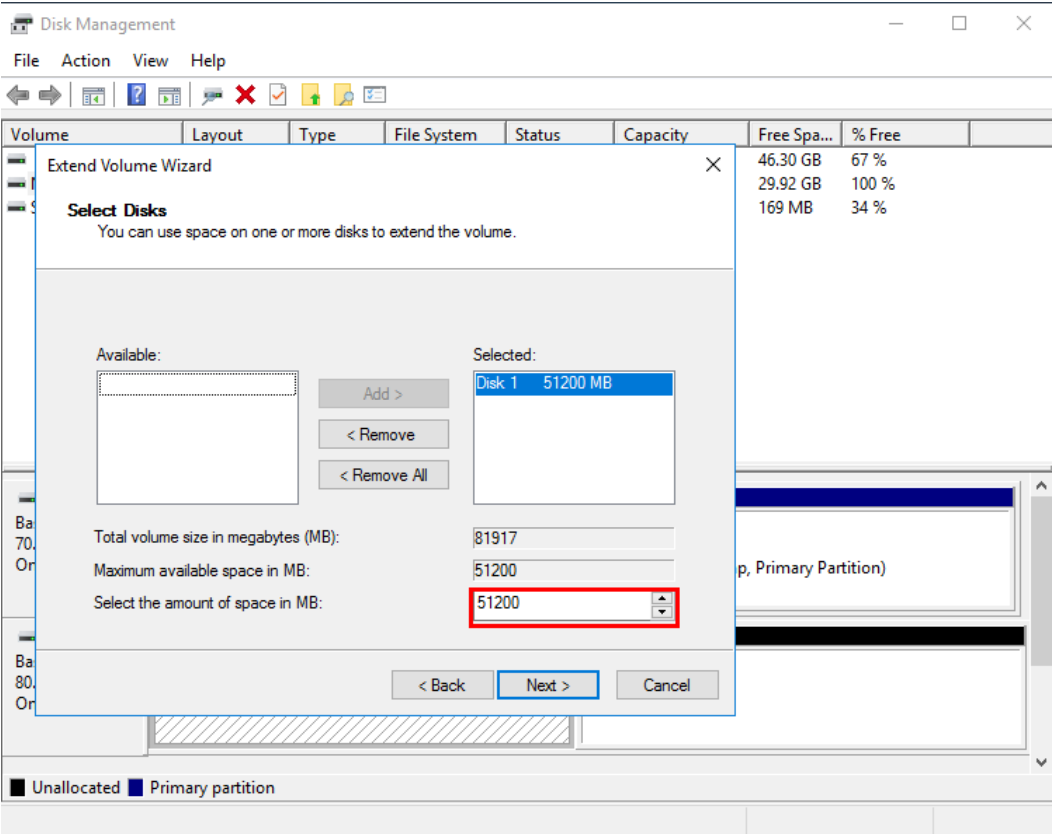
Step 4 On the displayed **Extend Volume Wizard** window, click **Next**.

Figure 3-55 Extend Volume Wizard (Windows Server 2016 operating system)

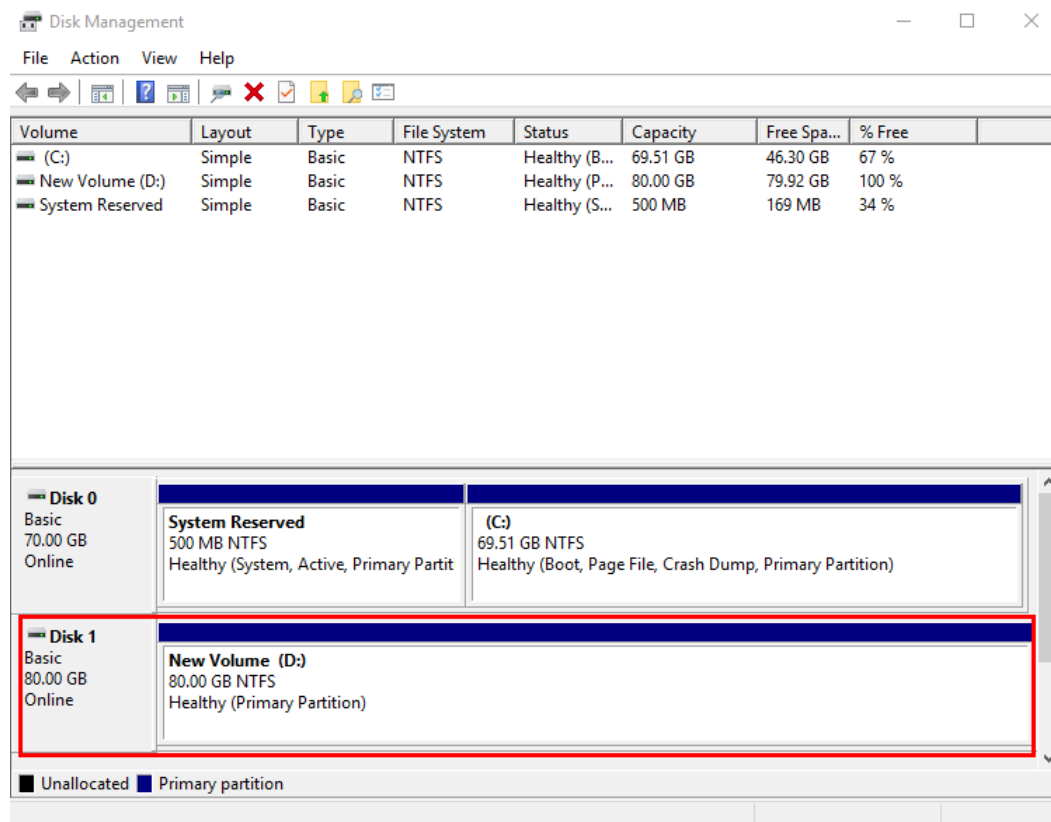


Step 5 In the text box to the right of **Select the amount of space in MB**, enter the amount of the additional space and click **Next**.

Figure 3-56 Selecting space (Windows Server 2016 operating system)



- Step 6** Click **Finish**.
- After the expansion succeeded, the partition size is larger than the original size.

Figure 3-57 Capacity expansion succeeded (Windows Server 2016 operating system)

-----End

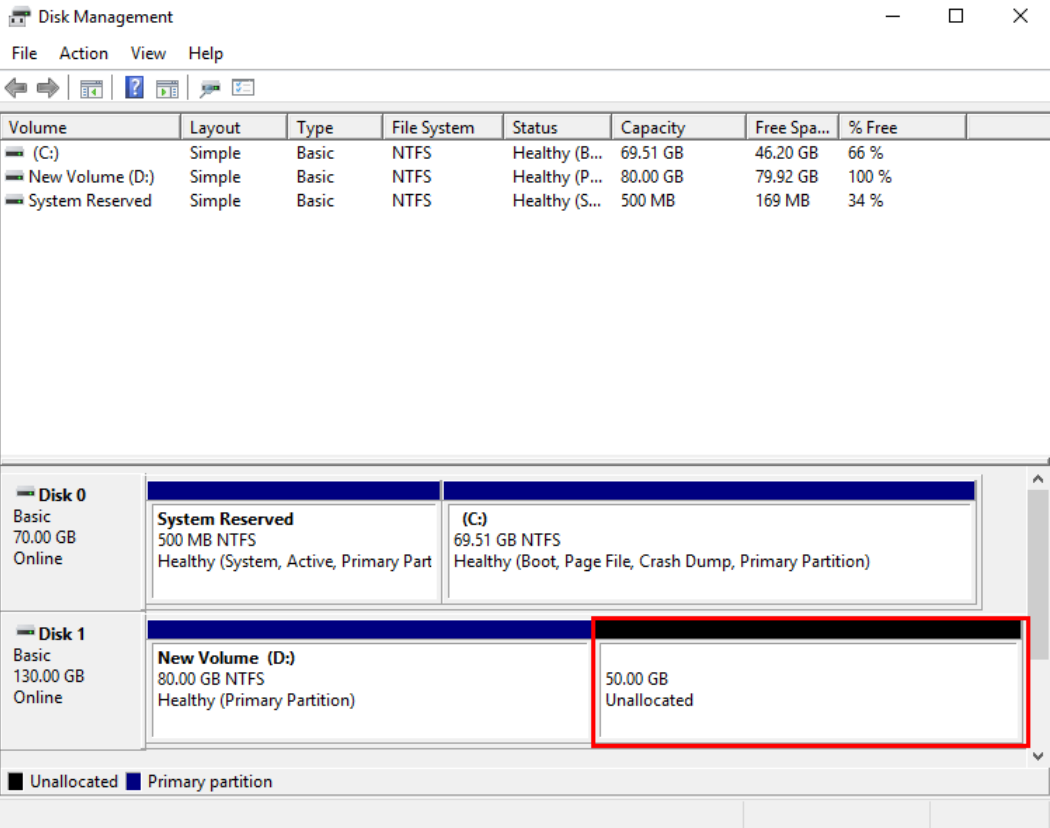
Data Disk: Create a New Volume with the Additional Space

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

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

The **Disk Management** window is displayed.

Figure 3-58 Unallocated disk space (Windows Server 2016 data disk)

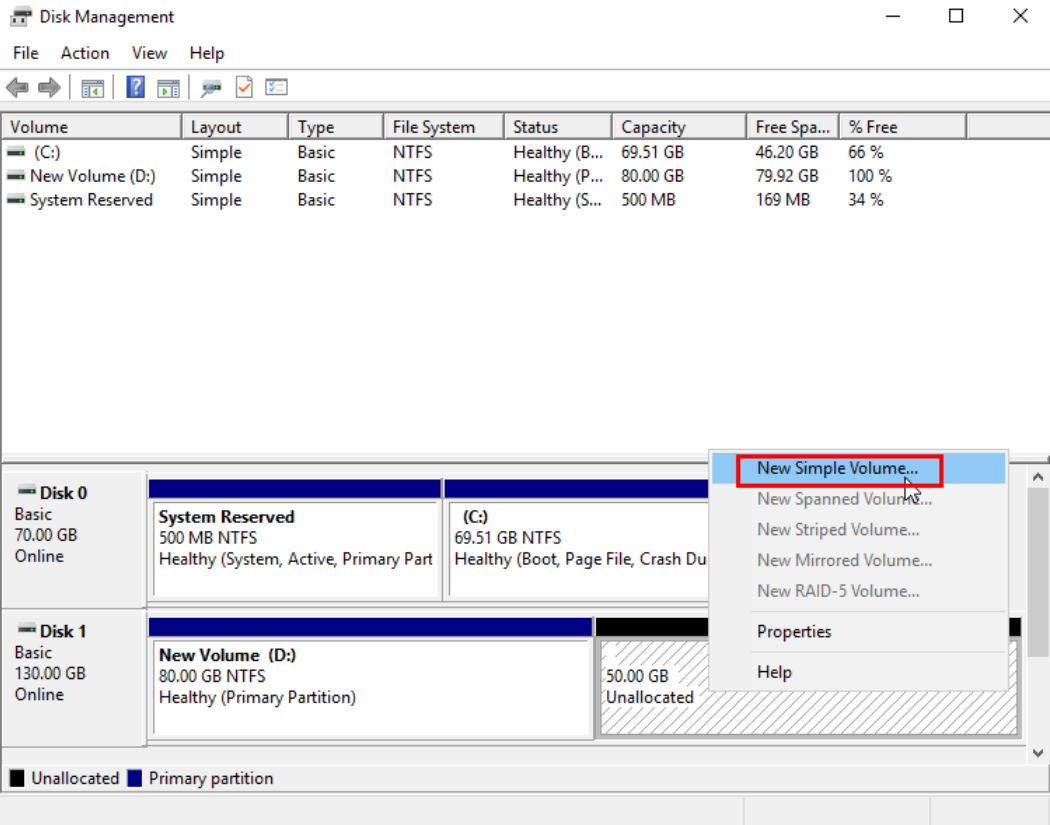


NOTE

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

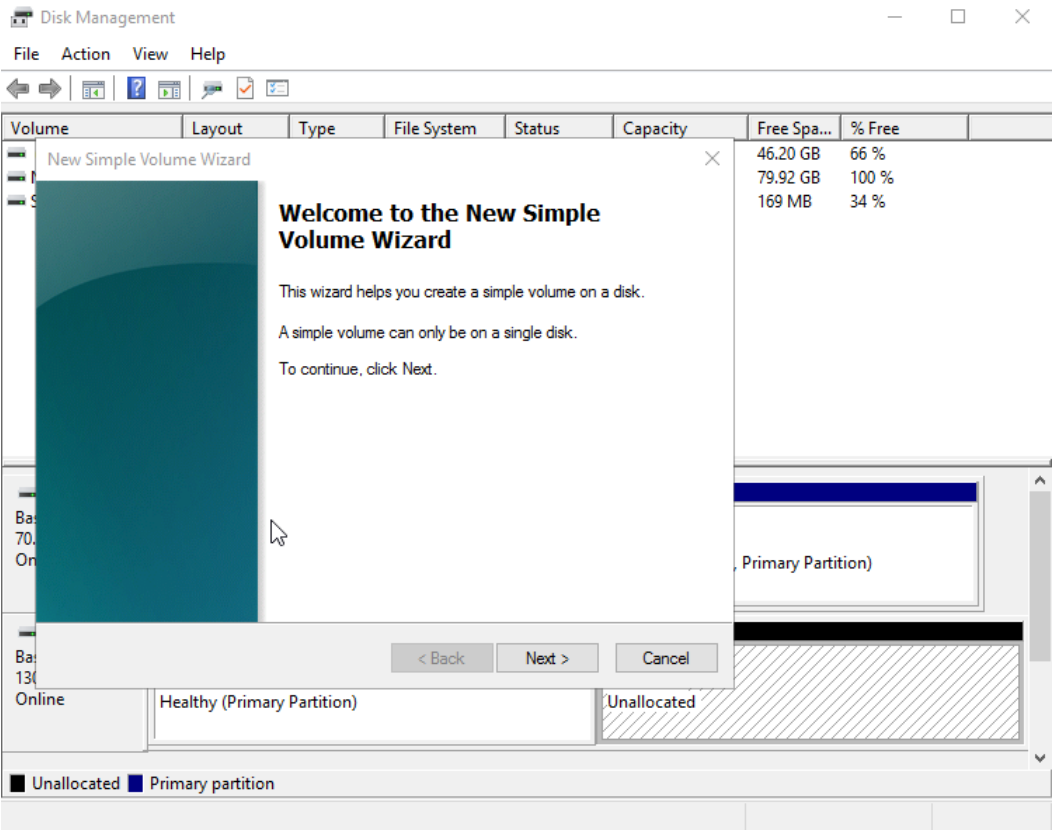
Step 2 In the **Unallocated** area of **Disk 1**, right-click the blank area and choose **New Simple Volume**.

Figure 3-59 New Simple Volume (Windows Server 2016 data disk)



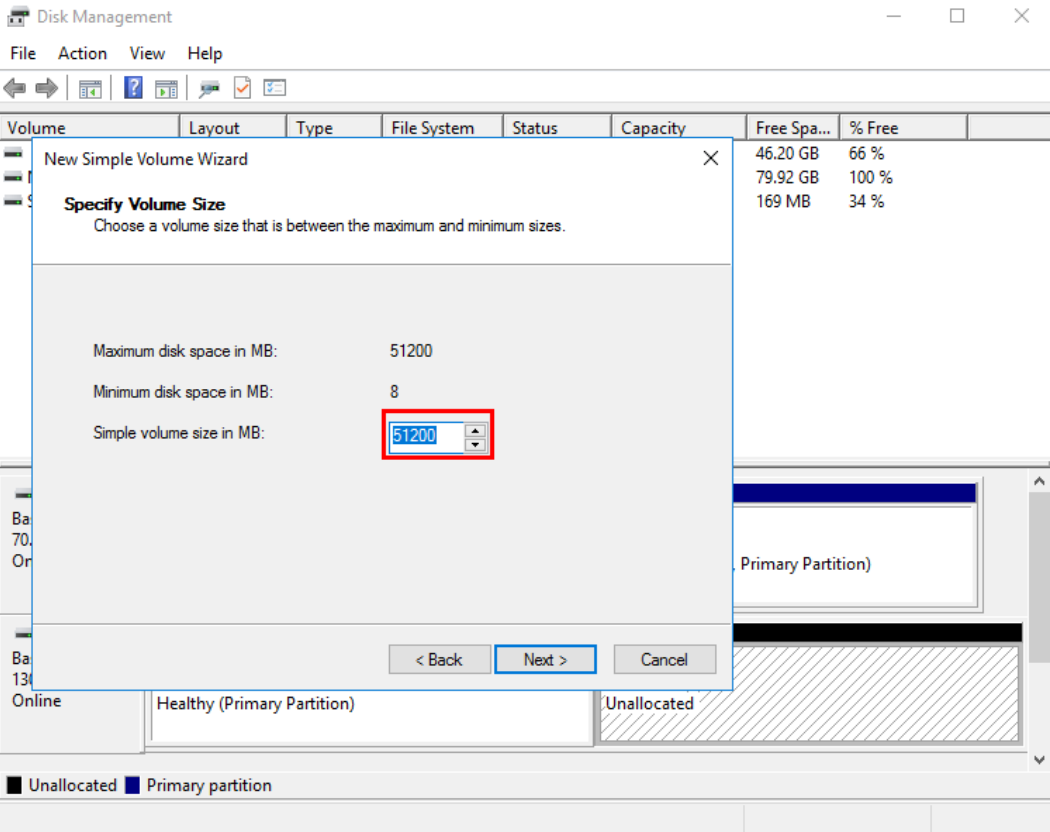
Step 3 On the displayed **New Simple Volume Wizard** window, click **Next**.

Figure 3-60 New Simple Volume Wizard (Windows Server 2016 data disk)




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

Figure 3-61 Specify Volume Size (Windows Server 2016 data disk)



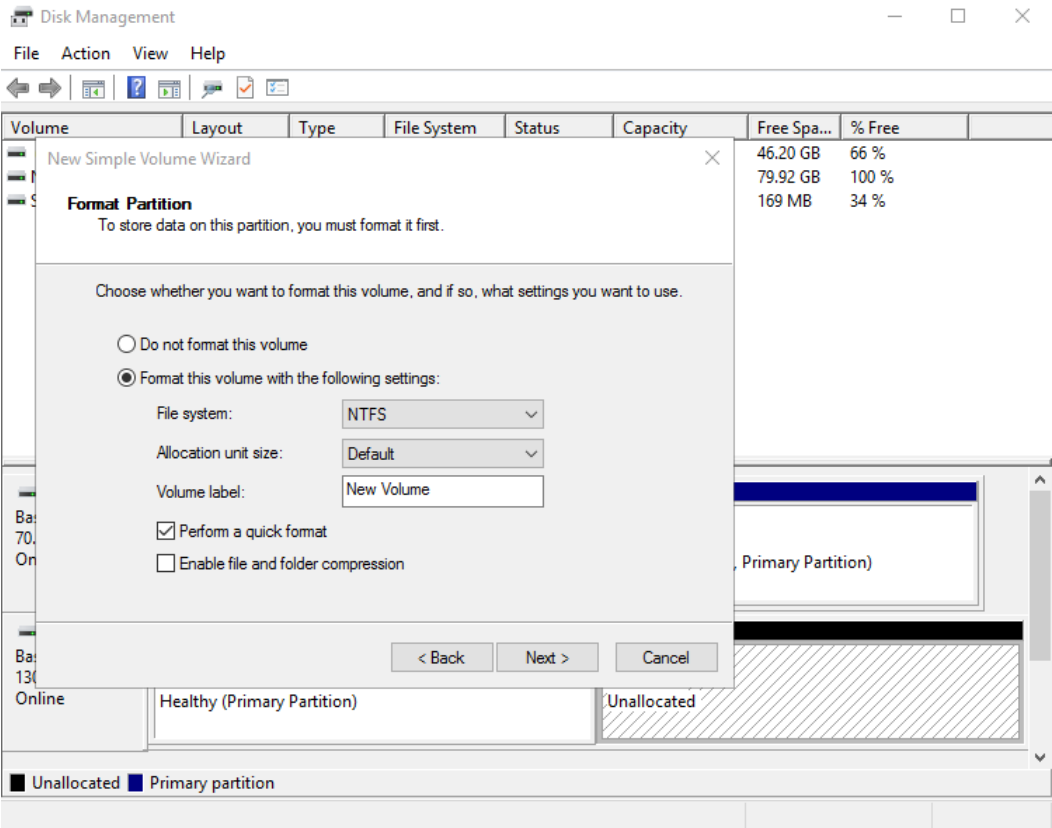
Step 5 On the displayed **Assign Drive Letter and Path** page, click **Assign the following drive letter**, select a drive letter, and click **Next**. In this example, drive letter **E** is selected.

Figure 3-62 Assign Drive Letter or Path (Windows Server 2016 data disk)



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

Figure 3-63 Format Partition (Windows Server 2016 data disk)



Step 7 Click **Finish**.

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

Figure 3-64 Completed

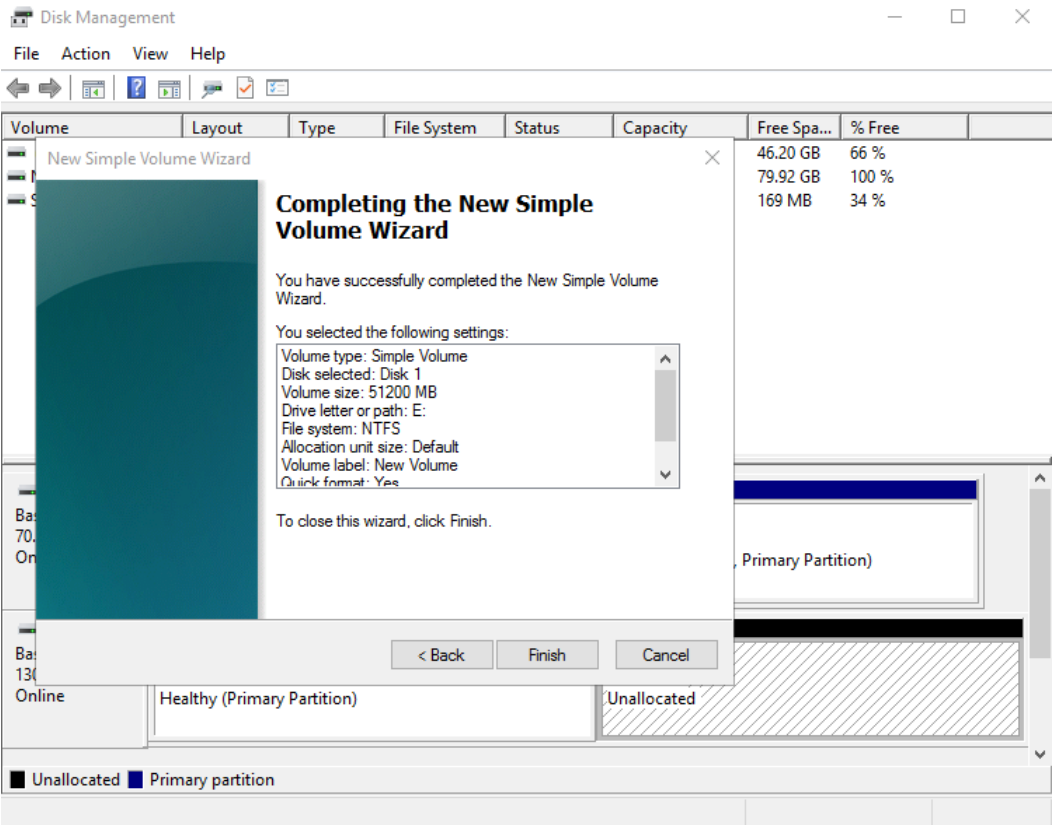
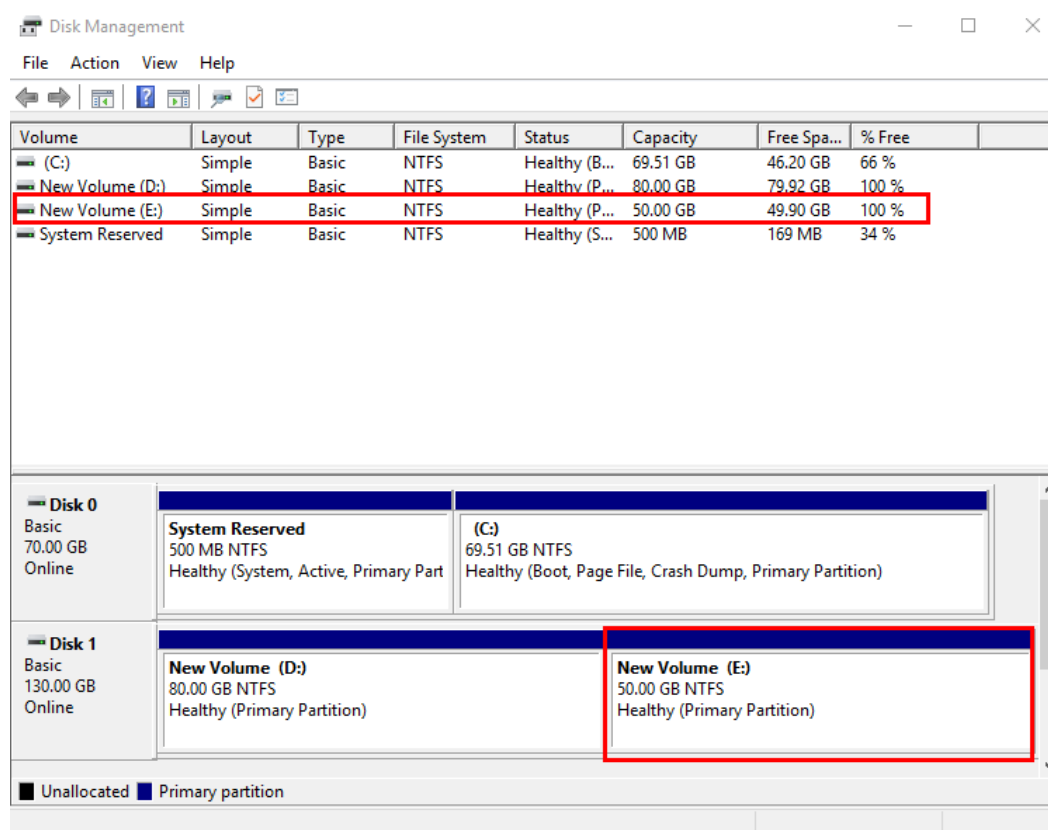


Figure 3-65 New Volume (E:)

----End

3.6 Extending Disk Partitions and File Systems (Linux)

3.6.1 Partition and File System Extension Preparations (Linux)

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

- To view the disk partition style, see the following methods:
 - [Method 1: Check Partition Style and File System Format Using fdisk](#)
 - [Method 2: Check Partition Style and File System Format Using parted](#)
- To extend disk partitions and file systems, see [Table 3-3](#).

NOTICE

- The following operation guide applies only to the Linux OS whose kernel version is 3.6.0 or later. You can run `uname -a` to check the Linux kernel version. If your kernel version is earlier than 3.6.0, see [Extending Disk Partitions and File Systems \(Linux Kernel Earlier Than 3.6.0\)](#).
- The additional space can only be added to the last partition of the disk.

Table 3-3 Disk partition and file system extension scenarios

Disk	Scenario	Method
System disk	Allocate the additional space to an existing MBR partition.	Extending an Existing MBR Partition
	Create a new MBR partition with the additional space.	Creating a New MBR Partition
Data disk	Allocate the additional space to an existing MBR partition.	Extending an Existing MBR or GPT Partition
	Create a new MBR partition with the additional space.	Creating a New MBR Partition
	Allocate the additional space to an existing GPT partition.	Extending an Existing MBR or GPT Partition
	Create a new GPT partition with the additional space.	Creating a New GPT Partition

 **NOTE**

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

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

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

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

lsblk

Information similar to the following is displayed:

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

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

 **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?](#) and expand the capacity. Otherwise, the additional space cannot be used after expansion.

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

fdisk -l

Information similar to the following is displayed:

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

Disk /dev/vda: 42.9 GiB, 42949672960 bytes, 83886080 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x000bcb4e

   Device Boot      Start         End      Blocks   Id  System
/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](#), existing partitions uses 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?](#) and expand the capacity. Otherwise, the additional space cannot be used after expansion.

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

parted Disk

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

parted /dev/vdb

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# parted /dev/vdb
GNU Parted 3.1
Using /dev/vdb
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) p
Error: The backup GPT table is not at the end of the disk, as it should be. This might mean that another
operating system believes the
disk is smaller. Fix, by moving the backup to the end (and removing the old backup)?
Fix/Ignore/Cancel? Fix
Warning: Not all of the space available to /dev/vdb appears to be used, you can fix the GPT to use all of
the space (an extra 104857600
blocks) or continue with the current setting?
Fix/Ignore? Fix
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 161GiB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

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

(parted)
```

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

- If the following error information is displayed, enter **Fix**.
Error: The backup GPT table is not at the end of the disk, as it should be. This might mean that another operating system believes the disk is smaller. Fix, by moving the backup to the end (and removing the old backup)?
The GPT partition table information is stored at the start of the disk. To reduce the risk of damage, a backup of the information is saved at the end of the disk. When you expand the disk capacity, the end of the disk changes accordingly. In this case, enter **Fix** to move the backup file of the information to new disk end.
- If the following warning information is displayed, enter **Fix**.
Warning: Not all of the space available to /dev/vdb appears to be used, you can fix the GPT to use all of the space (an extra 104857600 blocks) or continue with the current setting?
Fix/Ignore? Fix
Enter **Fix** as prompted. The system automatically sets the GPT partition style for the additional space.

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

----End

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

Scenarios

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

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

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

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

For how to extend the partitions and file systems on a BMS system disk, see [How Do I Increase the Size of the Root Partition of a BMS Which Is Quickly Provisioned?](#)

- [Extending an Existing MBR Partition](#)
- [Creating a New MBR Partition](#)

NOTICE

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

Prerequisites

- You have expanded the disk capacity and attached the disk to a server on the management console. For details, see [Expanding Capacity for an In-use EVS Disk](#) or [Expanding Capacity for an Available EVS Disk](#).
- You have logged in to the server.
 - For how to log in to an ECS, see [Logging In to an ECS](#).
 - For how to log in to a BMS, see [Logging In to a BMS](#).
- The Linux kernel version is 3.6.0 or later.

You can run **uname -a** to check the Linux kernel version. If your kernel version is earlier than 3.6.0, see [Extending Disk Partitions and File Systems \(Linux Kernel Earlier Than 3.6.0\)](#).

Constraints

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

Extending an Existing MBR Partition

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

Step 1 Run the **growpart** command to check whether growpart has been installed.

If the tool usage instructions are returned, the tool has been installed, and you do not need to install it again. Skip step 2.

Step 2 Run the following command to install the growpart tool:

yum install cloud-utils-growpart

Step 3 Run the following command to view the total capacity of the **/dev/vda** system disk:

fdisk -l

Information similar to the following is displayed:

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

```
Disk /dev/vda: 107.4 GiB, 107374182400 bytes, 209715200 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x000bcb4e
```

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

Step 4 Run the following command to view the capacity of the **/dev/vda1** partition:

df -TH

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/vda1       ext4      43G   2.0G   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   9.0M   2.0G    1% /run
tmpfs           tmpfs     2.0G    0   2.0G    0% /sys/fs/cgroup
tmpfs           tmpfs     398M    0   398M    0% /run/user/0
```

Step 5 Run the following command to extend the partition using growpart:

growpart *System disk Partition number*

In this example, run the following command:

growpart /dev/vda 1

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# growpart /dev/vda 1
CHANGED: partition=1 start=2048 old: size=83884032 end=83886080 new: size=209713119,end=209715167
```

NOTE

If error message "NOCHANGE:partition 1 is size xxxxxx. it cannot be grown" is returned, the expansion may be failed because the server disk is full (100% usage). Back up the disk data and clear necessary files or programs.

Step 6 Run the following command to extend the file system of the partition:

resize2fs *Disk partition*

In this example, run the following command:

resize2fs /dev/vda1

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# resize2fs /dev/vda1
resize2fs 1.42.9 (28-Dec-2013)
Filesystem at /dev/vda1 is mounted on /; on-line resizing required
old_desc_blocks = 5, new_desc_blocks = 13
The filesystem on /dev/vda1 is now 26214139 blocks long.
```

NOTE

If the error message "open: No such file or directory while opening /dev/vdb1" is returned, an incorrect partition is specified. Run **df -TH** to view the disk partitions.

Step 7 Run the following command to view the new capacity of the **/dev/vda1** partition:

df -TH

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/vda1       ext4     106G   2.0G   99G    2% /
devtmpfs        devtmpfs  2.0G    0   2.0G    0% /dev
tmpfs           tmpfs     2.0G    0   2.0G    0% /dev/shm
tmpfs           tmpfs     2.0G   9.0M   2.0G    1% /run
tmpfs           tmpfs     2.0G    0   2.0G    0% /sys/fs/cgroup
tmpfs           tmpfs     398M    0   398M    0% /run/user/0
```

----End

Creating a New MBR Partition

Originally, system disk **/dev/vda** has 40 GiB and one partition (**/dev/vda1**), and then 40 GiB is added to the disk. The following procedure shows you how to create a new MBR partition **/dev/vda2** with this 40 GiB.

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

fdisk -l

Information similar to the following is displayed:

```
[root@ecs-2220 ~]# fdisk -l

Disk /dev/vda: 85.9 GiB, 85899345920 bytes, 167772160 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x0008d18f
```

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

In the command output, the capacity of the **dev/vda** system disk is 80 GiB, in which the in-use **dev/vda1** partition takes 40 GiB and the additional 40 GiB has not been allocated.

Step 2 Run the following command to enter fdisk:

fdisk /dev/vda

Information similar to the following is displayed:

```
[root@ecs-2220 ~]# fdisk /dev/vda
Welcome to fdisk (util-linux 2.23.2).

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

Command (m for help):
```

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

Information similar to the following is displayed:

```
Command (m for help): n
Partition type:
  p   primary (1 primary, 0 extended, 3 free)
  e   extended
```

There are two types of disk partitions:

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

NOTE

If MBR is used, a maximum of 4 primary partitions, or 3 primary partitions plus 1 extended partition can be created. The extended partition must be divided into logical partitions before use.

Disk partitions created using GPT are not categorized.

Step 4 In this example, a primary partition is created. Therefore, enter **p** and press **Enter** to create a primary partition.

Information similar to the following is displayed:

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

- Step 5** Enter the serial number of the primary partition and press **Enter**. Partition number 2 is used in this example. Therefore, enter 2 and press **Enter**.

Information similar to the following is displayed:

```
Partition number (2-4, default 2): 2
First sector (83886080-167772159, default 83886080):
```

- Step 6** Enter the new partition's start sector and press **Enter**. In this example, the default start sector is used.

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

Information similar to the following is displayed:

```
First sector (83886080-167772159, default 83886080):
Using default value 83886080
Last sector, +sectors or +size{K,M,G} (83886080-167772159,default 167772159):
```

- Step 7** Enter the new partition's end sector and press **Enter**. In this example, the default end sector is used.

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

Information similar to the following is displayed:

```
Last sector, +sectors or +size{K,M,G} (83886080-167772159,
default 167772159):
Using default value 167772159
Partition 2 of type Linux and of size 40 GiB is set
```

Command (m for help):

- Step 8** Enter **p** and press **Enter** to view the new partition.

Information similar to the following is displayed:

```
Command (m for help): p

Disk /dev/vda: 85.9 GiB, 85899345920 bytes, 167772160 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x0008d18f
```

Device	Boot	Start	End	Blocks	Id	System
/dev/vda1	*	2048	83886079	41942016	83	Linux
/dev/vda2		83886080	167772159	41943040	83	Linux

Command (m for help):

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

Information similar to the following is displayed:

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

Calling ioctl() to re-read partition table.

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

The partition is created.

 **NOTE**

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

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

partprobe

Step 11 Run the following command to set the file system format for the new partition:

mkfs -t *File system Disk partition*

- Sample command of the ext* file system:

(The ext4 file system is used in this example.)

mkfs -t ext4 /dev/vda2

Information similar to the following is displayed:

```
[root@ecs-2220 ~]# mkfs -t ext4 /dev/vda2
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
2621440 inodes, 10485760 blocks
524288 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2157969408
320 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000, 7962624
```

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

- Sample command of the xfs file system:

mkfs -t xfs /dev/vda2

Information similar to the following is displayed:

```
[root@ecs-2220 ~]# mkfs -t xfs /dev/vda2
meta-data=/dev/vda2            isize=512    agcount=4, agsize=2621440 blks
        =                       sectsz=512   attr=2,    projid32bit=1
        =                       crc=1        finobt=0, sparse=0
data      =                       bsize=4096   blocks=10485760, imaxpct=25
        =                       sunit=0      swidth=0 blks
naming    =version2             bsize=4096   ascii-ci=0 ftype=1
log        =internal log        bsize=4096   blocks=5120, version=2
        =                       sectsz=512   sunit=0 blks, lazy-count=1
realtime  =none                 extsz=4096   blocks=0, rtextents=0
```

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

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

Perform this step if you want to mount the partition on a new mount point.

mkdir *Mount point*

In this example, run the following command to create the **/opt** mount point:

```
mkdir /opt
```

Step 13 Run the following command to mount the new partition:

```
mount Disk partition Mount point
```

In this example, run the following command to mount the new partition **/dev/vda2** on **/opt**:

```
mount /dev/vda2 /opt
```

NOTE

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

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

```
df -TH
```

Information similar to the following is displayed:

```
[root@ecs-2220 ~]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/vda1       ext4      43G   2.0G   39G   5% /
devtmpfs        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/vda2       ext4      43G   51M   40G   1% /opt
```

NOTE

If the server is restarted, the mounting will become invalid. You can modify the **/etc/fstab** file to configure automount at startup. See the following part for details.

----End

Configuring Automatic Mounting at System Start

The **fstab** file controls what disks are automatically mounted at server startup. You can configure the **fstab** file of a server that has data. This operation will not affect the existing data.

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

NOTE

UUIDs are the unique character strings for identifying partitions in Linux.

Step 1 Query the partition UUID.

blkid *Disk partition*

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

blkid /dev/vdb1

Information similar to the following is displayed:

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

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

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

vi /etc/fstab

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

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

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

The preceding information is used for reference only. The line starting with **UUID** is the information added. Edit this line from left to right to match the following format:

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

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

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

The system saves the configurations and exits the vi editor.

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

1. Unmount the partition.

umount *Disk partition*

In this example, run the following command:

```
umount /dev/vdb1
```

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

```
mount -a
```

3. Query the file system mounting information.

```
mount | grep Mount point
```

In this example, run the following command:

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

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

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

----End

3.6.3 Extending Partitions and File Systems for Data Disks (Linux)

Scenarios

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

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

This section uses CentOS 7.4 64bit as the sample OS to describe how to extend an MBR or GPT partition. The method for allocating the additional space varies with the server OS. This section is used for reference only. For detailed operations and differences, see the corresponding OS documents.

- [Extending an Existing MBR or GPT Partition](#)
- [Creating a New MBR Partition](#)
- [Creating a New GPT Partition](#)

NOTICE

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

Prerequisites

- You have expanded the disk capacity and attached the disk to a server on the management console. For details, see [Expanding Capacity for an In-use EVS Disk](#) or [Expanding Capacity for an Available EVS Disk](#).
- You have logged in to the server.
 - For how to log in to an ECS, see [Logging In to an ECS](#).

- For how to log in to a BMS, see [Logging In to a BMS](#).
- The Linux kernel version is 3.6.0 or later.
You can run **uname -a** to check the Linux kernel version. If your kernel version is earlier than 3.6.0, see [Extending Disk Partitions and File Systems \(Linux Kernel Earlier Than 3.6.0\)](#).

Constraints

- The additional space of a data disk cannot be added to the root partition. To extend the root partition, expand the system disk instead.
- The additional space can only be added to the last partition of the disk.

Extending an Existing MBR or GPT Partition

Originally, data disk **/dev/vdb** has 100 GiB and one partition (**/dev/vdb1**), and then 50 GiB is added to the disk. The following procedure shows you how to add this 50 GiB to the existing MBR or GPT partition (**/dev/vdb1**).

Step 1 Run the **growpart** command to check whether growpart has been installed.

If the tool usage instructions are returned, the tool has been installed, and you do not need to install it again. Skip step 2.

Step 2 Run the following command to install the growpart tool:

```
yum install cloud-utils-growpart
```

Step 3 Run the following command to install the gdisk software package:

```
yum install gdisk
```

The following information is displayed:

```
Is this ok [y/d/N]:
```

Enter **Y** and press **Enter** to complete the installation.

Step 4 Run the following command to view the disk partition information:

```
fdisk -l
```

Information similar to the following is displayed:

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

```
Disk /dev/vda: 42.9 GiB, 42949672960 bytes, 83886080 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x000bcb4e
```

Device	Boot	Start	End	Blocks	Id	System
/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

Step 5 Run the following command to view the size and file system type of the **/dev/vdb1** partition:

df -TH

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
devtmpfs        devtmpfs  509M   0 509M   0% /dev
tmpfs           tmpfs     520M   0 520M   0% /dev/shm
tmpfs           tmpfs     520M  7.1M 513M   2% /run
tmpfs           tmpfs     520M   0 520M   0% /sys/fs/cgroup
/dev/vda1       ext4      43G   2.3G 38G    6% /
tmpfs           tmpfs     104M   0 104M   0% /run/user/0
/dev/vdb1       ext4     106G   63M 101G   1% /mnt/sdc
```

NOTE

In the command output, the **Type** column shows the file system type, the **Size** column shows the total partition size, and the **Avail** column shows the partition's available space. In this example, the file system type of the **/dev/vdb1** partition is ext4, the total partition size is 106 GiB, and the partition's available space is 101 GiB.

Step 6 Run the following command to extend the partition using growpart:

growpart *Data disk Partition number*

In this example, run the following command:

growpart /dev/vdb 1

Information similar to the following is displayed:

```
[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 information similar to the following is displayed:
no tools available to resize disk with 'gpt'
FAILED: failed to get a resizer for id "
Install the gdisk software package according to [Step 3](#).
- If information similar to the following is displayed:
growpart /dev/vda 1 unexpected output in sfdisk --version [sfdisk is from util-linux 2.23.2]
Check that the system character set (language environment) is **en_US.UTF-8**.

Step 7 Extend the partition file system based on the file system type obtained in [Step 5](#). The commands vary depending on file system types.

- For the **ext*** file system, run the following command:

resize2fs *Disk partition*

In this example, run the following command:

resize2fs /dev/vdb1

Information similar to the following is displayed:

```
[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 **df -TH** to view the disk partitions.

- For the **xfs** file system, run the following command:

sudo xfs_growfs Mount point

In this example, run the following command:

sudo xfs_growfs /mnt/sdc

Information similar to the following is displayed:

```
[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
```

NOTICE

In the example command, **/mnt/sdc** is the mount point on which **/dev/vdb1** is mounted. Specify the mount point based on your site requirements.

Step 8 Run the following command to view the new capacity of the **/dev/vdb1** partition:

df -TH

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
devtmpfs        devtmpfs  509M   0 509M   0% /dev
tmpfs           tmpfs     520M   0 520M   0% /dev/shm
tmpfs           tmpfs     520M  7.1M 513M   2% /run
tmpfs           tmpfs     520M   0 520M   0% /sys/fs/cgroup
/dev/vda1       ext4      43G   2.3G 38G    6% /
tmpfs           tmpfs     104M   0 104M   0% /run/user/0
/dev/vdb1       ext4     156G   63M 151G   1% /mnt/sdc
```

----End

Creating a New MBR Partition

Originally, data disk **/dev/vdb** has 100 GiB and one partition (**/dev/vdb1**), and then 50 GiB is added to the disk. The following procedure shows you how to create a new MBR partition **/dev/vdb2** with this 50 GiB.

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

fdisk -l

Information similar to the following is displayed:

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

Disk /dev/vda: 42.9 GiB, 42949672960 bytes, 83886080 sectors
Units = sectors of 1 * 512 = 512 bytes
```

```
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x000bcb4e

   Device Boot      Start         End      Blocks   Id  System
/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
```

Step 2 Run the following command to enter fdisk:

fdisk *Disk*

In this example, run the following command:

fdisk /dev/vdb

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# fdisk /dev/vdb
Welcome to fdisk (util-linux 2.23.2).

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

Command (m for help):
```

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

Information similar to the following is displayed:

```
Command (m for help): n
Partition type:
   p   primary (1 primary, 0 extended, 3 free)
   e   extended
Select (default p):
```

There are two types of disk partitions:

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

NOTE

If MBR is used, a maximum of 4 primary partitions, or 3 primary partitions plus 1 extended partition can be created. The extended partition must be divided into logical partitions before use.

Disk partitions created using GPT are not categorized.

Step 4 In this example, a primary partition is created. Therefore, enter **p** and press **Enter** to create a primary partition.

Information similar to the following is displayed:

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

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

- Step 5** Enter the serial number of the primary partition and press **Enter**. Partition number **2** is used in this example. Therefore, enter **2** and press **Enter**.

Information similar to the following is displayed:

```
Partition number (2-4, default 2): 2
First sector (209715200-314572799, default 209715200):
```

First sector indicates the start sector. The value ranges from **209715200** to **314572799**, and the default value is **209715200**.

- Step 6** Enter the new partition's start sector and press **Enter**. In this example, the default start sector is used.

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

Information similar to the following is displayed:

```
First sector (209715200-314572799, default 209715200):
Using default value 209715200
Last sector, +sectors or +size{K,M,G} (209715200-314572799, default 314572799):
```

Last sector indicates the end sector. The value ranges from **209715200** to **314572799**, and the default value is **314572799**.

- Step 7** Enter the new partition's end sector and press **Enter**. In this example, the default end sector is used.

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

Information similar to the following is displayed:

```
Last sector, +sectors or +size{K,M,G} (209715200-314572799, default 314572799):
Using default value 314572799
Partition 2 of type Linux and of size 50 GiB is set
```

Command (m for help):

- Step 8** Enter **p** and press **Enter** to view the new partition.

Information similar to the following is displayed:

Command (m for help): p

```
Disk /dev/vdb: 161.1 GiB, 161061273600 bytes, 314572800 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x38717fc1
```

Device	Boot	Start	End	Blocks	Id	System
/dev/vdb1		2048	209715199	104856576	83	Linux
/dev/vdb2		209715200	314572799	52428800	83	Linux

Command (m for help):

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

Information similar to the following is displayed:

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

Calling ioctl() to re-read partition table.

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

NOTE

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

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

partprobe

Step 11 Run the following command to set the file system format for the new partition:

mkfs -t *File system Disk partition*

- Sample command of the ext* file system:

mkfs -t ext4 /dev/vdb2

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# mkfs -t ext4 /dev/vdb2
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
3276800 inodes, 13107200 blocks
655360 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2162163712
400 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000, 7962624, 11239424

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

- Sample command of the xfs file system:

mkfs -t xfs /dev/vdb2

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# mkfs -t xfs /dev/vdb2
meta-data=/dev/vdb2          isize=512    agcount=4, agsize=3276800 blks
       =                   sectsz=512   attr=2, projid32bit=1
       =                   crc=1        finobt=0, sparse=0
data      =                   bsize=4096   blocks=13107200, imaxpct=25
       =                   sunit=0       swidth=0 blks
naming    =version2          bsize=4096   ascii-ci=0 ftype=1
log        =internal log     bsize=4096   blocks=6400, version=2
       =                   sectsz=512   sunit=0 blks, lazy-count=1
realtime  =none              extsz=4096   blocks=0, rtextents=0
```

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

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

Perform this step if you want to mount the partition on a new mount point.

mkdir *Mount point*

In this example, run the following command to create the **/mnt/test** mount point:

```
mkdir /mnt/test
```

Step 13 Run the following command to mount the new partition:

```
mount Disk partition Mount point
```

In this example, run the following command to mount the new partition **/dev/vdb2** on **/mnt/test**:

```
mount /dev/vdb2 /mnt/test
```

NOTE

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

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

```
df -TH
```

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# df -TH
Filesystem      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
/dev/vdb2       ext4      53G   55M   50G    1% /mnt/test
```

NOTE

If the server is restarted, the mounting will become invalid. You can modify the **/etc/fstab** file to configure automount at startup. For details, see [Configuring Automatic Mounting at System Start](#).

----End

Creating a New GPT Partition

Originally, data disk **/dev/vdb** has 100 GiB and one partition (**/dev/vdb1**), and then 50 GiB is added to the disk. The following procedure shows you how to create a new GPT partition **/dev/vdb2** with this 50 GiB.

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

```
lsblk
```

Information similar to the following is displayed:

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

Step 2 Run the following command to enter parted:

parted *Disk*

In this example, run the following command:

parted */dev/vdb*

Information similar to the following is displayed:

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

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

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

Information similar to the following is displayed:

```
(parted) unit s
(parted) p
Error: The backup GPT table is not at the end of the disk, as it should be. This might mean that another
operating system believes the
disk is smaller. Fix, by moving the backup to the end (and removing the old backup)?
Fix/Ignore/Cancel? Fix
Warning: Not all of the space available to /dev/vdb appears to be used, you can fix the GPT to use all of
the space (an extra 104857600
blocks) or continue with the current setting?
Fix/Ignore? Fix
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 314572800s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start  End      Size    File system  Name  Flags
1       2048s  209713151s  209711104s  ext4         test

(parted)
```

In the command output, take note of the partition's end sector. In this example, the end sector of the **/dev/vdb1** partition is **209713151s**.

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

Error: The backup GPT table is not at the end of the disk, as it should be. This might mean that another operating system believes the disk is smaller. Fix, by moving the backup to the end (and removing the old backup)?

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

- If the following warning information is displayed, enter **Fix**.

Warning: Not all of the space available to /dev/vdb appears to be used, you can fix the GPT to use all of the space (an extra 104857600 blocks) or continue with the current setting?
Fix/Ignore? Fix

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

Step 5 Run the following command and press **Enter**:

mkpart *Partition name Start sector End sector*

In this example, run the following command:

mkpart data 209713152s 100%

In this example, the additional space is used to create a new partition. In [Step 4](#), the end sector of partition **dev/vdb1** is **209713151s**. Therefore, the start sector of the new partition **dev/vdb2** is set to **209713152s** and the end sector **100%**. This start and end sectors are for reference only. You can plan the number of partitions and partition size based on service requirements.

Information similar to the following is displayed:

```
(parted) mkpart data 209713152s 100%
(parted)
```

NOTE

The maximum sector can be obtained in either of the following ways:

- Query the disk's maximum end sector. For details, see [Step 2](#) to [Step 4](#).
- Enter **-1s** or **100%**, and the value displayed is the maximum end sector.

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

Information similar to the following is displayed:

```
(parted) p
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 314572800s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start      End          Size         File system  Name  Flags
  1      2048s     209713151s  209711104s  ext4         test
  2      209713152s 314570751s  104857600s               data

(parted)
```

Step 7 Run the following command to set the file system format for the new partition:

mkfs -t *File system Disk partition*

- Sample command of the ext* file system:

mkfs -t ext4 /dev/vdb2

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# mkfs -t ext4 /dev/vdb2
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
3276800 inodes, 13107200 blocks
655360 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2162163712
400 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000, 7962624, 11239424

Allocating group tables: done
Writing inode tables: done
```

```
Creating journal (32768 blocks): done  
Writing superblocks and filesystem accounting information: done
```

- Sample command of the xfs file system:

mkfs -t xfs /dev/vdb2

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# mkfs -t xfs /dev/vdb2  
meta-data=/dev/vdb2          isize=512    agcount=4, agsize=3276800 blks  
=                               sectsz=512   attr=2, projid32bit=1  
=                               crc=1      finobt=0, sparse=0  
data      =                   bsize=4096   blocks=13107200, imaxpct=25  
=                               sunit=0     swidth=0 blks  
naming    =version2           bsize=4096   ascii-ci=0 ftype=1  
log       =internal log      bsize=4096   blocks=6400, version=2  
=                               sectsz=512   sunit=0 blks, lazy-count=1  
realtime  =none              extsz=4096   blocks=0, rtextents=0
```

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

Step 8 (Optional) Run the following command to create a mount point:

Perform this step if you want to mount the partition on a new mount point.

mkdir *Mount point*

In this example, run the following command to create the **/mnt/test** mount point:

mkdir /mnt/test

Step 9 Run the following command to mount the new partition:

mount *Disk partition Mount point*

In this example, run the following command to mount the new partition **/dev/vdb2** on **/mnt/test**:

mount /dev/vdb2 /mnt/test

NOTE

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

Step 10 Run the following command to view the mount result:

df -TH

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# df -TH  
Filesystem      Type      Size  Used Avail Use% Mounted on  
/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  
/dev/vdb2       ext4      53G   55M   50G   1% /mnt/test
```

 NOTE

If the server is restarted, the mounting will become invalid. You can modify the `/etc/fstab` file to configure automount at startup. For details, see [Configuring Automatic Mounting at System Start](#).

----End

Configuring Automatic Mounting at System Start

The **fstab** file controls what disks are automatically mounted at server startup. You can configure the **fstab** file of a server that has data. This operation will not affect the existing data.

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

 NOTE

UUIDs are the unique character strings for identifying partitions in Linux.

Step 1 Query the partition UUID.

blkid *Disk partition*

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

blkid /dev/vdb1

Information similar to the following is displayed:

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

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

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

vi /etc/fstab

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

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

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

The preceding information is used for reference only. The line starting with **UUID** is the information added. Edit this line from left to right to match the following format:

- **UUID:** The UUID obtained in [Step 1](#).
- **Mount point:** The directory on which the partition is mounted. You can query the mount point using **df -TH**.
- **Filesystem:** The file system format of the partition. You can query the file system format using **df -TH**.
- **Mount option:** The partition mount option. Usually, this parameter is set to **defaults**.

- **Dump:** The Linux dump backup option.
 - **0:** Linux dump backup is not used. Usually, dump backup is not used, and you can set this parameter to **0**.
 - **1:** Linux dump backup is used.
- **fsck:** The fsck option, which means whether to use fsck to check the disk during startup.
 - **0:** not use fsck.
 - If the mount point is the root partition (/), this parameter must be set to **1**.
If this parameter is set to **1** for the root partition, this parameter for other partitions must start with **2** because the system checks the partitions in the ascending order of the values.

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

The system saves the configurations and exits the vi editor.

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

1. Unmount the partition.

umount *Disk partition*

In this example, run the following command:

umount /dev/vdb1

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

mount -a

3. Query the file system mounting information.

mount | grep *Mount point*

In this example, run the following command:

mount | grep /mnt/sdc

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

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

----End

4 Changing the EVS Disk Type (OBT)

Scenarios

If the performance of an existing disk no longer meets your service requirements, you can change the disk type to improve the disk performance.

NOTE

This function is in the OBT. [Submit a ticket](#) to apply for the OBT.

Constraints

Table 4-1 Constraints on the disk type change

Phase	Description
Before the change	<ul style="list-style-type: none">• The disk type can be changed only when the disk status is Available or In-use.• The disk type cannot be changed when any snapshot of the disk is being deleted.• Changing the disk type may affect the disk performance, so change the type during off-peak hours.

Phase	Description
During the change	<ul style="list-style-type: none">Some operations cannot be performed on the disk. Such operations include creating snapshots, creating backups, expanding the disk capacity, rolling back data from a snapshot, restoring data from a backup, attaching or detaching the disk, deleting the disk, transferring the disk, and creating an image from the ECS.Changing the disk type may take several hours or even longer, and cannot be stopped. The time depends on the throughput, storage space, and original disk type at the time of the change.In rare cases, the change may fail due to resource problems. In this case, you are advised to perform the change again.You can have a maximum of 10 disks with their types being changed at the same time.The OS cannot be changed if you are changing the disk type of a system disk.
After the change	In rare cases, the disk type may fail to be changed due to a backend issue. If this happens, try again later.

For details about the disk type change, see [Table 4-2](#).

Table 4-2 Supported changes between disk types

Source Disk Type	Description
Extreme SSD V2	Can be assigned with a different IOPS.
Extreme SSD	Can be changed to Ultra-high I/O or General Purpose SSD.
General Purpose SSD V2	Can be assigned with a different IOPS or throughput or both.
Ultra-high I/O	Can be changed to Extreme SSD or General Purpose SSD.
General Purpose SSD	Can be changed to Extreme SSD or Ultra-high I/O.
High I/O	Can be changed to Extreme SSD, Ultra-high I/O, or General Purpose SSD.
Common I/O (previous generation product)	Can be changed to Extreme SSD, Ultra-high I/O, General Purpose SSD, or High I/O.

Impact on the System

Read and write operations on the disk are not affected.

Procedure

Step 1 Log in to the [management console](#).

Step 2 Under **Storage**, click **Elastic Volume Service**.

The disk list page is displayed.

Step 3 In the disk list, locate the target disk, click **More** in the **Operation** column, and choose **Modify Specifications**.

The **Modify Specifications** page is displayed.

Step 4 Select a disk type from the drop-down list.

The system shows you the new price based on your selection.

Step 5 Click **Submit**.

The disk list is displayed, and the disk status is **Changing disk type**, indicating that the disk type is being changed. After the disk type changes to the target type, the operation is successful.

----End

5 Detaching an EVS Disk

5.1 Detaching a System Disk

Scenarios

If the file system on your system disk is damaged and your 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.

If you no longer need a system disk or want to replace it with a new one, you can detach the system disk.

A detached system disk will not be automatically deleted, and it will still be billed. To avoid unintended charges, you can delete or unsubscribe from the disk if it is no longer needed.

A system disk can only be detached offline, which means that its server must be in the **Stopped** state before the system disk is detached. To detach a system disk from a running server, you must first stop the server and then detach the disk.

NOTE

- For an attached system disk, the disk function is displayed as **System disk**, and the disk status is displayed as **In-use** in the disk list. After a system disk is detached from the server, the disk function changes to **Bootable disk**, and the status changes to **Available**.
- Bootable disks are the system disks detached from servers. A bootable disk can be re-attached to a server and be used as a system disk or data disk depending on the disk function selected.

Constraints

- A system disk can only be detached offline, which means that its server must be in the **Stopped** state before you can detach it. To detach a system disk from a running server, you must first stop the server and then detach the disk.
- After a system disk is detached, some operations cannot be performed on the original server and the system disk. The restricted operations 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
- System disk: changing disk billing mode

Procedure

Step 1 Log in to the [management console](#).

Step 2 Under **Compute**, click **Elastic Cloud Server**.

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

Step 3 In the server list, locate the row that contains the server whose system disk is to be detached, click **More** in the **Operation** column, and choose **Stop**.

When the server status changes to **Stopped**, the server has been stopped.

Step 4 Click the name of this server.

The server details page is displayed.

Step 5 Click the **Disks** tab to view the system disk attached to the server.

Step 6 Locate the row that contains the system disk and click **Detach**.

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

Step 7 Click **Yes** to detach the disk.

After the operation had succeeded, the detached system disk is no longer displayed under the **Disks** tab.

Step 8 (Optional) Re-attach the bootable disk to a server. You can use it as a system disk or data disk depending on the disk function you select.

- To re-attach and use it as a system disk, see [Attaching an Existing System Disk](#).
- To re-attach and use it as a data disk, see [Attaching an Existing Non-Shared Disk](#).

----End

Related Operations

For more detachment FAQs, see [Detachment](#).

5.2 Detaching a Data Disk

Scenarios

If you want to use a data disk on another server in the same region and AZ, you can detach the data disk and then attach it to that server.

If a data disk is no longer required, you can detach it and then delete it.

Data disks can be detached online or offline, meaning that the server using the to-be-detached data disk can either be in the **Running** or **Stopped** state.

- ECS

Detach an EVS disk from a running server. For details, see [Detaching an EVS Disk from a Running ECS](#).

- BMS

SCSI disks can be attached to BMSs and used as data disks. You can detach a data disk either from a running or stopped BMS.

NOTE

For an attached data disk, the disk function is displayed as **Data disk**, and the disk status is displayed as **In-use** in the disk list. After the data disk has been detached from the server, the disk function remains unchanged, the disk status changes to **Available** for a non-shared data disk, and the disk status changes to **Available** for a shared data disk after it is detached from all its servers.

Precautions

Data may be lost after you detach an encrypted disk. For more information, see [If I Detach a Disk, Will I Lose the Data on My Disk?](#)

Prerequisites

- Before detaching an EVS disk from a running Windows ECS, ensure that no programs are reading data from or writing data to the disk. Otherwise, data will be lost.
- Before detaching an EVS disk from a running Linux ECS, you must log in to the ECS and run the **umount** command to cancel the association between the disk and the file system. In addition, ensure that no programs are reading data from or writing data to the disk. Otherwise, detaching the disk will fail.

Detaching a Non-shared Disk

Step 1 Log in to the [management console](#).

Step 2 Under **Storage**, click **Elastic Volume Service**.

The disk list page is displayed.

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

- If yes, perform the following procedure:
 - a. In the disk list, click the name of the to-be-detached disk.
The disk details page is displayed.
 - b. Click the **Servers** tab to view the server where the target disk has been attached.
 - c. Click ☐ to select the server and click **Detach Disk**.
The **Detach Disk** dialog box is displayed.
 - d. Click **Yes** to detach the disk.

- If no, perform the following procedure:
 - a. In the disk list, locate the row that contains the target disk and choose **More > Detach** in the **Operation** column.
The **Detach Disk** dialog box is displayed.
 - b. Click **Yes** to detach the disk.

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

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

----End

6 Attaching an Existing Disk

6.1 Attaching an Existing System Disk

Scenarios

This section describes how to attach an existing system disk.

System disks can only be attached offline, which means that the server must be in the **Stopped** state.

You can view the disk function in the disk list. A disk can be attached to a server and used as the system disk only when its function is **Bootable disk** and its status is **Available**.

NOTE

- Bootable disks are the system disks detached from servers. A bootable disk can be re-attached to a server and be used as a system disk or data disk depending on the disk function selected.
- Non-shared, yearly/monthly disks purchased together with or later added to a yearly/monthly ECS cannot be attached to other ECSs.

Constraints

- Cloud servers created from ISO images are only used for OS installation. They have limited functions and cannot have EVS disks attached.
- A non-shared disk can be attached to one server only.
- The disk and the server must be in the same region and AZ.
- A detached, non-shared yearly/monthly data disk purchased together with a server can only be re-attached to the original server and be used as a data disk.
- A shared disk can be attached only when the servers' statuses are **Running** or **Stopped**.
- A frozen disk cannot be attached.
- A detached, yearly/monthly system disk purchased together with a server can be re-attached and used as a data disk for any server. If you want to use it again as a system disk, you must attach it to the original server.

- A detached, pay-per-use system disk purchased together with a server can be re-attached and used as a data disk for any server. If you want to use it again as a system disk, you must attach it to a server that uses the same image as the original server.

Procedure

Step 1 Log in to the [management console](#).

Step 2 Under **Storage**, click **Elastic Volume Service**.

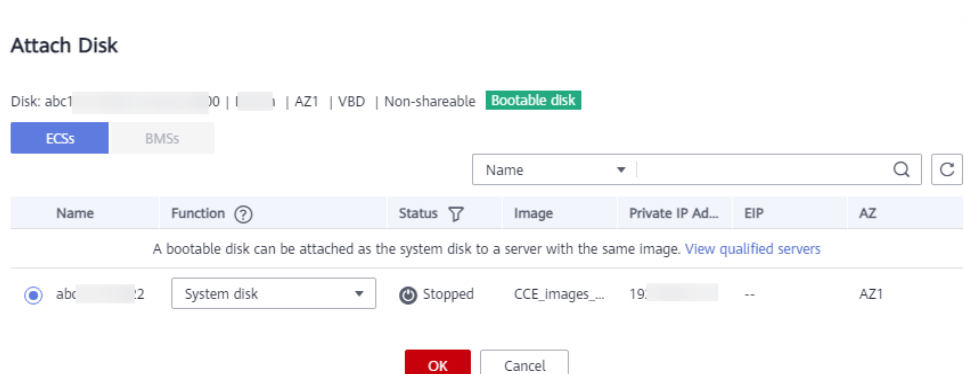
The disk list page is displayed.

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

Step 4 Select a server to attach the disk. Ensure that the disk and server are in the same AZ and that the server is in the **Stopped** state. After you select a server, **System disk** is automatically entered as the disk function for the server.

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

Figure 6-1 Attach Disk



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

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

----End

Related Operations

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

For more attachment FAQs, see [Attachment](#).

6.2 Attaching an Existing Non-Shared Disk

Scenarios

This section describes how to attach an existing non-shared disk to a server and use it as a data disk. A non-shared disk can be attached to one server only.

You can view the disk information in the disk list. A disk can be attached to a server and used as a data disk when all of the following conditions are met:

- Disk Sharing: Disabled
- Function: Bootable disk or Data disk
- Status: Available

NOTE

- Bootable disks are the system disks detached from servers. A bootable disk can be re-attached to a server and be used as a system disk or data disk depending on the disk function selected.
- Non-shared, yearly/monthly disks purchased together with or later added to a yearly/monthly ECS cannot be attached to other ECSs.

Constraints

- Cloud servers created from ISO images are only used for OS installation. They have limited functions and cannot have EVS disks attached.
- A non-shared disk can be attached to one server only.
- The disk and the server must be in the same region and AZ.
- A detached, non-shared yearly/monthly data disk purchased together with a server can only be re-attached to the original server and be used as a data disk.
- A shared disk can be attached only when the servers' statuses are **Running** or **Stopped**.
- A frozen disk cannot be attached.
- A detached, yearly/monthly system disk purchased together with a server can be re-attached and used as a data disk for any server. If you want to use it again as a system disk, you must attach it to the original server.
- A detached, pay-per-use system disk purchased together with a server can be re-attached and used as a data disk for any server. If you want to use it again as a system disk, you must attach it to a server that uses the same image as the original server.

Procedure

Step 1 Log in to the [management console](#).

Step 2 Under **Storage**, click **Elastic Volume Service**.

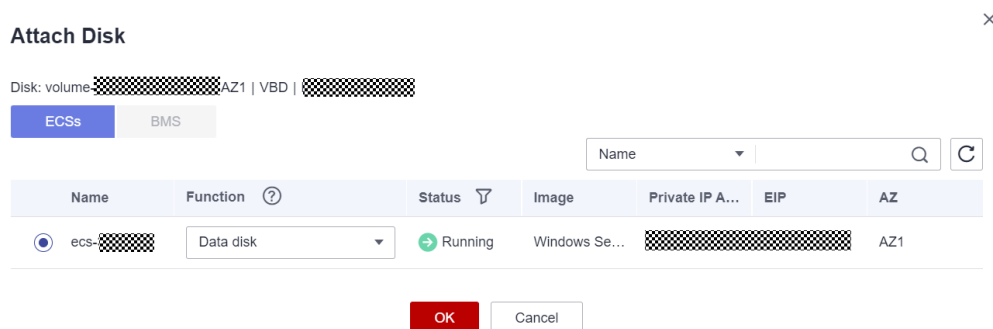
The disk list page is displayed.

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

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

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

Figure 6-2 Attach Disk



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

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

Step 6 (Optional) Mount the existing disk partition to a mount point if you have attached the disk to a Linux server. The mount command is as follows:

```
mount Disk partition Mount point
```

----End

Related Operations

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

For more attachment FAQs, see [Attachment](#).

6.3 Attaching an Existing Shared Disk

Scenarios

This section describes how to attach an existing shared disk to a server and use it as a data disk.

You can view the disk information in the disk list. A shared disk can be attached to servers and used as a data disk when all of the following conditions are met:

- Disk Sharing: Enabled
- Function: Data disk
- Status: In-use or Available

Constraints

NOTICE

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

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

Procedure

Step 1 Log in to the [management console](#).

Step 2 Under **Storage**, click **Elastic Volume Service**.

The disk list page is displayed.

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

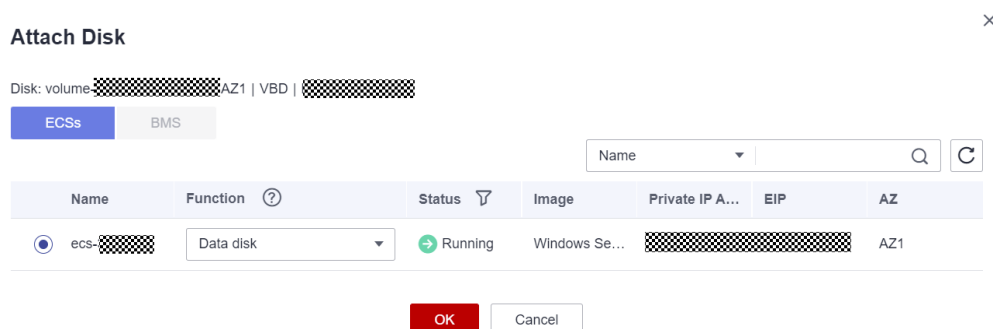
- Shared disks support batch attachment so that you can attach a shared disk to multiple servers. The left area in the **Attach Disk** dialog box shows the server list. After you select the target servers, the selected servers will be displayed in the right area.

- A shared disk can be attached to servers only when the disk status is **Available** or **In-use**.

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

One device name can be used for one disk only. If a device name has been used, it will no longer be displayed in the drop-down list and cannot be selected.

Figure 6-3 Attach Disk



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

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

Step 6 (Optional) Mount the existing disk partition to a mount point if you have attached the disk to a Linux server. The mount command is as follows:

```
mount Disk partition Mount point
```

----End

Related Operations

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

For more attachment FAQs, see [Attachment](#).

7 Deleting EVS Disks

Scenarios

If an EVS disk is no longer used, you can release the virtual resources by deleting it.

After a disk is deleted, you are no longer charged for it. EVS immediately destroys the metadata so that data can no longer be accessed. In addition, the physical storage space of the disk is reclaimed and cleared before being re-assigned. For any new disk created based on the re-assigned physical space, before data is written to the disk, EVS returns zero for all the read and write requests to the disk.

- Pay-per-use disk:

It can only be deleted when the following conditions are met:

- The disk status is **Available**, **Error**, **Expansion failed**, **Restoration failed**, or **Rollback failed**.
- The disk is not added to any replication pair in the Storage Disaster Recovery Service (SDRS). For any disk already added to a replication pair, you need to first [delete the replication pair](#) and then delete the disk.
- The disk is not locked by any service.
- The shared disk has been detached from all its servers.

- Yearly/Monthly disk:

It cannot be deleted, but you can unsubscribe from it if needed. For more information, see [Unsubscriptions](#).

The shared disk can be unsubscribed from when it 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.
-

Procedure

Step 1 Log in to the [management console](#).

Step 2 Under **Storage**, click **Elastic Volume Service**.

The disk list page is displayed.

Step 3 In the disk list, locate the row that contains the target disk, click **More** in the **Operation** column, and choose **Delete**.

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

Step 5 On the displayed page, confirm the deletion information.

- If operation protection is enabled, select a verification method and obtain and enter the verification code.

Supported verification methods include SMS, email, virtual MFA device. If none of these are associated, click **Associate**.

- If operation protection is not enabled, enter **DELETE** in the text box below.

For details about how to enable or disable operation protection, see [Operation Protection](#).

Step 6 Click **OK**.

----End

Related Operations

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

8 Unsubscribing from Yearly/Monthly EVS Disks

CAUTION

This function is available only in some regions. If you hover over the **Unsubscribe** button of a non-shared disk and the system displays , it means that yearly/monthly disks in this region must be unsubscribed from together with their servers and separate disk unsubscription is not supported.

Scenarios

This section describes how to unsubscribe from yearly/monthly EVS disks.

System disks must be unsubscribed from together with their servers.

For how to unsubscribe from data disks, see [Table 8-1](#).

Constraints

- Yearly/Monthly system disks cannot be unsubscribed from separately. They must be unsubscribed from together with their servers.
- Non-shared, yearly/monthly data disks purchased together with or later added to a yearly/monthly server have the same expiration time as the server. They can be unsubscribed from together with the server or separately when their statuses are **In-use**, **Available**, or **Error**.

Table 8-1 Unsubscription scenarios of data disks

Unsubscription Scenario	Sub-scenario	Reference
Unsubscribing from non-shared, yearly/ monthly data disks that were purchased together with or later added to a yearly/ monthly server	Unsubscribing from data disks when unsubscribing from the server	Unsubscribing from an ECS
	Unsubscribing from data disks separately	Unsubscribing from Disks on the EVS Console Unsubscribing from Disks on the ECS Console
Unsubscribing from shared, yearly/ monthly data disks that were purchased together with or later added to a yearly/ monthly server	Unsubscribing from data disks separately	Unsubscribing from Disks on the EVS Console

Unsubscribing from Disks on the EVS Console

Step 1 Log in to the [management console](#).

Step 2 Click . Under **Storage**, click **Elastic Volume Service**.

The disk list page is displayed.

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

NOTE

If the **Unsubscribe** button is grayed out, detach the disk and then unsubscribe from it.

Step 4 On the **Unsubscribe from Resource** page, confirm information again and select the reason for unsubscription. After confirming the information, select "After being unsubscribed from, the resources not in the recycle bin will be deleted immediately and cannot be restored. I've backed up data or no longer need the data" and click **Confirm**.

Step 5 Confirm the disks you want unsubscribe from and delete and click **Confirm**.

----End

Unsubscribing from Disks on the ECS Console

NOTE

- When you unsubscribe from a server, the yearly/monthly data disks purchased together with the server and those added later will also be unsubscribed from. For details, see [Unsubscribing from an ECS](#).
- To unsubscribe from non-shared data disks separately, perform the following steps:

Step 1 Log in to the [management console](#).

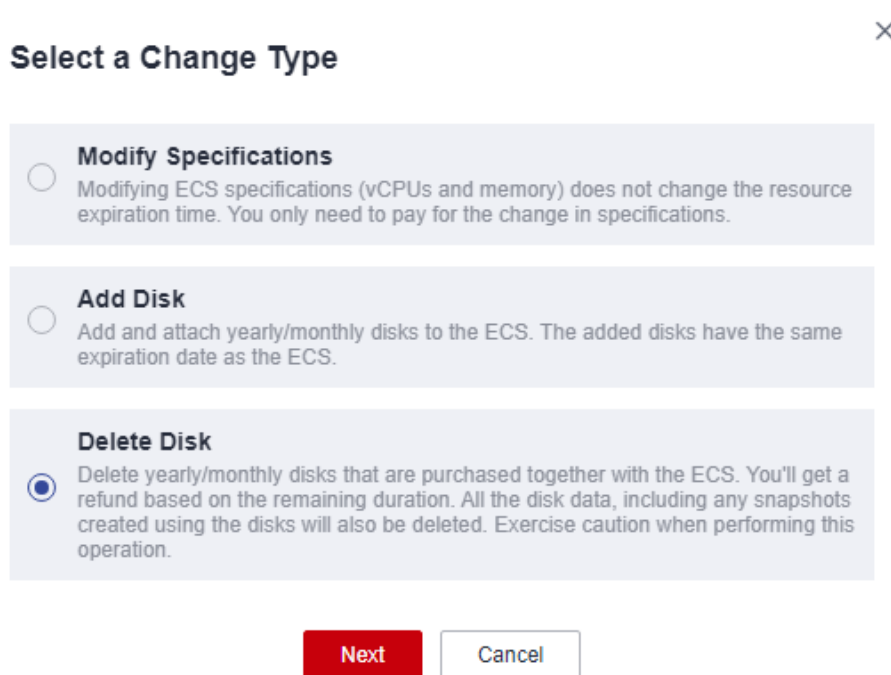
Step 2 Click . Under **Compute**, click **Elastic Cloud Server**.

The **ECS** console is displayed.

Step 3 In the server list, locate the target server, choose **More** > **Change** in the **Operation** column.

Step 4 In the displayed dialog box, select **Delete Disk**.

Figure 8-1 Selecting a change type



Step 5 Select the disks you want to delete and click **Next**.

Step 6 On the deletion page, confirm the information, select **I understand a handling fee will be charged for this unsubscription**, and click **Submit**.

----End

9 Managing EVS Recycle Bin

9.1 Recycle Bin Overview

EVS recycle bin is disabled by default. You need to manually enable it before you can use it.

If the recycle bin is enabled, EVS disks will be moved to the recycle bin upon deletion. This can help protect your disk data from accidental deletions.

To learn when deleted disks will be moved to the recycle bin, see [Recycle Bin Rules for Deleted EVS Disks](#).

You can configure a recycle bin policy to define when to move deleted disks to the recycle bin.

NOTE

EVS recycle bin is available in the following regions: LA-Lima1, LA-Sao Paulo1, LA-Buenos Aires1, and LA-Mexico City2

Constraints

- When you delete a disk, regardless of whether the disk will be moved to the recycle bin or not, the disk snapshots will always be deleted permanently.
- There are no limits on the capacity and quantity of disks in the recycle bin.
- In the recycle bin, disks can be stored for 7 days. After 7 days, the system will permanently delete the disks, and the disks cannot be recovered.

Recycle Bin Rules for Deleted EVS Disks

When will EVS disks be moved to the recycle bin?

You first have to configure a recycle bin policy to define when to move deleted disks to the recycle bin. Then, disks will be moved there if:

- You delete pay-per-use disks.
- You delete the cloud service resources that use the disks. The cloud services include ECS, BMS, CCE, and MRS.

- You reinstall the ECS OS, and the system automatically creates a new system disk and deletes the old system disk.

EVS disks will not be moved to the recycle bin if:

- You delete pay-per-use disks when your account is restricted or frozen.
- The system permanently deletes pay-per-use when your account is restricted or frozen.
- The number of days that the disks have been created is less than what you specified in the recycle bin policy.
- The pay-per-use disks you deleted are already in a grace or retention period.
- The system permanently deletes the pay-per-use disks whose retention period has expired.

Recycle Bin Operations

Table 9-1 Recycle bin operations

Operation	Description	Reference
Enable the recycle bin.	EVS recycle bin is disabled by default. You need to manually enable it before you can use it.	Enabling the Recycle Bin
Disable the recycle bin.	You can disable the recycle bin as needed. You must empty the recycle bin before disabling it.	Disabling the Recycle Bin
Configure a recycle bin policy.	You can configure a recycle bin policy to define when to move deleted disks to the recycle bin.	Configuring a Recycle Bin Policy
Recover disks from the recycle bin.	You can recover disks from the recycle bin.	Recovering Disks from the Recycle Bin
Permanently delete disks from the recycle bin.	In the recycle bin, disks can be stored for 7 days. Once expired, the system will permanently delete the disks.	Permanently Deleting Disks from the Recycle Bin

9.2 Enabling the Recycle Bin

Scenarios

EVS recycle bin is disabled by default. You need to manually enable it before you can use it.

If the recycle bin is enabled, EVS disks will be moved to the recycle bin upon deletion. This can help protect your disk data from accidental deletions.

Constraints

- To learn when will EVS disks be moved to the recycle bin after the recycle bin is enabled, see [Recycle Bin Rules for Deleted EVS Disks](#).
- When you delete a disk, regardless of whether the disk will be moved to the recycle bin or not, the disk snapshots will always be deleted permanently.
- There are no limits on the capacity and quantity of disks in the recycle bin.
- In the recycle bin, disks can be stored for 7 days. After 7 days, the system will permanently delete the disks, and the disks cannot be recovered.

Procedure

Step 1 Log in to the [management console](#).

Step 2 Under **Storage**, click **Elastic Volume Service** to go to the EVS console.

Step 3 Click the **Recycle Bin** tab.

Step 4 Click **Enable Recycle Bin**.

After the recycle bin is enabled, the **Recycle Bin** tab page shows the disks that are going to be permanently deleted. You can now use the recycle bin.

----End

9.3 Disabling the Recycle Bin

Scenarios

You can disable the recycle bin as needed.

Constraints

You must empty the recycle bin before disabling it. To empty the recycle bin, you can:

- Recover the disks from the recycle bin by referring to [Recovering Disks from the Recycle Bin](#).
- Permanently delete the disks from the recycle bin by referring to [Permanently Deleting Disks from the Recycle Bin](#).

Procedure

Step 1 Log in to the [management console](#).

Step 2 Under **Storage**, click **Elastic Volume Service**.

The disk list page is displayed.

Step 3 Click the **Recycle Bin** tab.

Step 4 In the upper right corner of the **Recycle Bin** tab page, click **Disable Recycle Bin**.

A dialog box is displayed.

Step 5 Click **Yes**.

The recycle bin has been disabled.

----End

9.4 Configuring a Recycle Bin Policy

Scenarios

If you have configured scaling policies for your workloads, the system may frequently delete EVS disks. At the same time, if you have also enabled EVS recycle bin, but do not want these frequently deleted disks to be moved to the recycle bin, you can configure a recycle bin policy to reduce unintended costs.

Example scenarios:

- Scenario 1: You have used Auto Scaling to dynamically scale services. The system may frequently delete disks based on the configured scaling policy. In this case, you can configure an appropriate recycle bin policy based on the scaling cycle to avoid moving any unintended disks to the recycle bin.
- Scenario 2: You have used Cloud Container Engine (CCE) to run workloads. The system may frequently delete disks based on the configured container scaling policy. In this case, you can configure an appropriate recycle bin policy based on the scaling cycle to avoid moving any unintended disks to the recycle bin.

These examples are for your reference only. You can configure an appropriate policy based on your own service scenario.

Constraints

The number of days you can set for the recycle bin policy ranges from 7 to 1,000 days and defaults to 7 days.

When the OS of an ECS is reinstalled, a new system disk will be created to replace the original one. EVS determines whether to move the original disk to the recycle bin by comparing the number of days elapsed since disk creation with the days configured in the recycle bin policy.

Procedure

Step 1 Log in to the [management console](#).

Step 2 Under **Storage**, click **Elastic Volume Service**.

The disk list page is displayed.

Step 3 Click the **Recycle Bin** tab.

Step 4 In the upper right corner of the **Recycle Bin** tab page, click **Configure Policy**.

The **Configure Recycle Bin Policy** page is displayed.

Step 5 Set the minimum number of days that must elapse after a disk was created before it can be moved to the recycle bin upon deletion or unsubscription.

 **NOTE**

For example, if you set 7 days for the recycle bin policy, EVS disks created within 7 days will not be moved to the recycle bin upon deletion and will be deleted permanently. EVS disks created at least 7 days ago will be moved to the recycle bin upon deletion and billed on a pay-per-use basis.

Step 6 Click **Confirm**.

----End

9.5 Recovering Disks from the Recycle Bin

Scenarios

In the recycle bin, disks can be stored for 7 days. During this period, you can recover them from the recycle bin if needed.

Constraints

- Disks in the recycle bin are all billed on a pay-per-use basis, regardless of their billing modes before deletion. Pay-per-use billing applies after disks are recovered from the recycle bin.

If you demand yearly/monthly billing for a recovered disk, attach it to an ECS or a BMS, and then change the server's billing mode to yearly/monthly.

To learn how to change pay-per-use billing to yearly/monthly, see [From Pay-per-Use to Yearly/Monthly](#).

- If your account is frozen or restricted, disks in the recycle bin cannot be recovered.

Procedure

Step 1 Log in to the [management console](#).

Step 2 Under **Storage**, click **Elastic Volume Service**.

The disk list page is displayed.

Step 3 Click the **Recycle Bin** tab.

Step 4 On the **Recycle Bin** tab page, locate the disk you want to recover and click **Recover** in the **Operation** column.

The **Recover Disk** page is displayed.

Step 5 Click **Submit**.

- If the recovery succeeds, the disk will be displayed in the disk list, and the disk status is **Available**.
- If the recovery fails, the disk remains in the recycle bin, and the disk status changes to **Recovery failed**.

----End

9.6 Permanently Deleting Disks from the Recycle Bin

Scenarios

You can permanently delete the EVS disks from the recycle bin at any time.

Constraints

- In the recycle bin, disks can be stored for 7 days. After 7 days, the system will permanently delete the disks, and the disks cannot be recovered.
- If you have disks already in the recycle bin and then your account goes in arrears, the disks will enter a grace period and then a retention period, and may be kept for less than 7 days before they are deleted permanently by the system.



Once disks are deleted from the recycle bin, data on them cannot be recovered.

Procedure

Step 1 Log in to the [management console](#).

Step 2 Under **Storage**, click **Elastic Volume Service**.

The disk list page is displayed.

Step 3 Click the **Recycle Bin** tab.

Step 4 Locate the disk you want to permanently delete and click **Delete** in the **Operation** column.

The confirmation dialog box is displayed.

Step 5 Click **Yes**.

If the disk disappears from the recycle bin, the disk has been permanently deleted.

----End

10 Viewing EVS Disk Details

Scenarios

This section describes how to view disk details, including the disk status and specifications. You can view disk details:

- [From the EVS Console](#)
- [From the Cloud Server Console](#)

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

NOTE

You can view the disks in the disk list even if your account is in arrears.

From the EVS Console

Step 1 Log in to the [management console](#).

Step 2 Under **Storage**, click **Elastic Volume Service**.

The disk list page is displayed.


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

You can filter disks by project, status, disk name, or tag.

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.

Step 5 (Optional) Export disk information.

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

----End

From the Cloud Server Console


Step 1 Log in to the [management console](#).

Step 2 Under **Compute**, click **Elastic Cloud Server**.

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

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

The server details page is displayed.

Step 4 On the **Disks** tab, click  in front of the row containing the target disk. In the unfolded area, click the disk ID.

The disk details page is displayed for you to view the disk details.

----End

11

Managing Encrypted EVS Disks

Encryption Scenarios

- **System disk encryption**

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

Table 11-1 Encryption relationship between images and system disks

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

- **Data disk encryption**

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

Table 11-2 Encryption relationship between backups, snapshots, images, and data disks

Purchased On	Method of Purchase	Whether Data Disk Will Be Encrypted	Description
The ECS console	Purchased together with the server	Yes/No	When a data disk is purchased together with a server, you can choose to encrypt the disk or not. For details, see Getting Started > Creating an ECS > Step 1: Configure Basic Settings in the <i>Elastic Cloud Server User Guide</i> .
The EVS console	No data source selected	Yes/No	When an empty disk is created, you can choose whether to encrypt the disk or not. The encryption attribute of the disk cannot be changed after the disk has been created.
	Creating from a backup	Yes/No	<ul style="list-style-type: none">When a disk is created from a backup, you can choose whether to encrypt the disk or not. The encryption attributes of the disk and backup do not need to be the same.When you create a backup for a system or data disk, the encryption attribute of the backup will be the same as that of the disk.
	Creating from a snapshot (The snapshot's source disk is encrypted.)	Yes	A snapshot created from an encrypted disk is also encrypted.
	Creating from a snapshot (The snapshot's source disk is not encrypted.)	No	A snapshot created from a non-encrypted disk is not encrypted.

Purchased On	Method of Purchase	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	-

Constraints

Table 11-3 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">• An encrypted disk cannot be changed to a non-encrypted disk.• A non-encrypted disk cannot be changed to an encrypted disk.
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">• If the user is the first user, the user needs to follow the prompt to create an agency, which grants KMS Administrator permissions to EVS. Then the user can create and obtain keys to encrypt and decrypt disks. <p>NOTE The first user must have the KMS Administrator permissions to create the agency. If the user does not have the KMS Administrator permissions, contact the account administrator to grant the permissions first.</p> <ul style="list-style-type: none">• If the user is not the first user, the user can use encryption directly.

Item	Description
Constraints on encrypted images	<ul style="list-style-type: none">• Encrypted images cannot be replicated across regions.• Encrypted images cannot be changed to non-encrypted images.• Encrypted images cannot be exported.

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.

Detaching an Encrypted EVS Disk

Before you detach a disk encrypted by a custom key, check whether the custom key is disabled or scheduled for deletion.

- If the custom key is available, the disk can be detached and re-attached, and data on the disk will not be lost.
- If the custom key 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 custom key.

The restoration method varies depending on the key status. For details, see [EVS Encryption](#).

For details about how to detach an encrypted disk, see [Detaching a Data Disk](#).

12 Managing Shared EVS Disks

How to Use Shared VBD and SCSI Disks

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

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

NOTICE

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

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

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

on ECSs that are in the same ECS group, thus having a working 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. The sharing function is not supported for 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, the sharing function will not work and data 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 EVS Disks](#).

Expanding a Shared EVS Disk

Shared EVS disks must be expanded when they are in the **Available** state. For details, see [Expanding Capacity for an In-use EVS Disk](#).

Related Operations

For more disk sharing FAQs, see [Sharing](#).

13 Managing EVS Backups

Scenarios

EVS backups are created using the CBR service.

You can configure a backup policy for disks. With backup policies configured, data on EVS disks can be periodically backed up to improve data security.

Constraints

- Backups can be created only when the disks are in the **Available** or **In-use** state.
- 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.

Purchasing a Disk Backup Vault and Applying a Backup Policy

Step 1 Log in to the CBR console.

1. Log in to the management console.
2. Choose **Storage > Cloud Backup and Recovery > Cloud Disk Backups**.

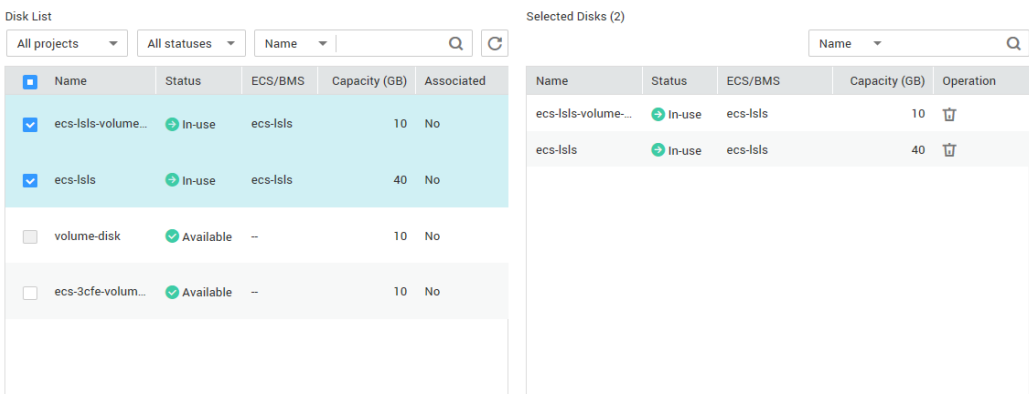
Step 2 In the upper right corner, click **Buy Disk Backup Vault**.

Step 3 Select a billing mode.

- Yearly/monthly is a prepaid billing mode. You are billed based on the subscription duration you specify. This mode provides lower prices and is ideal when the resource use duration is predictable.
- Pay-per-use is a postpaid billing mode. You are billed based on your resource usage. With this mode, you can increase or delete resources at any time. Fees are deducted from your account balance.

Step 4 (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 13-1 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 5 Specify a vault capacity ranging from the total sizes of disks to 10,485,760 GiB.

Step 6 Configure auto backup.

- If you select **Configure**, you must then select an existing backup policy or create a new one. After the vault is created, the system applies this backup policy to the vault, and all disks associated with this vault will be automatically backed up based on this policy.
- If you select **Skip**, disks associated with this vault are not automatically backed up.

Step 7 If you have subscribed to Enterprise Project, add the vault to an existing enterprise project.

An enterprise project facilitates project-level management and grouping of cloud resources and users. The default project is **default**.

Step 8 (Optional) Add tags to the vault.

A tag consists of a key-value pair. Tags are used to identify, classify, and search for vaults. Vault tags are used to filter and manage vaults only. You can add up to 10 tags for a vault.

Table 13-1 describes the tag parameters.

Table 13-1 Tag parameters

Parameter	Description	Example Value
Key	A tag key of a vault must be unique. You can customize the key or select the key of an existing tag created in TMS. A tag key: <ul style="list-style-type: none">• Can contain 1 to 36 Unicode characters.• Cannot be left blank, cannot start or end with spaces, or contain non-printable ASCII (0-31) characters or the following special characters: =*<>\ /	Key_0001
Value	A tag value can be repetitive or left blank. A tag value: <ul style="list-style-type: none">• Can contain 0 to 43 Unicode characters.• Can be an empty string, cannot start or end with spaces, or contain non-printable ASCII (0-31) characters or the following special characters: =*<>\ /	Value_0001

Step 9 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 10 Specify the subscription duration if you select **Yearly/Monthly** for **Billing Mode**. You can choose 1 month to 3 years for the usage duration.

Determine whether to enable auto renewal. If you select **Auto Renewal**:

- Your subscription will be renewed each month for monthly billing.
- Your subscription will be renewed each year for yearly billing.

Step 11 Click **Next**. Confirm the purchase details and continue.**Step 12** Complete the payment as prompted.**Step 13** Go back to the disk backup page. The vault you created is displayed in the list.

You can associate disks to the new vault or create backups for the disks. For details, see [Vault Management](#).

-----End

14 Managing EVS Snapshots (OBT)

14.1 Snapshot Overview (OBT)

What Is EVS Snapshot?

An EVS snapshot is a complete copy or image of the disk data at a specific point in time. Snapshots can be used as a disaster recovery (DR) approach, and you can use snapshots to fully restore data to the time when the snapshot was taken. You can create snapshots for disks on the console or via the API.

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

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

Application Scenarios

The snapshot function helps address your following needs:

- **Routine data backup**
You can create snapshots for disks on a timely basis and use snapshots to recover your data in case that data loss or data inconsistency occurred due to unintended operations, viruses, or attacks.
- **Rapid data restoration**
You can create a snapshot or multiple snapshots before an application software upgrade or a service data migration. If an exception occurs during the upgrade or migration, service data can be rapidly restored to the time point when the snapshot was created.
For example, a fault occurred on system disk A of server A, and therefore server A cannot be started. As system disk A is already faulty, data on system disk A cannot be restored by rolling back snapshots. But, you can create disk B using an existing snapshot of system disk A and attach disk B to a properly running server, for example server B. In this case, server B obtains the data of system disk A from disk B.

 **NOTE**

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

- Multi-service quick deployment

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

Operation Overview

You can create snapshots according to [Creating a Snapshot \(OBT\)](#) to rapidly save the disk data at specified points in time.

If a data loss happened, you can roll back the disk data to the time when the snapshot was created based on [Rolling Back Data from a Snapshot \(OBT\)](#). In addition, you can create a new disk from the snapshot so that the disk will contain the snapshot data in the beginning. For details, see [Creating an EVS Disk from a Snapshot \(OBT\)](#).

When a snapshot is no longer needed, delete it according to [Deleting a Snapshot \(OBT\)](#) to release the virtual resources.

Snapshot FAQ

For more snapshot FAQs, see [Snapshot](#).

14.2 Creating a Snapshot (OBT)

Scenarios

You can create an EVS snapshot on the management console to save the EVS disk data at a specific time point.

For details about the snapshot principle and application scenarios, see [EVS Snapshot \(OBT\)](#).

 **NOTE**

Creating snapshots does not affect the performance of the disk.

Constraints

- A maximum of 7 snapshots can be created for one disk.
- Snapshots can be created for both system disks and data disks.
- Snapshots can be created only for available or in-use disks.
- Snapshots of encrypted disks are stored encrypted, and those of non-encrypted disks are stored non-encrypted.
- The enterprise project of a snapshot is the same as that of the snapshot's source disk.

Creating a Snapshot on the Disks Page

Step 1 Log in to the [management console](#).

Step 2 Under **Storage**, click **Elastic Volume Service**.

The disk list page is displayed.

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

Configure the basic settings for the snapshot according to [Table 14-1](#).

Table 14-1 Snapshot parameter

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

Figure 14-1 Create Snapshot

The screenshot shows the 'Create Snapshot' dialog box. At the top, there is a note: 'You have created 1 snapshots for this disk and you can create 6 more snapshots. The disk you selected is not encrypted. Therefore, the snapshot data of this disk will not be encrypted.' Below the note, the 'Disk Specifications' section displays the following information:

Disk Name	pvc-0defe0cb-e980-440d-9eff-f6cc07deb000	Encrypted	No
Disk ID	336dfeac-871d-484b-a99d-7e431ca9d91b	Region	Shanghai1
Disk Type	High I/O	AZ	AZ1
Capacity (GB)	10		
Device Type	SCSI		

At the bottom, there is a 'Snapshot Name' field with the value 'snapshot-ee6d'.

Step 4 Click **Create Now**.

Step 5 Go back to the **Snapshots** page to view the snapshot creation information.

After the snapshot status changes to **Available**, the snapshot has been created.

----End

Creating a Snapshot on the Snapshots Page

Step 1 Log in to the management console.

Step 2 Under **Storage**, click **Elastic Volume Service**.

The disk list page is displayed.

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

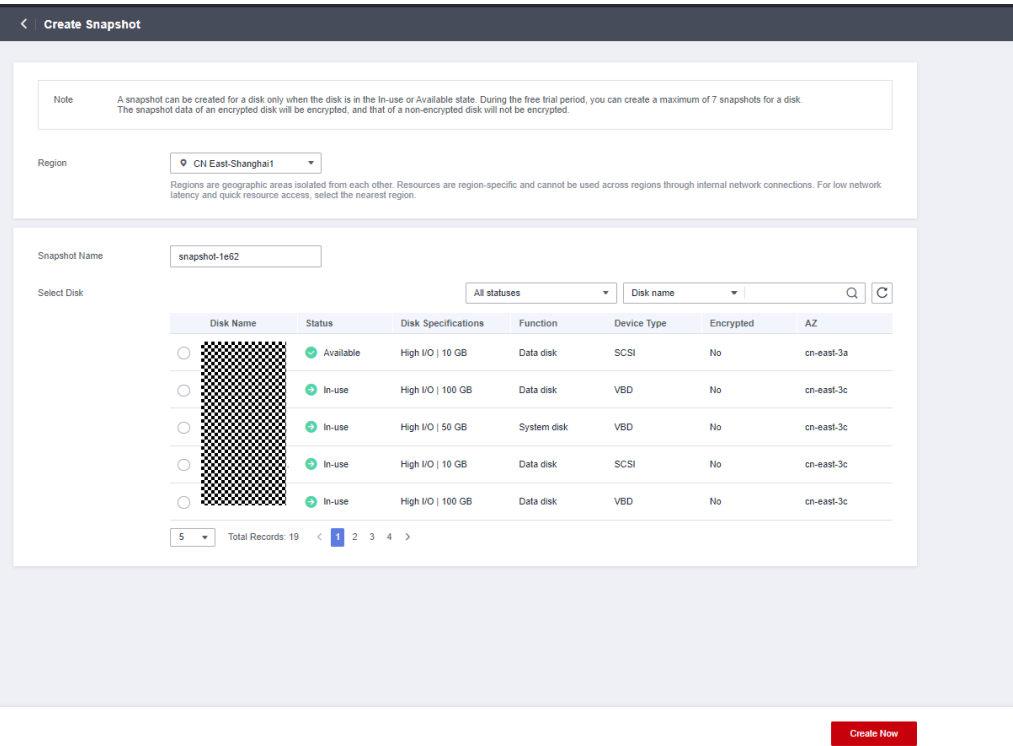
On the **Snapshots** page, click **Create Snapshot**.

Configure the basic settings for the snapshot according to [Table 14-2](#).

Table 14-2 Snapshot parameters

Parameter	Description	Example Value
Region	Mandatory After you select a region, disks in the selected region will be displayed for you to choose from.	-
Snapshot Name	Mandatory The name can contain a maximum of 64 characters.	snapshot-01
Select Disk	Mandatory Select a disk based on which the snapshot will be created.	volume-01

Figure 14-2 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

Snapshot FAQ

For more snapshot FAQs, see [Snapshot](#).

14.3 Deleting a Snapshot (OBT)

Scenarios

If a snapshot is no longer needed, you can delete it to release the virtual resources.

Constraints

- The snapshot status must be **Available** or **Error**.
- If a disk is deleted, all the snapshots created for this disk will also be deleted.
- If a snapshot is deleted, disks rolled back and created from this snapshot are not affected.
- If you have reinstalled or changed the server OS, snapshots of the system disk are automatically deleted. Snapshots of the data disks can be used as usual.
- A snapshot whose name starts with **autobk_snapshot_vbs_**, **manualbk_snapshot_vbs_**, **autobk_snapshot_csbs_**, or **manualbk_snapshot_csbs_** is automatically generated during backup. You can only view details of such snapshots and cannot delete them.

Procedure

Step 1 Log in to the [management console](#).

Step 2 Under **Storage**, click **Elastic Volume Service**.

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

The snapshot list page is displayed.

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

Step 5 (Optional) If multiple snapshots are to be deleted, select ☐ in front of each snapshot and click **Delete** in the upper area of the list.

Step 6 In the displayed dialog box, confirm the information and click **Yes**.

If the snapshot disappears from the snapshot list, the snapshot is deleted successfully.


----End

14.4 Viewing Details of a Snapshot (OBT)

Scenarios

This section describes how to view the details of a snapshot.

Procedure

- Step 1** Log in to the [management console](#).
- Step 2** Under **Storage**, click **Elastic Volume Service**.
- Step 3** In the navigation pane on the left, choose **Elastic Volume Service > Snapshots**.
The snapshot list page is displayed.
- Step 4** Locate the row that contains the target snapshot and click  in front of the snapshot name.
On the details page, view the snapshot information, such as the snapshot size.
- End

14.5 Rolling Back Data from a Snapshot (OBT)

Scenarios

If data on an EVS disk is incorrect or damaged, you can roll back data from a snapshot to the source disk.

Constraints

- A snapshot can be rolled back only to its source disk. Rollback to another disk is not possible.
- A snapshot can be rolled back only when the snapshot status is **Available** and its source disk status is **Available** (not attached to any server) or **Rollback failed**. If the source disk is attached, detach the disk first.
- A snapshot whose name starts with **autobk_snapshot_vbs_**, **manualbk_snapshot_vbs_**, **autobk_snapshot_csbs_**, or **manualbk_snapshot_csbs_** is automatically generated during backup. Such a snapshot can only be viewed. It cannot be used to roll back the disk data.
- If you first roll back the snapshot to the original disk, you cannot use the snapshot to create a new disk then.

Procedure

- Step 1** Log in to the [management console](#).
- Step 2** Under **Storage**, click **Elastic Volume Service**.
- Step 3** In the navigation pane on the left, choose **Elastic Volume Service > Snapshots**.
The snapshot list page is displayed.
- Step 4** In the snapshot list, locate the row that contains the target snapshot and click **Roll Back Disk** in the **Operation** column.
- Step 5** In the displayed dialog box, click **Yes**.

The snapshot list is displayed. After the snapshot status changes from **Rolling back** to **Available**, the data rollback is successful.

----End

14.6 Creating an EVS Disk from a Snapshot (OBT)

Scenarios

You can create new disks from snapshots by either selecting the target snapshot directly in the snapshot list or specifying parameter **Create from snapshot** on the disk creation page.

Constraints

- The disk type, device type, encryption, and snapshot attributes of the new disk are the same as those of the snapshot's source disk.
- Batch disk creation is not possible, and the quantity parameter must be set to 1.
- The AZ of the disk is the same as that of the snapshot's source disk and cannot be changed.
- A snapshot whose name starts with **autobk_snapshot_vbs_**, **manualbk_snapshot_vbs_**, **autobk_snapshot_csbs_**, or **manualbk_snapshot_csbs_** is automatically generated during backup. Such a snapshot can only be viewed. It cannot be used to create new disks.
- If you first roll back the snapshot to the original disk, you cannot use the snapshot to create a new disk then.

Procedure

Step 1 Log in to the [management console](#).

Step 2 Under **Storage**, click **Elastic Volume Service**.

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

The snapshot list page is displayed.

Step 4 In the snapshot list, locate the row that contains the target snapshot and click **Create Disk** in the **Operation** column.

Step 5 Configure the disk parameters.

NOTE

For details, see [Purchase an EVS Disk](#).

A maximum of 128 disks can be created from a snapshot.

- In the condition that you do not specify the disk capacity, if the snapshot size is smaller than 10 GiB, the default capacity 10 GiB will be used as the disk capacity; if the snapshot size is greater than 10 GiB, the snapshot size will be used as the disk capacity.
- To specify a disk capacity larger than the snapshot size, set the disk capacity in the **Disk Specifications** area.

Step 6 Click **Next**.

Step 7 Pay for the fees as prompted and click **OK**.

The disk list page is displayed.

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

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

----**End**

15 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 [EVS Transfer](#).

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.

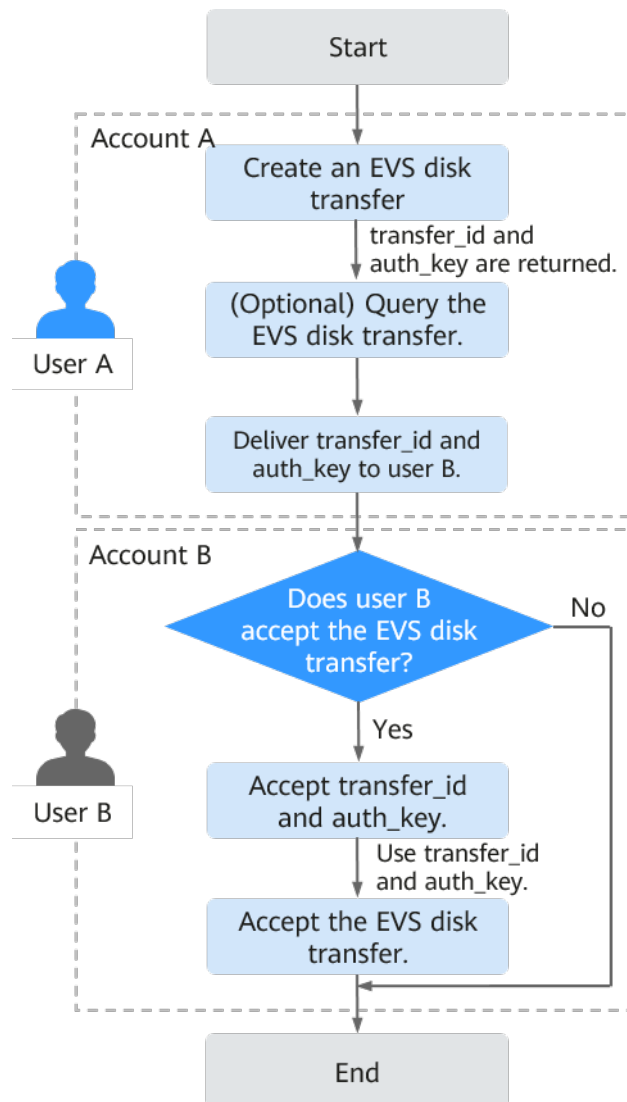
Procedure

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 15-1](#) 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 15-1 EVS disk transfer process



Step 1 User A creates an EVS disk transfer. For details, see [Creating a Disk Transfer](#).

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](#). 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](#).

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](#).

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](#).

----End

16 Managing a Tag

16.1 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 16-1 Tag overview

Operation	Scenario
Adding a Tag	Add tags for existing disks or during disk creations.
Modifying a Tag	Change tag values for existing disks. Tag keys of existing disks cannot be changed.
Deleting a Tag	Delete tags that are no longer needed for existing disks.
Searching for Disks by Tag	After tags are added, search for disks by tags.

16.2 Adding a Tag

Scenarios

You can add a tag for an existing EVS disk. You can also add tags during the disk creation.

- A tag is composed of a key-value pair.
 - A tag key can contain a maximum of 36 characters. It consists of letters, digits, underscores (_), hyphens (-), and Unicode characters (\u4E00-\u9FFF).

- A tag value can contain a maximum of 43 characters. It consists of letters, digits, underscores (_), periods (.), hyphens (-), and Unicode characters (\u4E00-\u9FFF).
- A maximum of 10 tags can be added for an EVS disk.

Procedure

- Step 1** Log in to the management console.
- Step 2** Under **Storage**, click **Elastic Volume Service**.
The disk list page is displayed.
- Step 3** In the disk list, locate the desired disk and click the disk name.
The disk details page is displayed.
- Step 4** Click the **Tags** tab.
- Step 5** Click **Add Tag**.
The **Add Tag** page is displayed.
- Step 6** Enter a key and a value for a tag and click **OK**.
 - **Key**: This parameter is mandatory.
 - **Value**: This parameter is optional.The **Tags** tab is displayed, and you can view the newly added tag.
- End

16.3 Modifying a Tag

Scenarios

You can change the value of a tag for an existing disk, but cannot change the key of a tag.

Procedure

- Step 1** Log in to the [management console](#).
- Step 2** Under **Storage**, click **Elastic Volume Service**.
The disk list page is displayed.
- Step 3** In the disk list, locate the desired disk and click the disk name.
The disk details page is displayed.
- Step 4** Click the **Tags** tab.
- Step 5** Locate the target tag and click **Edit** in the **Operation** column.
The **Edit Tag** page is displayed.
- Step 6** Change the value of the tag and click **OK**.

Return to the tag list. If the tag value is changed, the modification is complete.

-----End

16.4 Deleting a Tag

Scenarios

If an existing tag is no longer needed, you can delete it.

Procedure

- Step 1** Log in to the [management console](#).
 - Step 2** Under **Storage**, click **Elastic Volume Service**.
The disk list page is displayed.
 - Step 3** In the disk list, locate the desired disk and click the disk name.
The disk details page is displayed.
 - Step 4** Click the **Tags** tab.
 - Step 5** Locate the target tag and click **Delete** in the **Operation** column.
The **Delete Tag** page is displayed.
 - Step 6** Confirm the information and click **Yes**.
The tag is deleted if it disappears from the tag list.
- End


16.5 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 [management console](#).
- Step 2** Under **Storage**, click **Elastic Volume Service**.
The disk list page is displayed.
- Step 3** In the upper area of the disk list, click **Search by Tag**.
The **Search by Tag** page is displayed.
- Step 4** Enter or select an existing tag in the text box under **Search by Tag**.

Step 5 (Optional) If disks containing multiple tags need to be queried, click  to add tags.

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**

17 Changing EVS Disk Name

Scenarios

Disk names are used to identify disks. After a disk is created, you can perform operations in this section to change the disk name if needed.




Procedure

Step 1 Log in to the [management console](#).

Step 2 Under **Storage**, click **Elastic Volume Service**.

The disk list page is displayed.

Step 3 Change the disk name in either of the following ways:

- Perform the following steps to change the disk name in the disk list:
 - a. In the disk list, locate the target disk in the **Disk Name** column and click  to the right of the disk name.
The **Edit Disk Name** dialog box is displayed.
 - b. Enter a new name.
 - c. Click **OK**.
After the change is successful, the new disk name is displayed in the disk list.
- Perform the following steps to change the disk name on the disk details page:
 - a. In the disk list, locate the desired disk and click the disk name.
The disk details page is displayed.
 - b. Click  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

18 Viewing EVS Monitoring Data

Description

This section describes monitored metrics reported by EVS to Cloud Eye as well as their namespaces and dimensions. You can use the console or APIs provided by Cloud Eye to query the metrics of the monitored objects and alarms generated for EVS. For details about how to set alarms, see [Setting Alarm Rules](#).

Namespace

SYS.EVS

Metrics

Table 18-1 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	≥ 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	≥ 0 bytes/s	EVS disk	5 minutes in average
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

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period
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
disk_device_queue_length	Average Queue Length	Average number of read or write requests waiting for processing in the monitoring period for the monitored disk Unit: Count	≥ 0 Counts	EVS disk	5 minutes in average
disk_device_io_util	Disk I/O Utilization	Percentage of time spent during which read and write requests were sent to the monitored disk in the monitoring period Unit: Percent	0-100%	EVS disk	5 minutes in average
disk_device_write_bytes_per_operation	Avg Disk Bytes Per Write	Average number of bytes transmitted per I/O write for the monitored disk in the monitoring period Unit: Kbyte/operation	≥ 0 KiB/op	EVS disk	5 minutes in average
disk_device_read_bytes_per_operation	Avg Disk Bytes Per Read	Average number of bytes transmitted per I/O read for the monitored disk in the monitoring period Unit: Kbyte/operation	≥ 0 KiB/op	EVS disk	5 minutes in average
disk_device_write_await	Disk Write Await	Average await time per I/O write for the monitored disk in the monitoring period Unit: ms/operation	≥ 0 ms/operation	EVS disk	5 minutes in average

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period
disk_device_read_await	Disk Read Await	Average await time per I/O read for the monitored disk in the monitoring period Unit: ms/operation	≥ 0 ms/operation	EVS disk	5 minutes in average
disk_device_io_svctm	Disk I/O Service Time	Average service time per I/O read or write for the monitored disk in the monitoring period Unit: ms/operation	≥ 0 ms/operation	EVS disk	5 minutes in average
disk_device_io_iops_qos_num	IOPS Upper Limit Reached (Count)	Number of times that the IOPS of the monitored disk has reached the upper limit Unit: Count	≥ 0 Counts	EVS disk	5 minutes in average
disk_device_io_iobw_qos_num	Bandwidth Upper Limit Reached (Count)	Number of times that the bandwidth of the monitored disk has reached the upper limit Unit: Count	≥ 0 Counts	EVS disk	5 minutes in average

Dimension

Key	Value
disk_name	Server ID-drive letter, for example, 6f3c6f91-4b24-4e1b-b7d1-a94ac1cb011d-vda (vda is the drive letter)

Viewing Monitoring Data

Step 1 Log in to the [management console](#).

Step 2 Under **Storage**, click **Elastic Volume Service**.

The disk list page is displayed.

Step 3 In the EVS disk list, click the name of the disk you want to view the monitoring data.

The disk details page is displayed.

Step 4 On the **Servers** tab, locate the row that contains the server and click **View Metric** in the **Operation** column.

The **Monitoring Metrics** page is displayed.

Step 5 View the disk monitoring data by metric or monitored duration.

----End

19 Viewing EVS Monitoring Data (Agent Installed and Simplified Monitoring Metrics Used)

Description

This section describes the EVS-related metrics included in the OS metrics supported by ECS. The agent of the latest version is used with simplified monitoring metrics.

After installing the agent on an ECS, you can view its EVS-related metrics included in the OS monitoring metrics.

For instructions about how to install and configure the agent, see [Agent Installation and Configuration](#).

Monitoring Metrics

Table 19-1 EVS-related metrics

Metric	Name	Description	Value Range	Monitored Object	Monitoring Period
mountPointPrefix_disk_free	(Agent) Available Disk Space	<p>Available disk space on the monitored object</p> <p>Unit: GiB</p> <ul style="list-style-type: none">Linux: Run the df -h command to check the value in the Avail column. The path of the device name prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).Windows: Obtain the metric value using the WMI API GetDiskFreeSpaceExW. The path of the device name prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).	≥0 GiB	ECS	5 minutes in average

Metric	Name	Description	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: Disk Usage = Used Disk Space / Disk Storage Capacity.</p> <p>Unit: Percent</p> <ul style="list-style-type: none">Linux: Obtain the metric value using the following formula: Disk Usage = Used Disk Space / Disk Storage Capacity. The path of the device name prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).Windows: Obtain the metric value using the WMI API GetDiskFreeSpaceExW. The path of the device name prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).	0-100%	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">Linux: The disk I/O usage is calculated based on the data changes in the thirteenth column of the corresponding device in file /proc/diskstats in a collection period.The path of the device name prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).Windows does not support this metric.	0-100%	ECS	5 minutes in average

Dimensions

Key	Value
instance_id	Specifies the ECS ID.

20 Auditing

Scenarios

EVS supports the recording of EVS operations through CTS. You can query EVS traces and use them for historical operation audits and backtracks.

Prerequisites

CTS has been enabled.

Key EVS Operations Recorded by CTS

Table 20-1 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
Create disk tag	evs	createVolumeTag

Viewing Traces

To view audit logs, see [Querying Real-Time Traces](#).

21 Managing Quotas

21.1 Querying EVS Resource Quotas

Scenarios

Quotas are enforced for service resources on the platform to prevent unforeseen spikes in resource usage. Quotas can limit the number or amount of resources available to users, such as the number of EVS disks, the capacity of EVS disks, and the number of EVS snapshots.

Users can perform the following operations to view the resource quota details.

Procedure


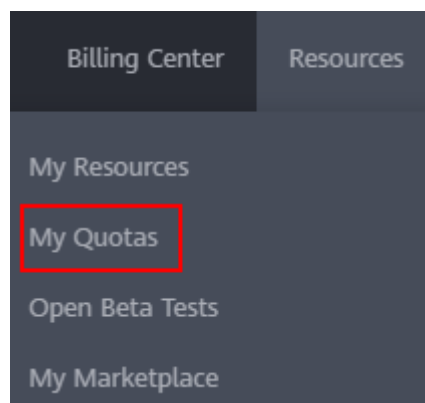
- Step 1** Log in to the [management console](#).
- Step 2** Click  in the upper left corner and select the desired region and project.
- Step 3** In the upper right corner of the page, choose **Resources** > **My Quotas**.
The **Service Quota** page is displayed.

Figure 21-1 My Quotas



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

21.2 Increasing EVS Resource Quotas

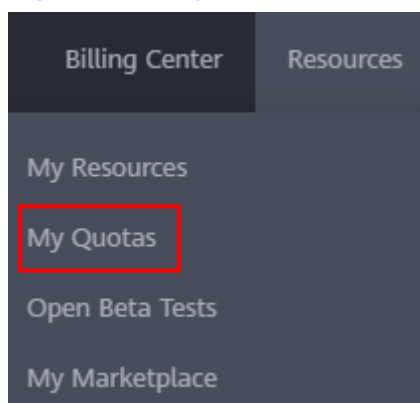
Scenarios

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.

Figure 21-2 My Quotas



3. Click **Increase Quota** in the upper right corner of the page.

Figure 21-3 Increasing quota

Service Quota ⓘ				Increase Quota
Service	Resource Type	Used Quota	Total Quota	
Auto Scaling	AS group	0		
	AS configuration	0		
Image Management Service	Image	0		
Cloud Container Engine	Cluster	0		
FunctionGraph	Function	0		
	Code storage(MB)	0		
	Disk	3		
Elastic Volume Service	Disk capacity(GB)	120		
	Snapshots	4		
Storage Disaster Recovery Service	Protection group	0		
	Replication pair	0		
Cloud Server Backup Service	Backup Capacity(GB)	0		
	Backup	0		
Scalable File Service	File system	0		
	File system capacity(GB)	0		
CDN	Domain name	0		
	File URL refreshing	0		
	Directory URL refreshing	0		
	URL preheating	0		

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**.

A Appendix

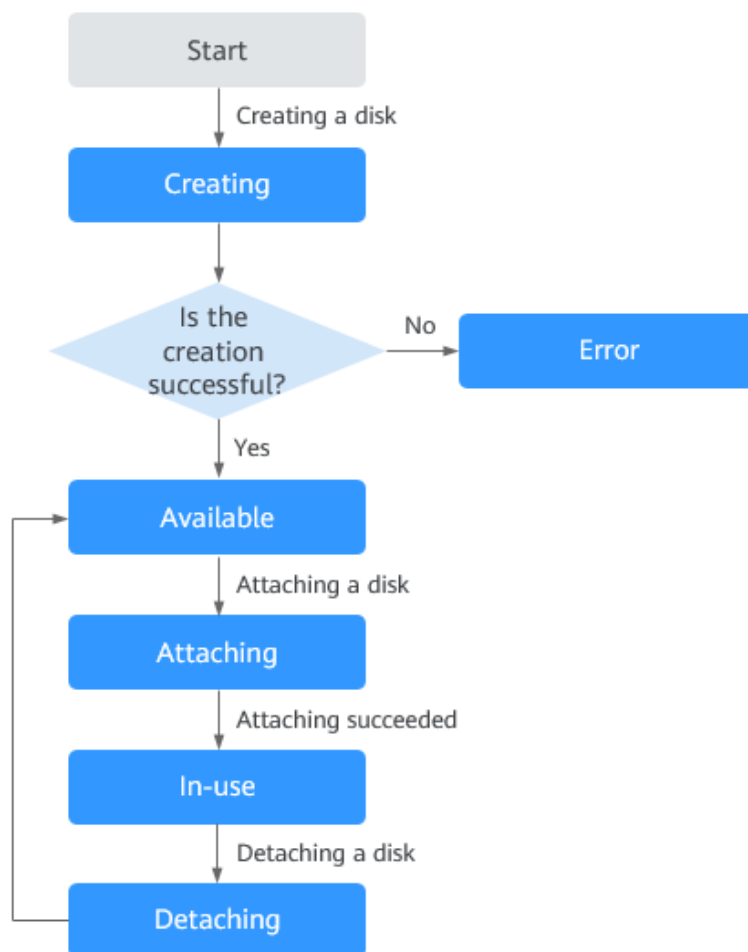
A.1 EVS Disk Status

An EVS disk has several statuses. [Table A-1](#) lists EVS disk statuses, the meaning of each status, and the operations a disk in each status allows.

Table A-1 Disk status details

EVS Disk Status	Status Description	Allowed Operation
In-use	The EVS disk is attached to a server and is in use.	<ul style="list-style-type: none">• Detaching• Creating backups• Expanding capacity
Available	The EVS disk has not been attached to any server and can be attached.	<ul style="list-style-type: none">• Attaching• Expanding capacity• Deleting• Creating backups• Rolling back snapshots to EVS disks
Creating	The EVS disk is being created.	None
Attaching	The EVS disk is being attached to a server.	None
Detaching	The EVS disk is being detached from a server.	None
Deleting	The EVS disk is being deleted.	None
Restoring	A backup is being used to restore the EVS disk.	None

EVS Disk Status	Status Description	Allowed Operation
Expanding	The capacity of the EVS disk is being expanded.	None
Uploading	Data on the EVS disk is being uploaded to an image. This status occurs when you create an image from a server.	None
Downloading	Data is being downloaded from an image to the EVS disk. This status occurs when you create a server.	None
Error	An error occurs when you try to create an EVS disk.	Deleting
Deletion failed	An error occurs when you try to delete the EVS disk.	None
Expansion failed	An error occurs when you try to expand the capacity of the EVS disk.	Deleting
Restoration failed	An error occurs when you try to restore the EVS disk from a backup.	Deleting
Rolling back	Data on the EVS disk is being restored from a snapshot. NOTE <ul style="list-style-type: none">When you roll back a snapshot, you can only restore the data to the original disk. Data restoration to a specific disk is not possible.A snapshot can only be rolled back when the original disk is in the available or error_rollbacking state.	None
Rollback failed	An error occurs when the EVS disk is rolled back from a snapshot.	<ul style="list-style-type: none">DeletingRolling back snapshots to EVS disks
Awaiting transfer	The EVS disk is awaiting for a transfer.	None

Figure A-1 Change of some EVS disk statuses**NOTE**

If an EVS disk status is **Error**, **Deletion failed**, **Expansion failed**, **Restoration failed**, or **Rollback failed**, you can rectify the error by following the steps provided in "What Can I Do If an Error Occurs on My EVS Disk" in FAQs.

A.2 EVS Snapshot Status

An EVS snapshot has several statuses. [Table A-2](#) lists the EVS snapshot statuses, the meaning of each status, and the operations a snapshot in each status allows.

Table A-2 Snapshot status details

Snapshot Status	Description	Allowed Operation
Creating	The snapshot is being created.	No operations are allowed.

Snapshot Status	Description	Allowed Operation
Available	The snapshot is successfully created.	<ul style="list-style-type: none"> Creating EVS disks using snapshots Deleting snapshots Rolling back snapshots to source EVS disks
Deleting	The snapshot is being deleted.	No operations are allowed.
Error	An error occurs when you try to create a snapshot.	Deleting
Deletion failed	An error occurs when you try to delete a snapshot.	No operations are allowed.
Rolling back	<p>The snapshot is rolling back data.</p> <p>NOTE</p> <ul style="list-style-type: none"> When you roll back a snapshot to an EVS disk, you can only roll back the snapshot to the source EVS disk. Rollback to a specified disk is not supported. You can roll back an EVS disk from a snapshot only when the disk is in the Available or Rollback failed state. 	No operations are allowed.
Backing up	<p>This status is available only to temporary snapshots. When you create a backup for an EVS disk, a temporary snapshot is automatically created. This status indicates that a temporary snapshot is being created during the backup creation.</p> <p>NOTE</p> <p>Temporary snapshots are created through the CBR service. Do not perform any operation on these snapshots.</p>	No operations are allowed.

B Change History

Released On	Description
2023-11-14	This issue is the thirty-first official release, which incorporates the following change: Modified the supported changes between disk types in section Changing the EVS Disk Type (OBT) .
2023-11-01	This issue is the sixteenth official release, which incorporates the following change: Updated or added certain permission restrictions in sections Managing Encrypted EVS Disks and Managing Shared EVS Disks .
2023-07-27	This issue is the fifteenth official release, which incorporates the following change: Added section Unsubscribing from Yearly/Monthly EVS Disks .
2023-07-20	This issue is the fourteenth official release, which incorporates the following change: Updated and added constraints in sections Detaching a System Disk , Detaching a Data Disk , Managing Encrypted EVS Disks , Managing Shared EVS Disks , Changing EVS Disk Name , and Changing the EVS Disk Type (OBT) .
2023-06-15	This issue is the thirteenth official release, which incorporates the following change: Added support for the General Purpose SSD V2 disk type in section Changing the EVS Disk Type (OBT) .
2023-05-09	This issue is the twelfth official release, which incorporates the following change: Added a constraint that the disk type cannot be changed when any snapshot of the disk is being deleted in Changing the EVS Disk Type (OBT) .

Released On	Description
2023-01-19	This issue is the eleventh official release, which incorporates the following change: Added a note about the language environment check in Extending an Existing MBR or GPT Partition .
2022-08-30	This issue is the tenth official release, which incorporates the following change: <ul style="list-style-type: none">Added section Changing EVS Disk Type.
2021-01-25	This issue is the ninth official release, which incorporates the following change: <ul style="list-style-type: none">Added Extending Disk Partitions and File Systems (Windows Server 2016).
2019-06-30	This issue is the eighth official release, which incorporates the following change: Added support for the TMS tag function.
2018-10-30	This issue is the seventh official release, which incorporates the following changes: <ul style="list-style-type: none">Added Changing Disk Name.Modified the parameter list for purchasing disks according to the modifications made on the console.
2018-08-15	This issue is the sixth official release, which incorporates the following changes: <ul style="list-style-type: none">Added support for the EVS snapshot function.Added support for expanding In-use EVS disks.
2018-08-03	This issue is the fifth official release, which incorporates the following changes: <ul style="list-style-type: none">Added precautions for using shared EVS disks together with SCSI.Modified disk performance metrics.
2018-06-30	This issue is the fourth official release, which incorporates the following change: <ul style="list-style-type: none">Added Managing EVS Disk Transfer.
2018-01-22	This issue is the third official release, which incorporates the following changes: <ul style="list-style-type: none">Added support for SCSI EVS disks.Added support for shared EVS disks.
2018-01-15	This issue is the second official release, which incorporates the following change: <ul style="list-style-type: none">Made modifications to the document according to changes made on the console.

Released On	Description
2017-12-31	This issue is the first official release.