

Elastic Volume Service(EVS)

User Guide

Issue 01
Date 2024-11-26



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
2.1 Overview.....	5
2.2 Purchasing an EVS Disk.....	6
2.3 Attaching an EVS Disk.....	16
2.3.1 Attaching a Non-Shared Disk.....	16
2.3.2 Attaching a Shared Disk.....	19
2.4 Initializing EVS Data Disks.....	22
2.4.1 Initialization Overview.....	22
2.4.2 Initializing a Linux Data Disk (Less Than or Equal to 2 TiB).....	25
2.4.3 Initializing a Linux Data Disk (Greater Than 2 TiB).....	31
2.4.4 Initializing a Windows Data Disk.....	36
3 Viewing EVS Disk Details.....	45
4 Changing the EVS Disk Type (OBT).....	49
5 Expanding EVS Disk Capacity.....	52
5.1 Expansion Overview.....	52
5.2 Step 1: Expand Disk Capacity.....	53
5.3 Step 2: Extend Disk Partitions and File Systems.....	56
5.3.1 Extending Disk Partitions and File Systems (Linux).....	56
5.3.2 Extending Disk Partitions and File Systems (Windows).....	74
6 Detaching and Deleting an EVS Disk.....	92
6.1 Detaching an EVS Disk.....	92
6.2 Unsubscribing from or Deleting an EVS Disk.....	95
6.3 Unsubscribing from Yearly/Monthly EVS Disks.....	98
7 Managing EVS Recycle Bin.....	101
7.1 Recycle Bin Overview.....	101
7.2 Enabling the Recycle Bin.....	103
7.3 Configuring a Recycle Bin Policy.....	104
7.4 Recovering Disks from the Recycle Bin.....	105

7.5 Permanently Deleting Disks from the Recycle Bin.....	106
7.6 Disabling the Recycle Bin.....	107
8 Managing EVS Snapshots.....	108
8.1 EVS Snapshot Overview.....	108
8.2 Using EVS Snapshots.....	116
8.2.1 Creating an EVS Snapshot.....	117
8.2.2 Rolling Back Disk Data from a Snapshot.....	125
8.2.3 Creating a Disk from a Snapshot.....	125
8.2.4 Enabling or Disabling Instant Snapshot Restore (for Snapshots in Commercial Use).....	127
8.2.5 Checking the EVS Snapshot Storage Usage (for Snapshots in Commercial Use).....	128
8.2.6 Checking EVS Snapshot Details.....	129
8.2.7 Deleting an EVS Snapshot.....	131
9 Managing Encrypted EVS Disks.....	133
10 Managing Shared EVS Disks.....	139
11 Managing EVS Disk Backups.....	144
11.1 CBR Overview.....	144
11.2 Backing Up EVS Disks.....	146
12 Managing EVS Transfers.....	149
13 Managing EVS Tags.....	152
13.1 Tag Overview.....	152
13.2 Adding a Tag.....	152
13.3 Modifying a Tag.....	154
13.4 Deleting a Tag.....	155
13.5 Searching for Disks by Tag.....	155
14 Managing EVS Quotas.....	157
14.1 Querying EVS Resource Quotas.....	157
14.2 Increasing EVS Resource Quotas.....	158
15 Cloud Eye Monitoring.....	160
15.1 Viewing Basic EVS Monitoring Data.....	160
15.2 Viewing EVS Monitoring Data Included in OS Metrics (with Agent Installed).....	163
16 Recording EVS Operations Using CTS.....	167

1 Permissions Management

1.1 Creating a User and Granting EVS Permissions

You can use [IAM](#) for fine-grained permissions control for your EVS resources. With IAM, you can:

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

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

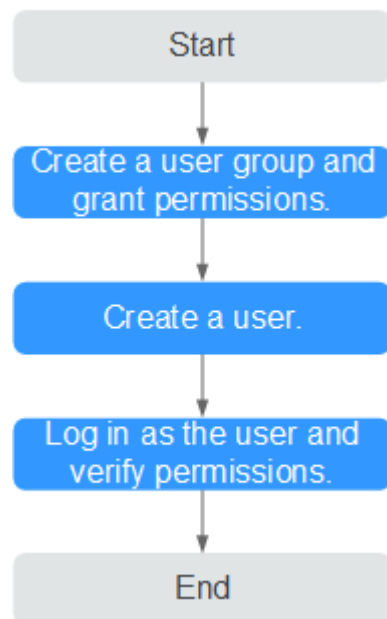
This section describes the procedure for granting permissions (see [Figure 1-1](#)).

Prerequisites

Before granting permissions to user groups, learn about system-defined permissions in [EVS Permissions](#) for EVS.

Process Flow

Figure 1-1 Process for granting EVS permissions



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

1.2 EVS Custom Policies

You can create custom policies to supplement the system-defined policies of EVS. For the actions supported for custom policies, see [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 3: Grant permissions to forcibly create encrypted disks.

You can create a custom policy to force users to create only encrypted disks.

```
{
  "Version": "5.0",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": [
        "evs:volumes:create"
      ],
      "Condition": {
        "Bool": {
          "evs:Encrypted": [
            "false"
          ]
        }
      }
    }
  ]
}
```



```
}  
  }  
} ]  
}
```

- Example 4: Grant permissions to forcibly create backups for disks.
You can create a custom policy to force users to use cloud backup when creating disks.

 **NOTE**

When forcible backup is configured and you are creating a yearly/monthly disk, you must choose an existing backup vault.

Example policy:

```
{  
  "Version": "5.0",  
  "Statement": [  
    {  
      "Effect": "Deny",  
      "Action": [  
        "evs:volumes:create"  
      ],  
      "Condition": {  
        "Null": {  
          "cbr:VaultId": [  
            "true"  
          ]  
        }  
      }  
    }  
  ]  
}
```

2 Purchasing and Using an EVS Disk

2.1 Overview

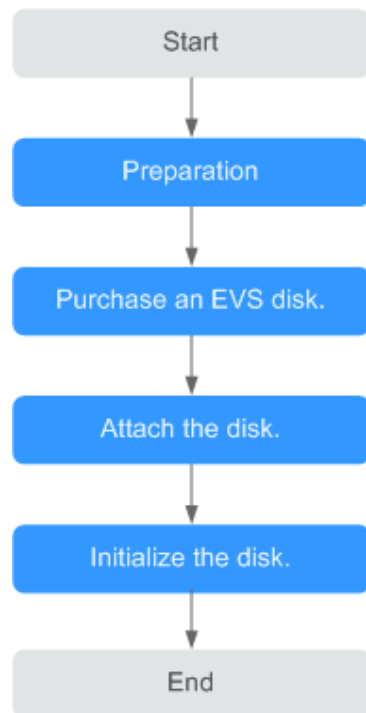
EVS disks can be attached to servers to be used as system disks or data disks. For details, see [Table 2-1](#).

Table 2-1 Method for purchasing disks

Function	Description	Method
System disk	System disks are purchased together with servers. You cannot purchase them separately.	<ul style="list-style-type: none">• Purchasing an ECS• Creating a BMS
Data disk	You can purchase data disks together with servers or separately.	<ul style="list-style-type: none">• Purchasing an ECS• Creating a BMS• Purchasing an EVS Disk

[Figure 2-1](#) shows the process of purchasing and using a data disk.

Figure 2-1 Process overview



1. **Make preparations:** [Sign up for a HUAWEI ID](#), [enable Huawei Cloud services](#), and [top up your Huawei Cloud account](#).
2. **Buy an EVS disk:** Configure the disk parameters, including the disk type, size, name, and other information by referring to [Purchasing an EVS Disk](#).
3. **Attach the data disk:** Attach the separately purchased disk to an ECS by referring to [Attaching an EVS Disk](#).
4. **Initialize the data disk:** After the data disk is attached, log in to the ECS and initialize the disk before using it. For details about how to initialize the disk, see the following sections:
 - [Initialization Overview](#)
 - [Initializing a Linux Data Disk \(Less Than or Equal to 2 TiB\)](#)
 - [Initializing a Linux Data Disk \(Less Than or Equal to 2 TiB\)](#)
 - [Initializing a Linux Data Disk \(Less Than or Equal to 2 TiB\)](#)

2.2 Purchasing an EVS Disk

Scenarios

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

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


Notes and Constraints

Table 2-2 Constraints on purchasing disks

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

Procedure

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

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

Step 3 In the upper right corner, click **Buy Disk**.

Step 4 Configure disk parameters according to [Table 2-3](#).

Table 2-3 Disk parameters

Parameter	Sub-Parameter	Description	Example Value
Region	-	Mandatory Resources are region-specific and cannot be used across regions through internal network connections. For low network latency and quick resource access, select the nearest region.	-

Parameter	Sub-Parameter	Description	Example Value
AZ	-	<p>Mandatory</p> <p>The availability zone (AZ) where you want to create the disk.</p> <p>NOTE</p> <ul style="list-style-type: none">• Disks can only be attached to the servers in the same AZ.• The AZ of a disk cannot be changed after the disk has been created.	AZ1
Attach to Server	-	<p>Optional</p> <ul style="list-style-type: none">• Now: If you select this option, you need to select a server to attach the disk. The billing mode of the disk will be the same as the selected server.• Later: When no server is available, you can select this option to create the disk first and attach the disk after the purchase. <p>NOTE</p> <p>This parameter is available only in some regions. Whether it is displayed depends on the region where you use EVS.</p>	-

Parameter	Sub-Parameter	Description	Example Value
Billing Mode	-	<p>Mandatory</p> <p>You can pay for EVS disks in two ways:</p> <ul style="list-style-type: none"> ● Yearly/Monthly ● Pay-per-use <p>NOTICE</p> <ul style="list-style-type: none"> ● Selecting Now for Attach To Server: <ul style="list-style-type: none"> - If you select a yearly/monthly server, only yearly/monthly billing is available for the disk. If you want to buy a pay-per-use disk for the yearly/monthly server, select Later for Attach To Server, buy a pay-per-use disk, and attach it to the yearly/monthly server after the purchase. - If you select a pay-per-use server, only pay-per-use billing is available for the disk. If you want to buy a yearly/monthly disk for the pay-per-use server, select Later for Attach To Server, buy a yearly/monthly disk, and attach it to the pay-per-use server after the purchase. ● Selecting Later for Attach To Server: If you buy a yearly/monthly disk and attach it after the purchase, the disk cannot be renewed or unsubscribed from together with its server and may have an expiration time different from the server. 	Pay-per-use

Parameter	Sub-Parameter	Description	Example Value
Data Source (Optional)	Create from <ul style="list-style-type: none"> • Backup • Snapshot • Image 	Optional <ul style="list-style-type: none"> • Create from Backup: The backup data is used to create the disk. Click Create from and choose Backup. On the displayed page, select the target backup and click OK. NOTE <ul style="list-style-type: none"> - One backup cannot be used for concurrent disk creation operations at the same time. For example, if you are creating disk A from a backup, this backup can be used to create another disk only after disk A has been created. - If a disk is created from a backup of a system disk, the new disk can be used as a data disk only. • Create from Snapshot: The snapshot data is used to create the disk. Click Create from and choose Snapshot. On the displayed page, select the target snapshot and click OK. NOTE <p>For details about how to create disks from snapshots, see Creating a Disk from a Snapshot.</p> • Create from Image: The image data is used to create the disk. Click Create from and choose Image. On the displayed page, select the target image and click OK. NOTE <ul style="list-style-type: none"> - The device type of the new disk is the same as that of the image's source disk. - The encryption attribute of the new disk is the same as that of the image's source disk. 	<ul style="list-style-type: none"> • Create from Backup: autobackup-001

Parameter	Sub-Parameter	Description	Example Value
Disk Specifications	Disk Type	<p>Mandatory</p> <p>EVS disk types vary depending on regions. See the EVS types displayed on the console.</p> <p>To learn more about disk types, see Disk Types and Performance.</p> <p>NOTE General Purpose SSD V2 disks allow you to specify the disk IOPS and throughput. You can change the disk type after a disk is created. For details, see Changing the EVS Disk Type (OBT).</p>	Ultra-high I/O
Disk Specifications	Capacity (GiB)	<p>Mandatory</p> <p>The disk size. Only data disks can be created on the current page, and the disk size ranges from 10 GiB to 32,768 GiB.</p> <p>NOTE</p> <ul style="list-style-type: none"> When you use a backup to create a disk, the disk capacity must be greater than or equal to the backup size. In the condition that you do not specify a disk capacity, if the backup size is smaller than 10 GiB, the default capacity 10 GiB will be used as the disk capacity; if the backup size is greater than 10 GiB, the disk capacity will be consistent with the backup size. When you use a snapshot to create a disk, the disk capacity must be greater than or equal to the snapshot size. In the condition that you do not specify a disk capacity, if the snapshot size is smaller than 10 GiB, the default 10 GiB will be used as the disk capacity; if the snapshot size is greater than 10 GiB, the disk capacity will be consistent with the snapshot size. The system shows you the maximum number of disks as well as the maximum disk capacity allowed to purchase. To ensure effective resource usage, if the disk capacity you need exceeds the upper limit, click Increase Quota to obtain a higher quota. You can purchase the disk capacity you need after the request is approved. 	100 GiB

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

Parameter	Sub-Parameter	Description	Example Value
More	Advanced Settings <ul style="list-style-type: none"> • Share • SCSI • Encryption 	<p>Optional</p> <ul style="list-style-type: none"> • Share <p>If you select Share, a shared disk is created. A shared disk can be attached to up to 16 servers. If you do not select Share, a non-shared disk is created, and the disk can be attached to one a server only.</p> <p>If you select both SCSI and Share, a shared SCSI disk is created.</p> <p>NOTE The sharing attribute of a disk cannot be changed after the disk has been created.</p> • SCSI <p>If you select SCSI, a SCSI disk is created. Such disks allow the server OS to directly access the underlying storage media and send SCSI commands to the disks. If you do not select SCSI, a VBD disk is created. That said, the disk device type is VBD, the default device type.</p> <p>NOTE The device type of a disk cannot be changed after the disk has been created.</p> • Encryption <p>This option is only used to encrypt data disks, and you need to create an agency to grant KMS access rights to EVS.</p> <p>After the access rights are granted, configure the following parameters on the Encryption Settings page displayed:</p> <ul style="list-style-type: none"> - Select an existing key <p>If you select Select an existing key, select a key from the drop-down list. You can select one of the following keys:</p> <p>Default Key: After the KMS access rights have been granted to EVS, the system automatically creates a Default Key evs/default.</p> 	-

Parameter	Sub-Parameter	Description	Example Value
		<p>An existing or new custom key. For details about how to create a key, see Creating a Key.</p> <ul style="list-style-type: none"> - Enter a key ID <p>If you select Enter a key ID, enter an ID of a key shared with you by another account. Ensure that the shared key is in the same region where you want to create the disk. For details, see Creating a Grant.</p> <p>NOTE</p> <ul style="list-style-type: none"> - System disk encryption relies on the image. For details, see Creating Encrypted Images. - Before using the encryption function, you need to create an agency to grant KMS access rights to EVS. If you have the right to grant the permission, grant the KMS access rights to EVS directly. After the KMS access rights have been granted, follow-up operations do not require the rights to be granted again. If you do not have this permission, contact a user with the security administrator permissions to grant KMS access rights to EVS, then repeat the preceding operations. - The encryption attribute of a disk cannot be changed after the disk has been created. 	
More	Tag	<p>Optional</p> <p>You can add tags when creating disks. Tags can help you identify, classify, and search for your disks. For details about tag rules, see Adding a Tag.</p> <p>NOTE</p> <ul style="list-style-type: none"> • Tag rules vary depending on regions. See the rules displayed on the console. • Except for tagging the disk during disk creation, you can also add, modify, or delete tags for existing disks. For more information about tags, see Managing EVS Tags. 	-

Parameter	Sub-Parameter	Description	Example Value
Disk Name	-	<p>Mandatory</p> <ul style="list-style-type: none"> If you create a single disk, the name you entered will be used as the disk name. The name can contain a maximum of 64 characters. If you create multiple disks in a batch, the name you entered will be used as the prefix of disk names. An actual disk name will be composed of the name you entered and a four-digit number. The name can contain a maximum of 59 characters. 	<p>For example, if you create two disks and set volume for Disk Name, the EVS disk names will be volume-0001 and volume-0002.</p>
Quantity	-	<ul style="list-style-type: none"> Usage Duration: This parameter is mandatory if you select Yearly/Monthly for Billing Mode. You can choose from 1 month to 3 years for the usage duration. Quantity: This parameter is optional. The preset disk quantity is 1, which means one disk will be created. You can create a maximum of 100 disks at a time. <p>NOTE</p> <ul style="list-style-type: none"> If the disk is created from a backup, batch creation is not possible, and this parameter must be set to 1. If the disk is created from a snapshot, batch creation is not possible, and this parameter must be set to 1. The system shows you the maximum number of disks as well as the maximum disk capacity allowed to purchase. To ensure effective resource usage, if the number of disks you need exceeds the upper limit, click Increase Quota to obtain a higher quota. You can purchase the disks you need after the request is approved. 	<p>Disk validity period: 1 year Disk quantity: 1</p>

Step 5 Click **Next**.

- If you select **Yearly/Monthly** for **Billing Mode**:
 - Check the disk details on the **Confirm** page.
 - Confirm the information and click **Submit**.

- c. On the **Pay** page, select a desired payment method and confirm the payment. The system displays a message indicating payment processed successfully.
 - d. Click **Back to Elastic Volume Service** to return to the **Elastic Volume Service** page.
- If you select **Pay-per-use** for **Billing Mode**:
 - a. Check the disk details on the **Confirm** page.
 - b. Confirm the information and click **Submit**. The system displays a message indicating request submitted successfully.
 - c. Click **Back to Disk List** to return to the **Elastic Volume Service** page.

Step 6 Click **Back to Disk List**.

The disk list page is displayed.

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

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

----End

2.3 Attaching an EVS Disk

2.3.1 Attaching a Non-Shared Disk

Scenarios

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

- Separately created data disks
- Detached data disks
- Detached system disks

 **NOTE**

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

Prerequisites

- The non-shared disk status is **Available**.
- To attach a data disk, the status of the server must be **Running** or **Stopped**.
- To attach a system disk, the status of the server must be **Stopped**.
- The account is not in arrears.


Notes and Constraints

- Cloud servers created from ISO images are only used for OS installation. They have limited functions and cannot have EVS disks attached.

- A non-shared disk can only be attached to one server.
- The disk and the server must be in the same region and AZ.
- A detached, non-shared yearly/monthly data disk purchased together with a server can only be re-attached to the original server to 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 system disk purchased together with a yearly/monthly 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 system disk purchased together with a pay-per-use 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.

Attaching the Disk on the EVS Console

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

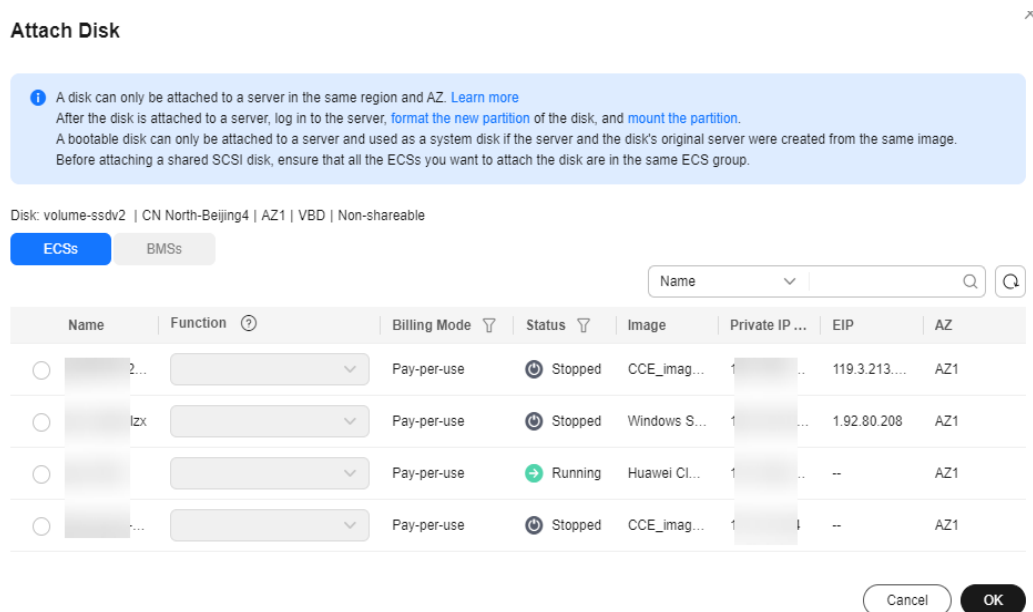
Step 2 Click  in the upper left corner and choose **Storage > Elastic Volume Service**.
The **Elastic Volume Service** page is displayed.

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

Step 4 Select the server and then select the disk function from the drop-down list. Ensure that the disk and server are in the same AZ.

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

Figure 2-2 Attach Disk



Step 5 Click **OK**.

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

NOTICE

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

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

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

----End

Attaching the Disk on the ECS Console



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

Figure 2-3 Attach Disk (KVM)



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

 **NOTE**

- If no disks are available, click **Create Disk** in the lower part of the list.
 - For the restrictions on attaching disks, see [What Are the Requirements for Attaching an EVS Disk to an ECS?](#)
8. Click **OK**.
After the disk is attached, you can view information about it on the **Disks** tab.

Follow-Up Operations

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

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

```
mount Disk partition Mount point
```

Helpful Links

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

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

2.3.2 Attaching a Shared Disk

Scenarios

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

- Separately created data disks
- Detached data disks

Prerequisites

- The shared disk status is **In-use** or **Available**.
- The statuses of servers are **Running** or **Stopped**.
- The account is not in arrears.

Notes and Constraints

NOTICE

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


- A shared disk can be attached to a maximum of 16 servers. These servers and the shared disk must be in the same AZ of a region.
- A shared, **In-use** disk can only be attached to servers when the maximum number of servers that the disk can be attached to has not been reached.
- A shared disk can only be attached to servers running the same type of OS (either Windows or Linux).

For example, if you attach a shared disk to multiple Windows servers and then detach it, the shared disk cannot be attached to Linux servers later. This is because Windows and Linux support different file systems. Improper operations may damage the original file system.

- A shared disk can only be used as a data disk. It cannot be used as a system disk.
- Cloud servers created from ISO images are only used for OS installation. They have limited functions and cannot have EVS disks attached.
- A frozen disk cannot be attached.
- A detached system disk purchased together with a yearly/monthly 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 system disk purchased together with a pay-per-use 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.

Attaching the Disk on the EVS Console

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

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

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

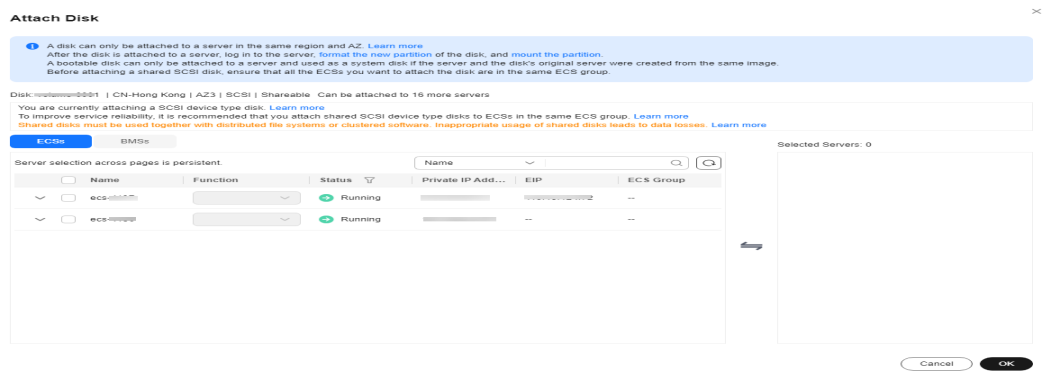
Step 3 In the disk list, locate the disk and click **Attach**.

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

Step 4 Select the target servers to attach the shared disk. Ensure that the disk and servers are in the same AZ. After you select servers, **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 show up in the drop-down list and cannot be selected.

Figure 2-4 Attach Disk



Step 5 Click **OK**.


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

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

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

----End

Attaching the Disk on the ECS Console

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

 **NOTE**

- If no disks are available, click **Create Disk** in the lower part of the list.
 - For the restrictions on attaching disks, see [What Are the Requirements for Attaching an EVS Disk to an ECS?](#)
8. Click **OK**.
- After the disk is attached, you can view information about it on the **Disks** tab.

Follow-Up Operations

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

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

```
mount Disk partition Mount point
```

Helpful Links

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

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

2.4 Initializing EVS Data Disks

2.4.1 Initialization Overview

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

Scenarios

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

For detailed operation instructions, see [Table 2-4](#).

- **Existing data disk**

An existing data disk is a disk created from a snapshot, a backup, or an image, or a disk detached from another server.

- You can choose not to initialize the disk and use the disk existing partitions.
 - In Linux, **mount the partitions on desired mount points** and **configure auto mount at system start**.
 - In Windows, no further action is required. You can simply use the existing partitions.
- You can also re-initialize the data disk.

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

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

For detailed initialization operations, see [Table 2-4](#).

 NOTE

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

Operation Instructions

Table 2-4 Disk initialization instructions

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

NOTICE

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

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

Scenarios

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

Table 2-5 Initialization instructions

OS	Partition Style	File System Format	Partitioning Tool	Example Configuration
Not limited	<ul style="list-style-type: none">• GPT• MBR	ext* (such as ext2, ext3, and ext4), xfs, and btrfs	<ul style="list-style-type: none">• fdisk• parted	<ul style="list-style-type: none">• Partitioning tool: fdisk• Device name: /dev/vdb• File system format: ext4• Mount points: /mnt/sdc and /mnt/sdd• Partition 1: /dev/vdb1<ul style="list-style-type: none">- Size: 40 GiB- Partition style: MBR• Partition 2: /dev/vdb2<ul style="list-style-type: none">- Size: 60 GiB- Partition style: MBR

Prerequisites

You have attached the disk to a server.

Notes and Constraints

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

Initializing a Data Disk Manually

NOTE

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

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

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

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

Step 1 Log in to the server.

For how to log in to an ECS, see [Logging In to an ECS](#).

For how to log in to a BMS, see [Logging In to a BMS](#).

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

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

`lsblk`

```
[root@ecs-centos76 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda 253:0 0 40G 0 disk
├vda1 253:1 0 1G 0 part /boot
├vda2 253:2 0 39G 0 part /
└vdb 253:16 0 100G 0 disk
```

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

`fdisk /dev/vdb`

`n`

`p`

`1`

NOTE

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

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

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

```
Device does not contain a recognized partition table
Building a new DOS disklabel with disk identifier 0x38717fc1.
```

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

```
Partition type:
```

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

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

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

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

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

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

n

p

2

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

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

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

NOTE

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

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

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

First sector = 2048

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

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

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

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

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

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

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

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

1. Check whether the partitioning is successful.

p

```
Command (m for help): p
Disk /dev/vdb: 107.4 GB, 107374182400 bytes, 209715200 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x994727e5
```

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

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


 NOTE

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

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

w

partprobe

 NOTE

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

3. Confirm that the partition style is MBR.

parted /dev/vdb

p

 NOTE

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

```
[root@ecs-test-0001 ~]# parted /dev/vdb
GNU Parted 3.1
Using /dev/vdb
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) p
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 107GB
Sector size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:

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

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

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

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

mkfs -t ext4 /dev/vdb1

mkfs -t ext4 /dev/vdb2

 NOTE

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

```
[root@ecs-test-0001 ~]# mkfs -t ext4 /dev/vdb1
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
2621440 inodes, 10485504 blocks
524275 blocks (5.00%) reserved for the super user
```

```
First data block=0
Maximum filesystem blocks=2157969408
320 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000, 7962624

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

Check whether the file system format is ext4.

parted /dev/vdb

p

```
[root@ecs-test-0001 ~]# parted /dev/vdb
GNU Parted 3.1
Using /dev/vdb
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) p
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 107GB
Sector size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:

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

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

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

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

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

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

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

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

```
lsblk
```

View the mount results.

```
[root@ecs-test-0001 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda  253:0  0 40G 0 disk
├─vda1 253:1  0 40G 0 part /
vdb  253:16 0 100G 0 disk
├─vdb1 253:17 0 40G 0 part /mnt/sdc
└─vdb2 253:18 0 60G 0 part /mnt/sdd
```

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

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

 NOTE

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

1. Query the partition UUIDs.

```
blkid /dev/vdb1
```

```
blkid /dev/vdb2
```

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

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

2. Configure auto mount at startup.

```
vi /etc/fstab
```

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

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

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

Table 2-6 Parameter description

Example Value	Description
UUID=0b3040e2-1367-4abb-841d-ddb0b92693df	The UUID of the partition.
/mnt/sdc	The mount point of the partition.
ext4	The file system format of the partition.
defaults	The partition mount option. Normally, this parameter is set to defaults .
0	<ul style="list-style-type: none"> – The Linux dump backup option. <ul style="list-style-type: none"> ▪ 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.

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

Step 7 Verify that auto mount takes effect.

```
umount /dev/vdb1
```

```
umount /dev/vdb2
```

```
mount -a
```

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

Query file system mounting information.

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

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

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

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

----End

2.4.3 Initializing a Linux Data Disk (Greater Than 2 TiB)

Scenarios

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

Partition Style	OS	File System Format	Partitioning Tool	Example Configuration
GPT	Not limited	ext* (such as ext2, ext3, and ext4), xfs, and btrfs	parted	<ul style="list-style-type: none"> • Device name: /dev/vdb • File system format: ext4 • Mount point: /mnt/sdc • Partition name: /dev/vdb1 • Partition style: GPT • Size: 3 TiB

Prerequisites

You have attached the disk to a server.

Notes and Constraints

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

Initializing a Data Disk Greater Than 2 TiB

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

Step 1 Log in to the server.

For how to log in to an ECS, see [Logging In to an ECS](#).

For how to log in to a BMS, see [Logging In to a BMS](#).

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

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

lsblk

```
[root@ecs-centos76 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda 253:0 0 40G 0 disk
└─vda1 253:1 0 40G 0 part /
vdb 253:16 0 3T 0 disk
```

2. Create the /dev/vdb1 partition.

parted /dev/vdb

```
p
[root@ecs-centos74 ~]# parted /dev/vdb
GNU Parted 3.1
Using /dev/vdb
```

```
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) p
Error: /dev/vdb: unrecognised disk label
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 3299GB
Sector size (logical/physical): 512B/512B
Partition Table: unknown
Disk Flags:
(parted)
```

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

NOTE

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

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

mklabel gpt

unit s

p

```
(parted) mklabel gpt
(parted) unit s
(parted) p
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 6442450944s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number Start End Size File system Name Flags
(parted)
```

NOTE

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

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

mkpart /dev/vdb1 2048s 100%

p

NOTE

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

```
(parted) mkpart /dev/vdb1 2048s 100%
(parted) p
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 6442450944s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
```

```
Disk Flags:
```

Number	Start	End	Size	File system	Name	Flags
1	2048s	6442448895s	6442446848s		/dev/vdb1	

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

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

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

mkfs -t ext4 /dev/vdb1

 **NOTE**

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

```
[root@ecs-test-0001 ~]# mkfs -t ext4 /dev/vdb1
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
201326592 inodes, 805305856 blocks
40265292 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2952790016
24576 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000, 7962624, 11239424, 20480000, 23887872, 71663616, 78675968,
    102400000, 214990848, 512000000, 550731776, 644972544

Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
[root@ecs-test-0001 ~]#
```

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

```
[root@ecs-test-0001 ~]# parted /dev/vdb
GNU Parted 3.1
Using /dev/vdb
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) p
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 3299GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

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

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

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

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

```
mkdir -p /mnt/sdc
mount /dev/vdb1 /mnt/sdc
```

lsblk

```
[root@ecs-test-0001 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda 253:0 0 40G 0 disk
├─vda1 253:1 0 40G 0 part /
vdb 253:16 0 3T 0 disk
└─vdb1 253:17 0 3T 0 part /mnt/sdc
```

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

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

 **NOTE**

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

1. Query the partition UUID.

blkid /dev/vdb1

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

The UUID of the **/dev/vdb1** partition is **0b3040e2-1367-4abb-841d-ddb0b92693df**.

2. Configure auto mount at startup.

vi /etc/fstab

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

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

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

Table 2-7 Parameter description

Example Value	Description
UUID=0b3040e2-1367-4abb-841d-ddb0b92693df	The UUID of the partition.
/mnt/sdc	The mount point of the partition.
ext4	The file system format of the partition.
defaults	The partition mount option. Normally, this parameter is set to defaults .

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

Step 6 Verify that auto mount takes effect.

```
umount /dev/vdb1
```

```
mount -a
```

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

Query file system mounting information.

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

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

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

----End

2.4.4 Initializing a Windows Data Disk

Scenarios

This section uses the example configurations below to describe how to use Disk Management Tool or a script to initialize a Windows data disk. The initialization operations may vary depending on the server OS. Perform initialization operations based on your server OS.

Partition Style	Example Configuration
<ul style="list-style-type: none">• GPT• MBR	<ul style="list-style-type: none">• Version: Windows Server 2019 Standard (64-bit)• Disk name: Disk 1• Size: 100 GiB• After the initialization:<ul style="list-style-type: none">- Partition name: New volume (D:)- Partition style: GPT- File system format: NTFS

Prerequisites

You have attached the disk to a server.

Notes and Constraints

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

Initializing a Data Disk Manually

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

Step 1 Log in to the server.

For how to log in to an ECS, see [Logging In to an ECS](#).

For how to log in to a BMS, see [Logging In to a BMS](#).

----End

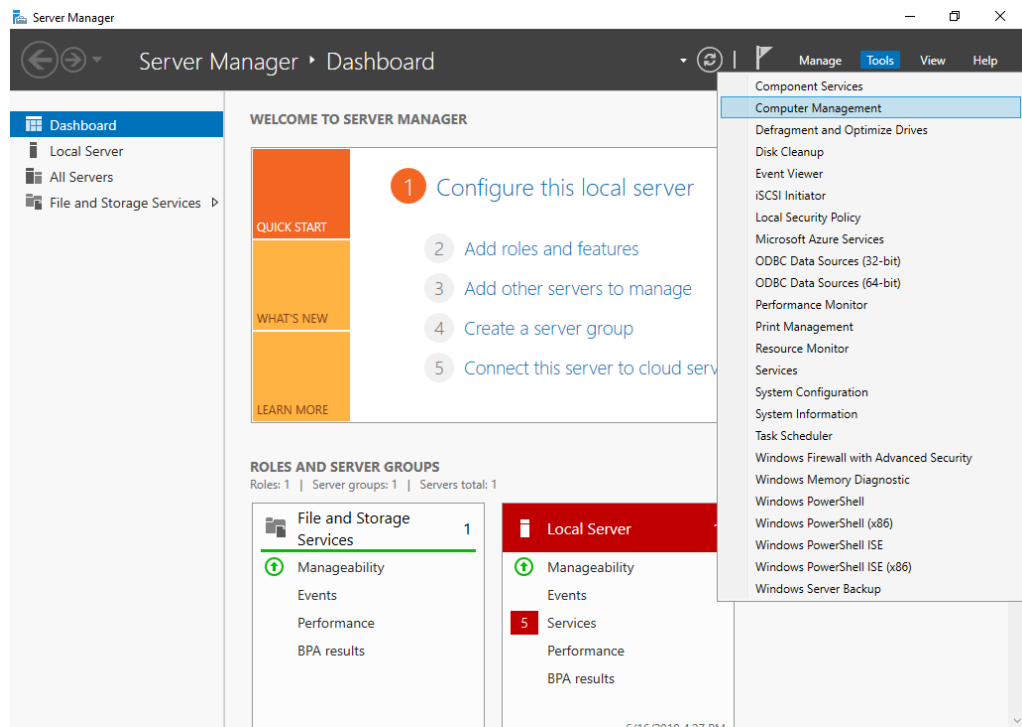
Step 1 On the desktop of the server, click the start icon in the lower left corner.

The **Windows Server** window is displayed.

Step 2 Click **Server Manager**.


The **Server Manager** window is displayed.

Figure 2-5 Server Manager



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

Figure 2-6 Computer Management



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

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

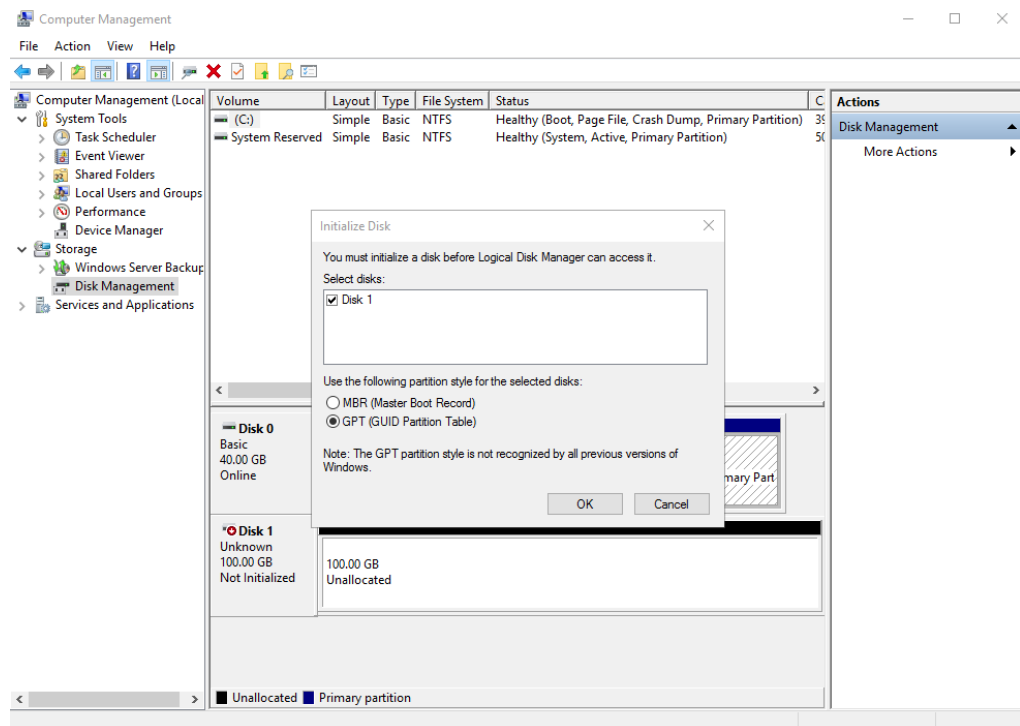
In the **Initialize Disk** dialog box, the to-be-initialized disk is selected. Select a partition style and click **OK**. In this example, **GPT (GUID Partition Table)** is selected.

NOTICE

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


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

Figure 2-7 Disk list




Step 5 In the **Unallocated** area of **Disk 1**, right-click the blank area and choose **New Simple Volume**.

Figure 2-8 Computer Management



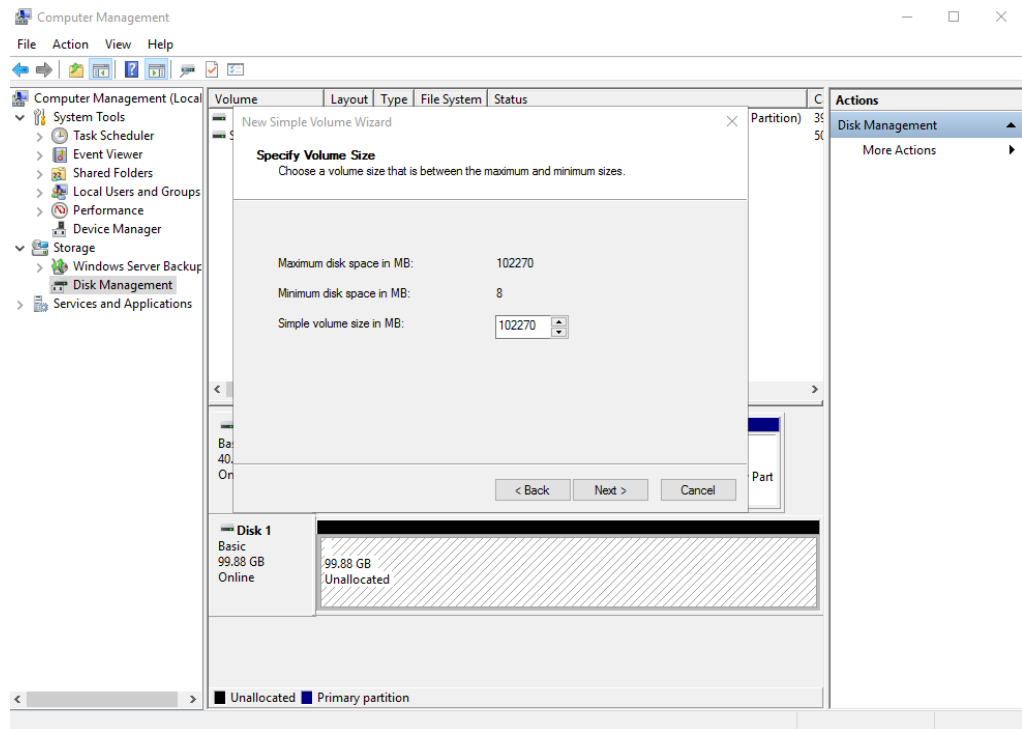
The **New Simple Volume Wizard** window is displayed.

Figure 2-9 New Simple Volume Wizard



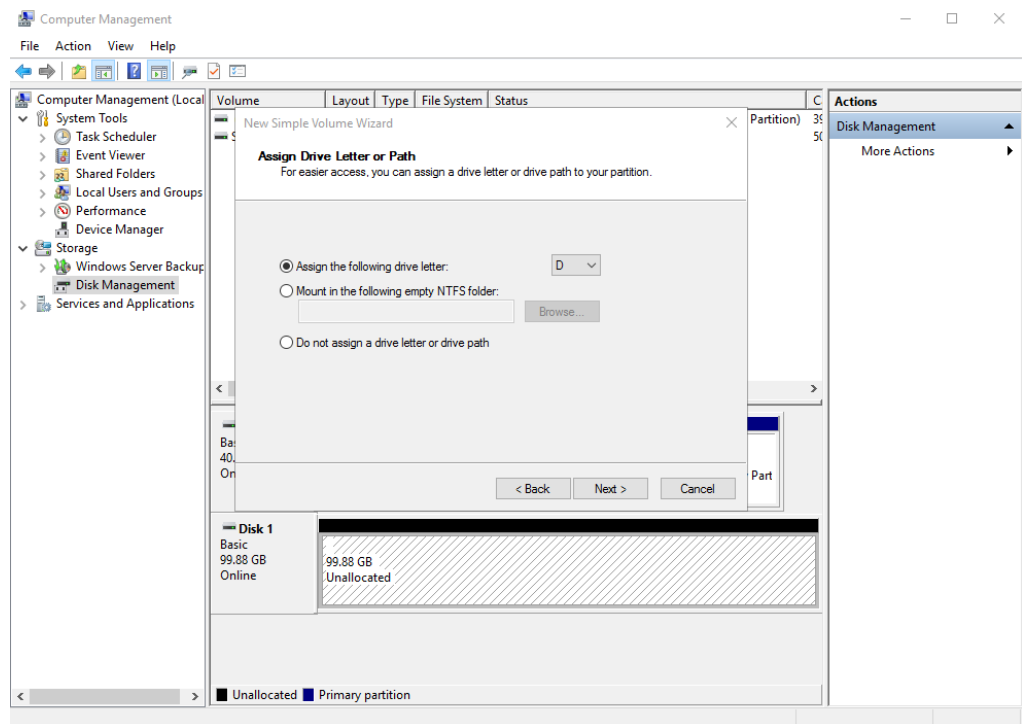
Step 6 Click **Next** to go to the **Specify Volume Size** page.

Figure 2-10 Specify Volume Size




Step 7 Specify the volume size and click **Next**. The system selects the maximum volume size by default. You can specify the volume size as required. In this example, the default setting is used.

Figure 2-11 Assign Drive Letter or Path




Step 8 Assign a drive letter or path to your partition and click **Next**. The system assigns drive letter D by default. In this example, the default setting is used.

Figure 2-12 Format Partition



Step 9 Specify format settings and click **Next**. The system selects the NTFS file system by default. You can specify a file system type as required. In this example, the default setting is used.

Figure 2-13 Completing the New Simple Volume Wizard



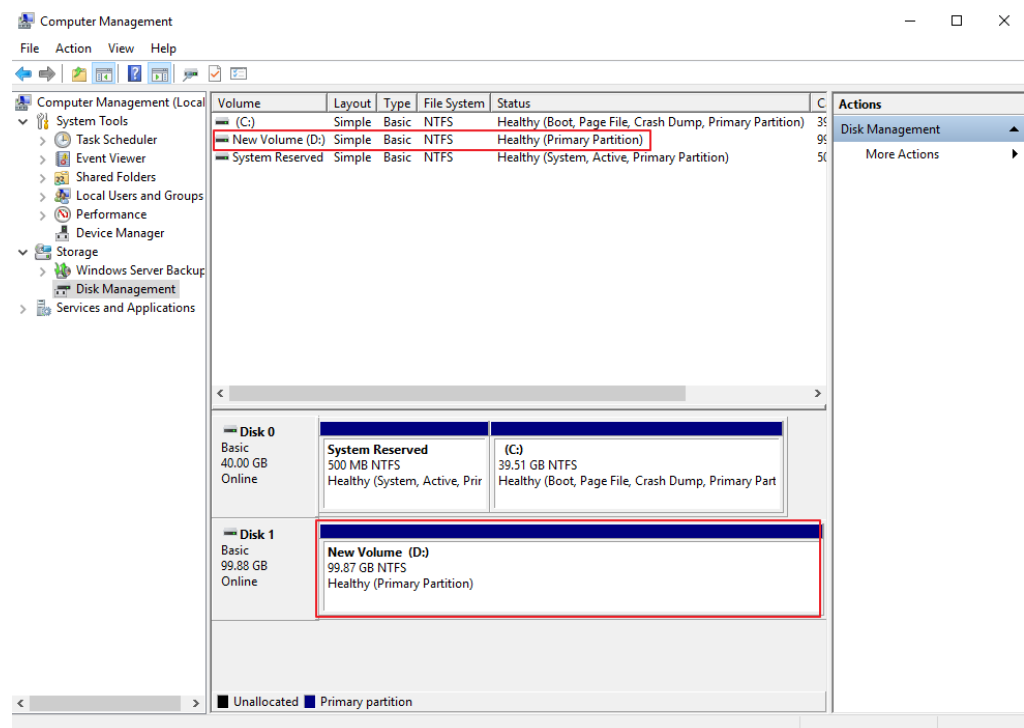
NOTICE


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

Step 10 Click **Finish**.

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

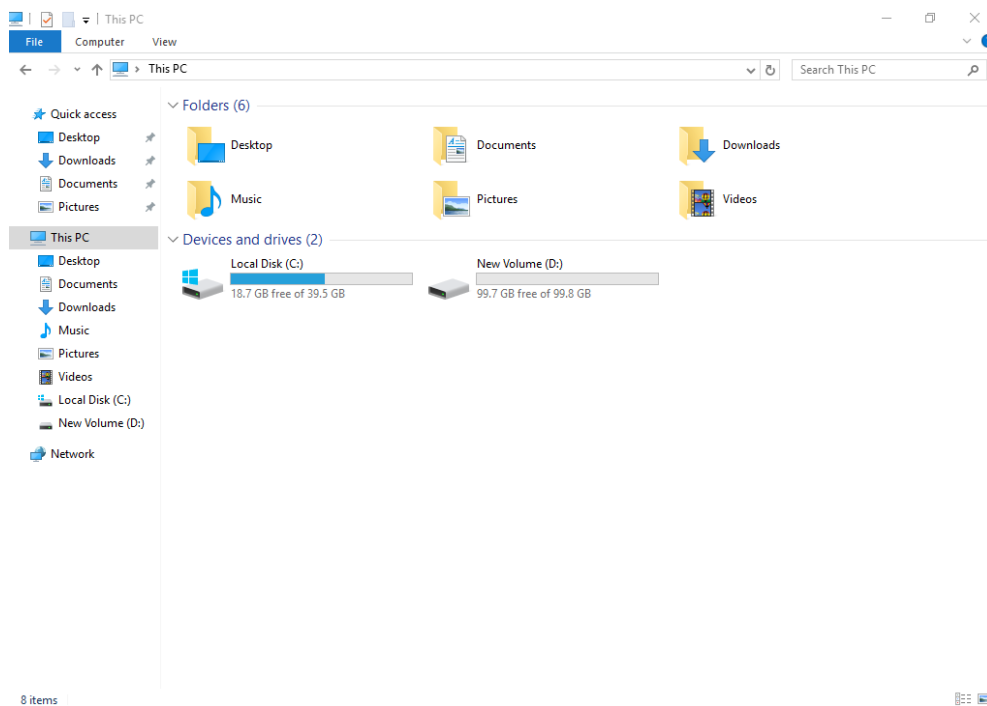
Figure 2-14 Disk initialized



Step 11 After the volume is created, click  on the task bar and check whether a new volume appears in the File Explorer. In this example, New Volume (D:) is the new volume.

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

Figure 2-15 File Explorer



-----End

3 Viewing EVS Disk Details

Scenarios

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

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


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

NOTE

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

Viewing Disk Details from the EVS Console

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

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

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


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

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

Viewing Disk Details from the Cloud Server Console

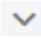
Step 1 Log in to the [console](#).

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

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

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

The server details page is displayed.

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

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

----End

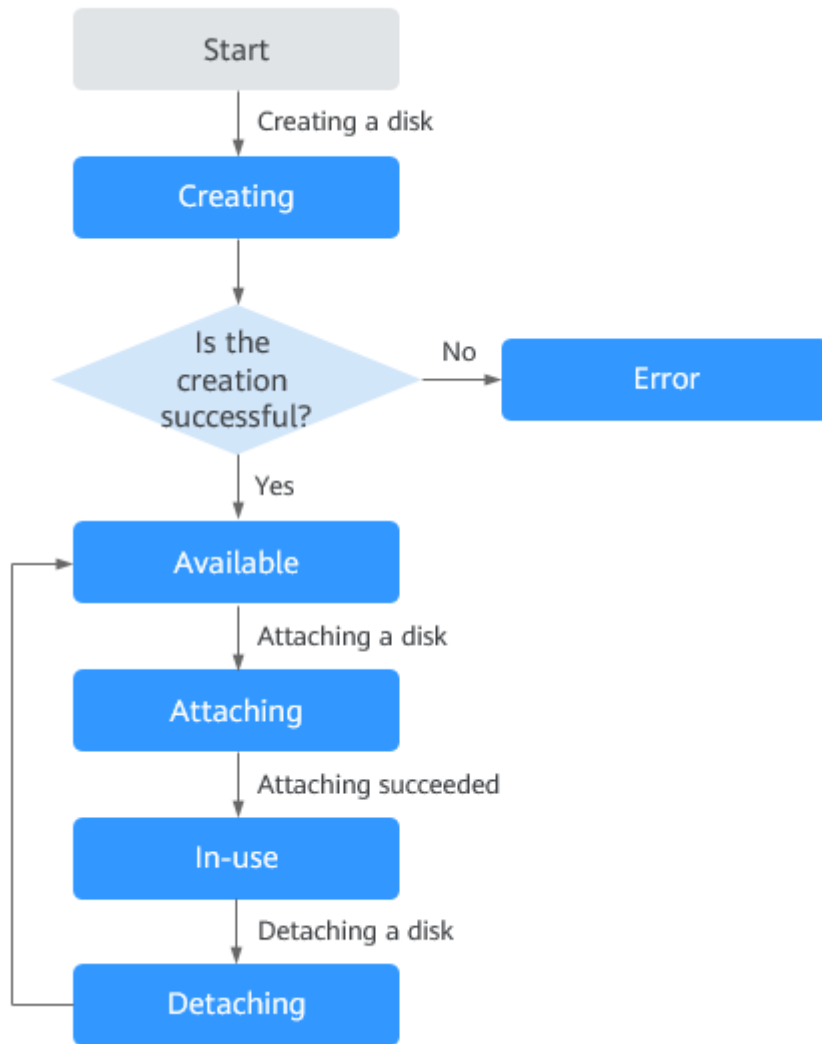
EVS Disk Status

Table 3-1 EVS disk status details

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

Status	Description	Allowed Operation
Uploading	Data on the EVS disk is being uploaded to an image. This status occurs when you create an image from a server.	-
Downloading	Data is being downloaded from an image to the EVS disk. This status occurs when you create a server.	-
Error	An error occurs when you try to create an EVS disk.	Deleting
Deletion failed	An error occurs when you try to delete the EVS disk.	None
Expansion failed	An error occurs when you try to expand the capacity of the EVS disk.	Deleting
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 data from a snapshot, you can only roll back data to the source EVS disk. Rollback to a specific disk is not supported. A snapshot can only be used for rollback when its source disk is in the Available or Rollback failed state. 	-
Rollback failed	An error occurs when the EVS disk data is rolled back from a snapshot.	<ul style="list-style-type: none"> Deleting Rolling back data to EVS disks using snapshots
Awaiting transfer	The EVS disk is awaiting for a transfer.	-

Figure 3-1 Change between some of 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.

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 OBT. [Submit a ticket](#) to apply for OBT.

Notes and Constraints

Table 4-1 Constraints on the disk type change

Phase	Description
Before the change	<ul style="list-style-type: none">You can only change the disk type 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.

The following table shows the supported changes between disk types.

 **NOTE**

Supported changes between disk types vary depending on regions. See the allowed changes on the console.

Table 4-2 Supported changes between disk types

Source Disk Type	New Disk Type
General Purpose SSD V2	General Purpose SSD V2 (IOPS or throughput, or both changed), Ultra-high I/O, General Purpose SSD, or Extreme SSD
Extreme SSD	General Purpose SSD V2 (IOPS or throughput, or both changed), Ultra-high I/O, or General Purpose SSD
Ultra-high I/O	General Purpose SSD V2 (IOPS or throughput, or both changed), Extreme SSD, or General Purpose SSD
General Purpose SSD	General Purpose SSD V2 (IOPS or throughput, or both changed), Extreme SSD, or Ultra-high I/O
High I/O	General Purpose SSD V2 (IOPS or throughput, or both changed), Extreme SSD, Ultra-high I/O, or General Purpose SSD


Source Disk Type	New Disk Type
Common I/O (previous generation product)	General Purpose SSD V2 (IOPS or throughput, or both changed), 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 [console](#).

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

Step 3 In the disk list, locate the target disk, 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.
To change to the General Purpose SSD V2 type, you also need to specify the disk IOPS and throughput.
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 Expanding EVS Disk Capacity

5.1 Expansion Overview

What Is EVS Capacity Expansion?

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

Expansion Upper Limits

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

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

 **NOTE**

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

How Do I Expand the Disk Capacity?

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

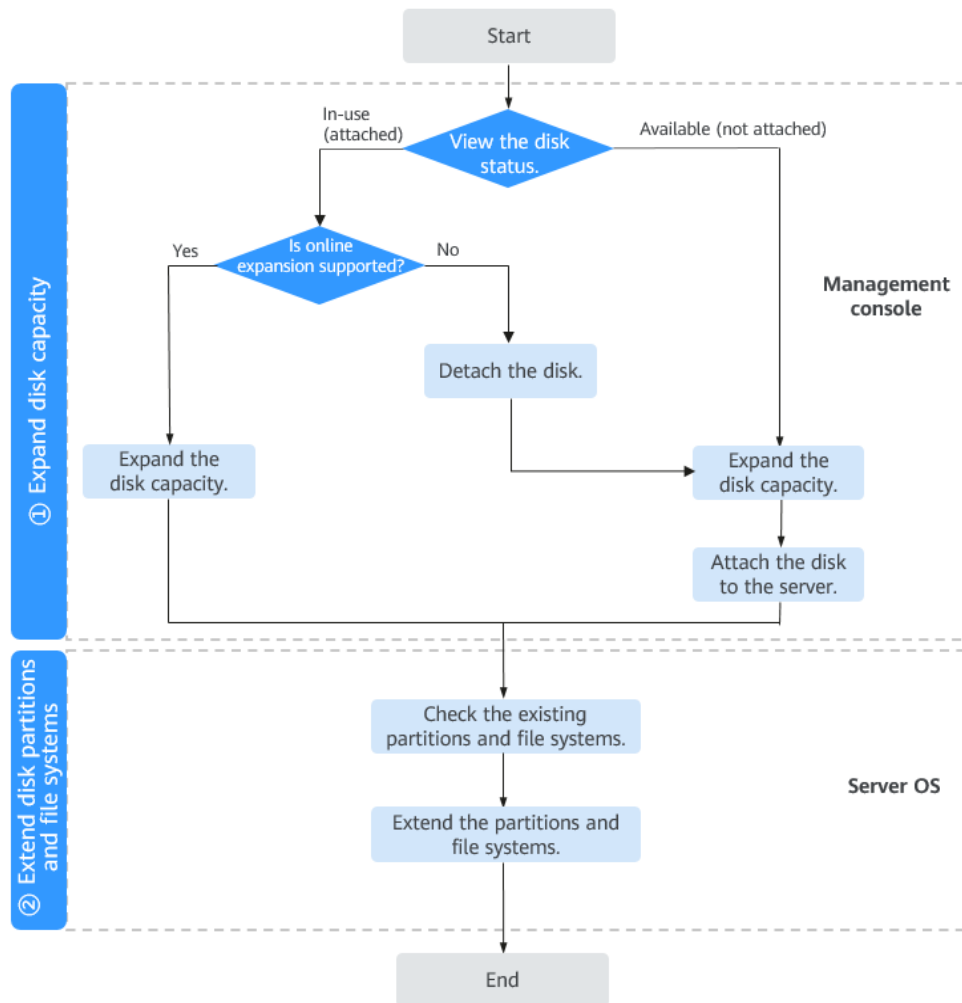
Step 1: Expand Disk Capacity

Step 2: Extend Disk Partitions and File Systems

 **NOTE**

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

Figure 5-1 Capacity expansion procedure



Capacity Expansion Charges

You will be billed for the additional capacity of a disk after you expand the disk capacity. The billing mode of the additional capacity is the same as that of the disk.

- For a pay-per-use disk: The new capacity takes effect immediately, so you will be billed for the new capacity of the disk immediately.
- For a yearly/monthly disk: You need to pay for the price difference after expanding the disk capacity. The disk expiration time remains unchanged.

For details about EVS billing, see [Billing for EVS Disks](#).

5.2 Step 1: Expand Disk Capacity

Scenarios

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

Prerequisites

NOTE

For how to view the disk status, see [Viewing EVS Disk Details](#).

Ensure that the disk meets the following conditions:

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

Ensure that the server meets the following conditions:

- If the disk status is **In-use**, the server status must be **Running** or **Stopped**.
- If the disk status is **In-use**, the server OS must meet the requirements describes in [Related Operations](#).

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

Notes and Constraints

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

Procedure

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

Step 2 Choose an expansion entry.

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

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

Step 4 In the displayed **Note** dialog box, read the note, and click **Expand Capacity**.

Step 5 On the **Expand Capacity** page, check the disk configuration.

- Click **Submit** to start the expansion of a pay-per-use disk. For a yearly/monthly disk, make the payment before you can continue.

- Click **Previous** to change the settings.

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

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

 **NOTE**

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

Step 7 (Optional) Skip this step if the disk status is **In-use** (attached to a server). Attach the disk to a server if the disk status is **Available**. For details, see [Attaching an EVS Disk](#).

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

The operations vary depending on the server OS.


- For Linux, see [Extending Disk Partitions and File Systems \(Linux\)](#).
- For Windows, see [Extending Disk Partitions and File Systems \(Windows\)](#).

----End

Related Operations

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

1. Check your server image. Certain public images and private images same as those public images allow you to expand **In-use** disks, and do not require the servers to be stopped and then started after the expansion.

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

2. If your server OS is not in the image list, check whether it is included in [Table 5-1](#).

If it is included in [Table 5-1](#), you can expand capacity while the disk is in use without the need to stop and start the server after the expansion. Otherwise, you must detach the disk and then expand its capacity, or stop and start the server after the expansion.

Table 5-1 OSs that support the capacity expansion of **In-use** disks

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

OS	Version
Debian	8.5.0 64-bit or later
Fedora	24 64-bit or later
SUSE 12	SUSE Linux Enterprise Server 12 64-bit or later
SUSE 11	SUSE Linux Enterprise Server 11 SP4 64-bit
OpenSUSE	42.1 64-bit or later
Oracle Linux Server release 7	7.2 64-bit or later
Oracle Linux Server release 6	6.7 64-bit or later
Ubuntu Server	14.04 64-bit or later
Red Hat Enterprise Linux 7	7.3 64bit
Red Hat Enterprise Linux 6	6.8 64bit
EulerOS	2.2 64-bit or later
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

5.3 Step 2: Extend Disk Partitions and File Systems

5.3.1 Extending Disk Partitions and File Systems (Linux)

Scenarios

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

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

Table 5-2 Operation instructions of extending partitions and file systems in Linux

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

Scenario	Partition Style	Disk Function	OS	File System Format	Capacity Expansion Tool	Example Configuration
Extending an Existing MBR Partition (for System Disks Whose Kernel Version Is Earlier Than 3.6.0)	MBR	System disk	The kernel version is earlier than 3.6.0.	ext* (such as ext2, ext3, and ext4), xfs, and btrfs	dracut-modules - growroot	<ul style="list-style-type: none"> • Device name: /dev/vda • File system format: ext4 • Mount point: /mnt/sda • Partition name: /dev/vda1 • Space added: 60 GiB • Partition style: MBR
Creating a New MBR Partition	MBR	System disk Data disk	Not limited	ext* (such as ext2, ext3, and ext4), xfs, and btrfs	<ul style="list-style-type: none"> • fdisk • parted 	<ul style="list-style-type: none"> • Partitioning tool: fdisk • Device name: /dev/vdb • File system format: ext4 • Mount points: /mnt/sdc and /mnt/sdd • Partition 1: /dev/vdb1 <ul style="list-style-type: none"> - Size: 100 GiB - Partition style: MBR • Partition 2: /dev/vdb2 <ul style="list-style-type: none"> - Size: 50 GiB - Partition style: MBR

Scenario	Partition Style	Disk Function	OS	File System Format	Capacity Expansion Tool	Example Configuration
Creating a New GPT Partition	GPT	Data disk	Not limited	ext* (such as ext2, ext3, and ext4), xfs, and btrfs	parted	<ul style="list-style-type: none"> ● Partitioning tool: parted ● Device name: /dev/vdb ● File system format: ext4 ● Mount points: /mnt/sdc and /mnt/sdd ● Partition 1: /dev/vdb1 <ul style="list-style-type: none"> - Size: 2 TiB - Partition style: GPT ● Partition 2: /dev/vdb2 <ul style="list-style-type: none"> - Size: 1 TiB - Partition style: GPT

 NOTE

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

For how to extend 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?](#)

If the disk is not partitioned, see [How Do I Extend the File System of an Unpartitioned Data Disk in Linux?](#)

Notes and Constraints

- The additional space of a data disk cannot be added to the root partition. To extend the root partition, expand the system disk instead.
- During an expansion, the additional space is added to the end of the disk. If the disk has multiple partitions, the additional space can only be allocated to the partition at the disk end.
- If a disk uses MBR, the storage space in excess of 2 TiB cannot be used because the maximum capacity that MBR supports is 2 TiB. If your disk

already uses MBR for partitioning and you require more than 2 TiB after the capacity expansion, do as follows:

- (Recommended) Create a new EVS disk and use GPT.
- Back up the disk data, perform the expansion, and then change the partition style from MBR to GPT. During this change, services will be interrupted and data on the disk will be erased.

Prerequisites

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

Procedure

Extending an Existing Partition

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

Step 1 Check whether the growpart tool is installed.

growpart

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

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

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

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

yum install cloud-utils-growpart

```
Loaded plugins: fastestmirror
Determining fastest mirrors
```

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

The installation is successful.

Step 2 Check the partition information of the `/dev/vdb` disk.

lsblk

```
[root@ecs-centos76 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda 253:0 0 40G 0 disk
└─vda1 253:1 0 40G 0 part /
vdb 253:16 0 150G 0 disk
└─vdb1 253:17 0 100G 0 part /mnt/sdc
```

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

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

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

growpart /dev/vdb 1

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

NOTE

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

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

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

parted /dev/vdb

P

```
[root@ecs-centos74 ~]# parted /dev/vdb
GNU Parted 3.1
Using /dev/vdb
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) p
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 107GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:
```

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

(parted)

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

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

resize2fs /dev/vdb1

```
[root@ecs-test-0001 ~]# resize2fs /dev/vdb1
resize2fs 1.42.9 (28-Dec-2013)
Filesystem at /dev/vdb1 is mounted on /mnt/sdc; on-line resizing required
old_desc_blocks = 13, new_desc_blocks = 19
The filesystem on /dev/vdb1 is now 39321339 blocks long.
```

NOTE

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

sudo xfs_growfs /mnt/sdc

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

Step 5 Check the partition size after extension.

lsblk

```
[root@ecs-centos76 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda 253:0 0 40G 0 disk
├vda1 253:1 0 40G 0 part /
vdb 253:16 0 150G 0 disk
├vdb1 253:17 0 150G 0 part /mnt/sdc
```

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

NOTE

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

----End

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

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

NOTICE

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

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

yum install dracut-modules-growroot

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

NOTE

Skip this step if the tool is already installed.

Step 2 Regenerate the `initramfs` file.

dracut -f

NOTE

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

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

lsblk

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

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

Step 4 Restart the server.

reboot

Reconnect to the server after it is restarted.

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

lsblk

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

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

----End

Creating a New MBR Partition

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

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

1. Check disk partition sizes.

lsblk

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

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

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

2. Check the disk partition style.

parted /dev/vdb

p

```
[root@ecs-test-0001 ~]# parted /dev/vdb
GNU Parted 3.1
Using /dev/vdb
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) p
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 161GiB
Sector size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:
```

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

(parted)

In this example, the disk uses MBR.

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

 **NOTE**

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

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

1. Create the partition.

fdisk /dev/vdb

n

p

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

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

```
Device does not contain a recognized partition table  
Building a new DOS disklabel with disk identifier 0x38717fc1.
```

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

```
Partition type:
```

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

```
  e  extended
```

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

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

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

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

 **NOTE**

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

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

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

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

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

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

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

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

```
using default value 83886080
```

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

```
using default value 209715199
```

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

Command (m for help):

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

NOTE

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

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

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

- Sector value of the data disk **/dev/vdb** (100 GiB) is **209715200** ($100 \times 1073741824/512$), so the disk's last sector is **209715199** ($209715200 - 1$).
In the preceding example, **First sector (2048-209715199, default 2048)** means that the first sector of the disk ranges from **2048** to **209715199**.
- Sector value of the first primary partition **/dev/vdb1** (40 GiB) is **83886080** ($40 \times 1073741824/512$), so the partition's last sector is **83886079** ($83886080 - 1$).
In this example, the default first sector is used as first sector of this partition, which is **2048**.
- Sector value of the second primary partition **/dev/vdb2** (60 GiB) is **125829120** ($60 \times 1073741824/512$), so the partition's last sector is **125829119** ($125829120 - 1$).

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

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

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

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

Command (m for help): p

```
Disk /dev/vdb: 107.4 GB, 107374182400 bytes, 209715200 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x994727e5
```

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

Command (m for help):

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

NOTE

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

3. Synchronize the new partition table to the OS.

partprobe

 NOTE

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

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

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

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

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

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

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

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

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

 NOTE

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

```
[root@ecs-test-0001 ~]# mkfs -t ext4 /dev/vdb2
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
2621440 inodes, 10485504 blocks
524275 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2157969408
320 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000, 7962624

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

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

```
[root@ecs-test-0001 ~]# parted /dev/vdb
GNU Parted 3.1
Using /dev/vdb
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) p
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 107GiB
Sector size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:
```



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

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

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

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

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

mkdir -p /mnt/sdd

mount /dev/vdb2 /mnt/sdd

lsblk

View the mount results.

```
[root@ecs-test-0001 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda 253:0 0 40G 0 disk
├vda1 253:1 0 40G 0 part /
vdb 253:16 0 150G 0 disk
├vdb1 253:17 0 100G 0 part /mnt/sdc
└vdb2 253:18 0 50G 0 part /mnt/sdd
```

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

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

 **NOTE**

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

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

blkid /dev/vdb2

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

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

2. Configure auto mount at startup.

vi /etc/fstab

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

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

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

Table 5-3 Parameter description

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

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

Step 7 Verify that auto mount takes effect.

```
umount /dev/vdb2
```

```
mount -a
```

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

Query file system mounting information.

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

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

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

----End

Creating a New GPT Partition

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

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

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

1. Check disk partition sizes.

lsblk

```
[root@ecs-test-0001 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda 253:0 0 40G 0 disk
└─vda1 253:1 0 40G 0 part /
vdb 253:16 0 3T 0 disk
└─vdb1 253:17 0 2T 0 part /mnt/sdc
```

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

2. Check the disk partition style.

parted /dev/vdb

p

```
[root@ecs-test-0001 ~]# parted /dev/vdb
GNU Parted 3.1
Using /dev/vdb
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) p
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 3299GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:
```

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

(parted)

In this example, the disk uses GPT.

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

NOTE

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

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

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

parted /dev/vdb

unit s

p

```
[root@ecs-centos74 ~]# parted /dev/vdb
GNU Parted 3.1
Using /dev/vdb
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) unit s
(parted) p
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 6442450944s
Sector size (logical/physical): 512B/512B
```

```
Partition Table: gpt
Disk Flags:

Number Start      End              Size            File system  Name      Flags
 1    2048s    4294965247s    4294963200s    ext4        /dev/vdb1
(parted)
```

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

NOTE

- If error message **-bash: parted: command not found** is returned, the system cannot identify the command. In this case, run **yum install -y parted** to install the command. Then run the command again.
- If the following error information is displayed, enter **Fix**.
Error: The backup GPT table is not at the end of the disk, as it should be. This might mean that another operating system believes the disk is smaller. Fix, by moving the backup to the end (and removing the old backup)?
Fix/Ignore/Cancel?

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

- If the following warning information is displayed, enter **Fix**.
Warning: Not all of the space available to /dev/vdb appears to be used, you can fix the GPT to use all of the space (an extra 104857600 blocks) or continue with the current setting?
Fix/Ignore?

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

2. Set the partition name and size.

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

```
p
```

NOTE

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

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

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

```
lsblk
```

```
[root@ecs-centos74 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda   253:0  0  40G  0 disk
├─vda1 253:1  0  40G  0 part /
vdb   253:16  0   3T  0 disk
├─vdb1 253:17  0   2T  0 part /mnt/sdc
└─vdb2 253:18  0   1T  0 part
```

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

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

NOTE

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

```
[root@ecs-test-0001 ~]# mkfs -t ext4 /dev/vdb2
mke2fs 1.42.9 (28-Dec-2013)
```

```
Filesystem label=  
OS type: Linux  
Block size=4096 (log=2)  
Fragment size=4096 (log=2)  
Stride=0 blocks, Stripe width=0 blocks  
67108864 inodes, 268435456 blocks  
13421772 blocks (5.00%) reserved for the super user  
First data block=0  
Maximum filesystem blocks=2415919104  
8192 block groups  
32768 blocks per group, 32768 fragments per group  
8192 inodes per group  
Superblock backups stored on blocks:  
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,  
    4096000, 7962624, 11239424, 20480000, 23887872, 71663616, 78675968,  
    102400000, 214990848  
  
Allocating group tables: done  
Writing inode tables: done  
Creating journal (32768 blocks): done  
Writing superblocks and filesystem accounting information: done  
[root@ecs-test-0001 ~]#
```

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

```
[root@ecs-test-0001 ~]# parted /dev/vdb  
GNU Parted 3.1  
Using /dev/vdb  
Welcome to GNU Parted! Type 'help' to view a list of commands.  
(parted) p  
Model: Virtio Block Device (virtblk)  
Disk /dev/vdb: 3299GB  
Sector size (logical/physical): 512B/512B  
Partition Table: gpt  
Disk Flags:  
  
Number  Start  End    Size  File system  Name      Flags  
1       1049kB 2199GB 2199GB ext4        /dev/vdb1  
2       2199GB 3299GB 1100GB ext4        /dev/vdb2  
  
(parted) q  
[root@ecs-test-0001 ~]#
```

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

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

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

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

```
lsblk
```

```
[root@ecs-test-0001 ~]# lsblk  
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT  
vda 253:0 0 40G 0 disk  
└vda1 253:1 0 40G 0 part /  
vdb 253:16 0 3T 0 disk  
└vdb1 253:17 0 2T 0 part /mnt/sdc  
└vdb2 253:18 0 1T 0 part /mnt/sdd
```

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

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

 **NOTE**

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

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

blkid /dev/vdb2

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

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

2. Configure auto mount at startup.

vi /etc/fstab

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

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

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

Table 5-4 Parameter description

Example Value	Description
UUID=0b3040e2-1367-4abb-841d-ddb0b92693df	The UUID of the partition.
/mnt/sdc	The mount point of the partition.
ext4	The file system format of the partition.
defaults	The partition mount option. Normally, this parameter is set to defaults .
0	<ul style="list-style-type: none"> - The Linux dump backup option. <ul style="list-style-type: none"> ▪ 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.

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

Step 6 Verify that auto mount takes effect.

```
umount /dev/vdb2
```

```
mount -a
```

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

Query file system mounting information.

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

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

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

----End

5.3.2 Extending Disk Partitions and File Systems (Windows)

Scenarios

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

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

- Extending an Existing Partition
- Creating a New Partition

Notes and Constraints

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

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

Prerequisites

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

Procedure

Extending an Existing Partition

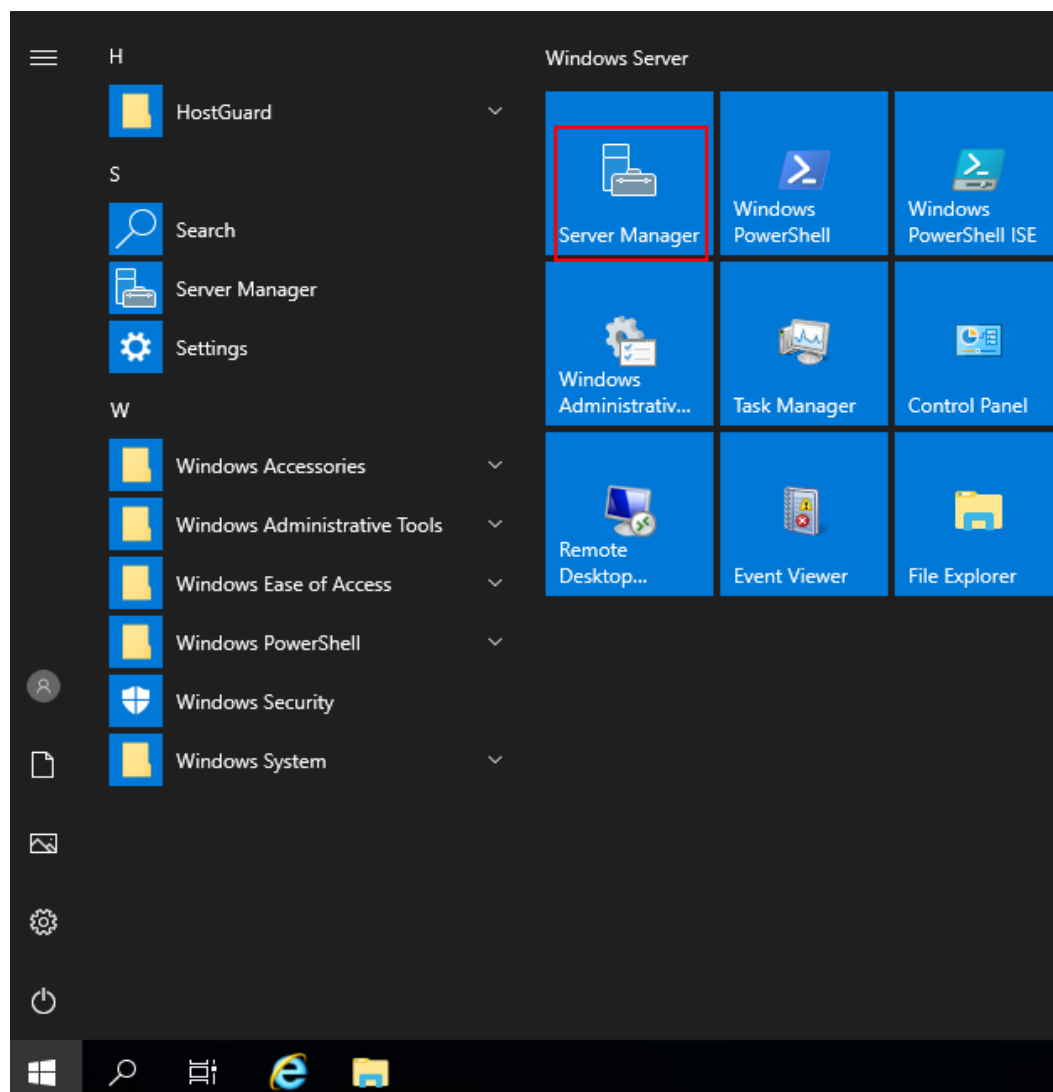
Originally, the D drive in the Windows Server 2019 has 60 GiB, and then 30 GiB is added to the disk. The following example shows you how to allocate the additional 30 GiB to the D drive.

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

The **Windows Server** window is displayed.

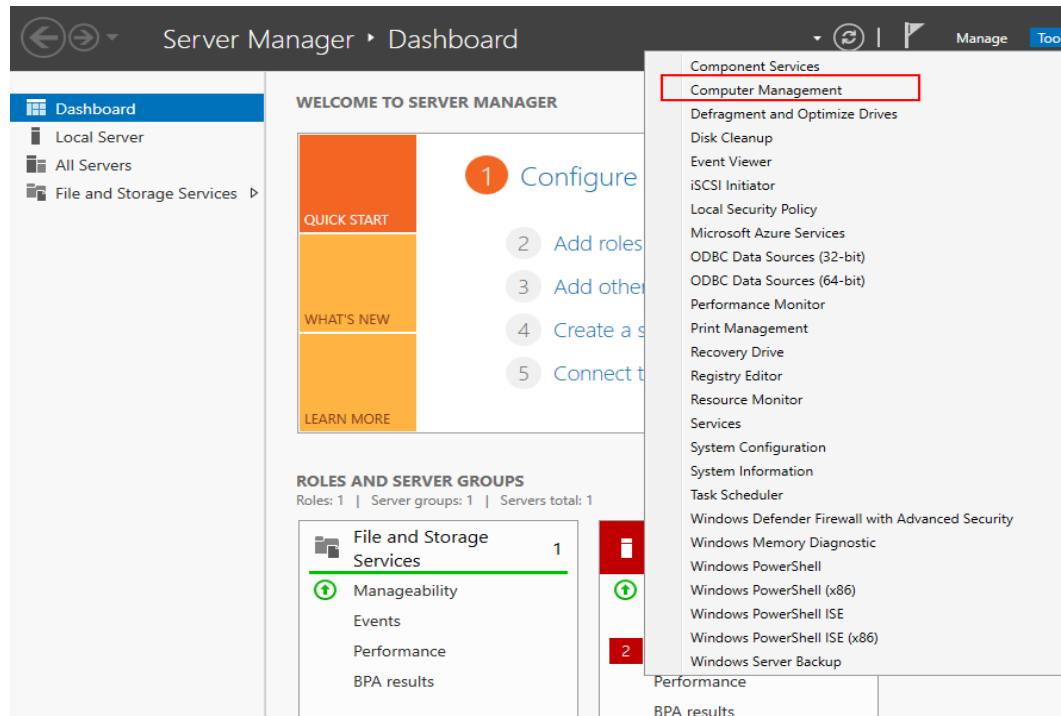
Step 2 Click **Server Manager** to open **Server Manager**.

Figure 5-2 Server Manager



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


Figure 5-3 Computer Management



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

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

Figure 5-4 Disk expanded but additional space not allocated




 **NOTE**

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

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

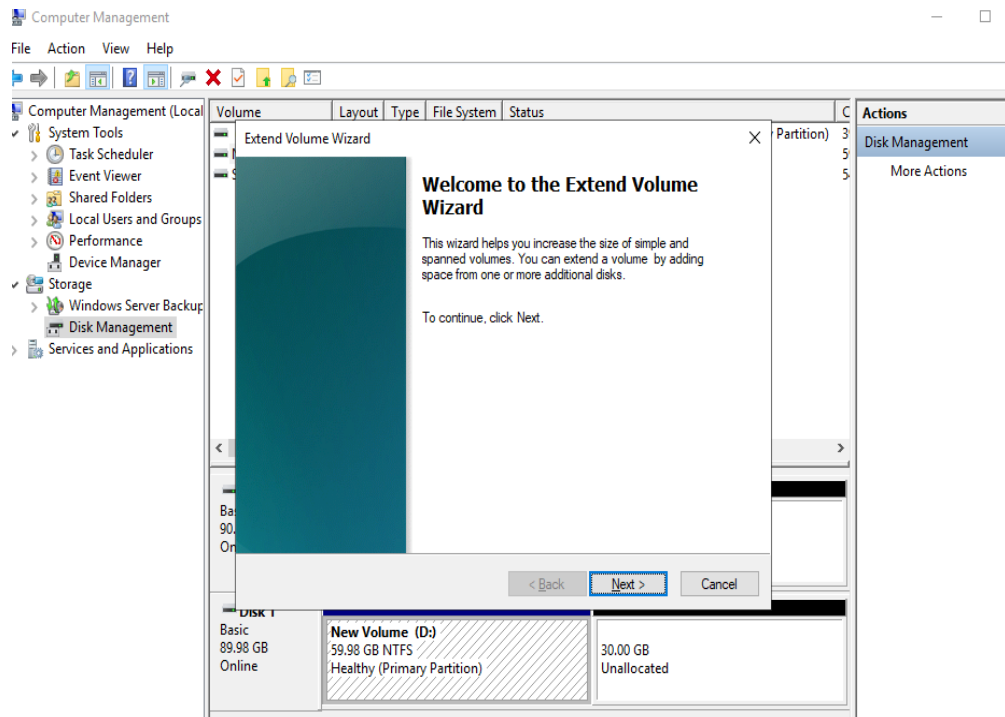
Step 6 Right-click the volume and choose **Extend Volume** from the shortcut menu. In this example, right-click **New Volume (D:)**.

Figure 5-5 Choosing Extend Volume



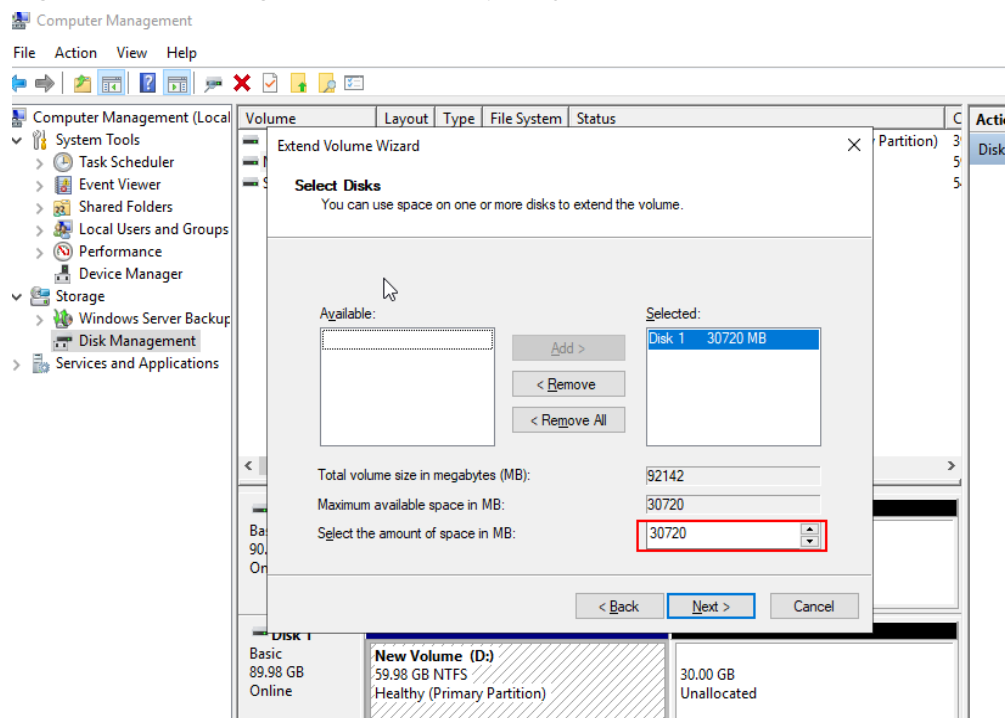
On the displayed **Extend Volume Wizard** windows, click **Next**.

Figure 5-6 Extend Volume Wizard



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

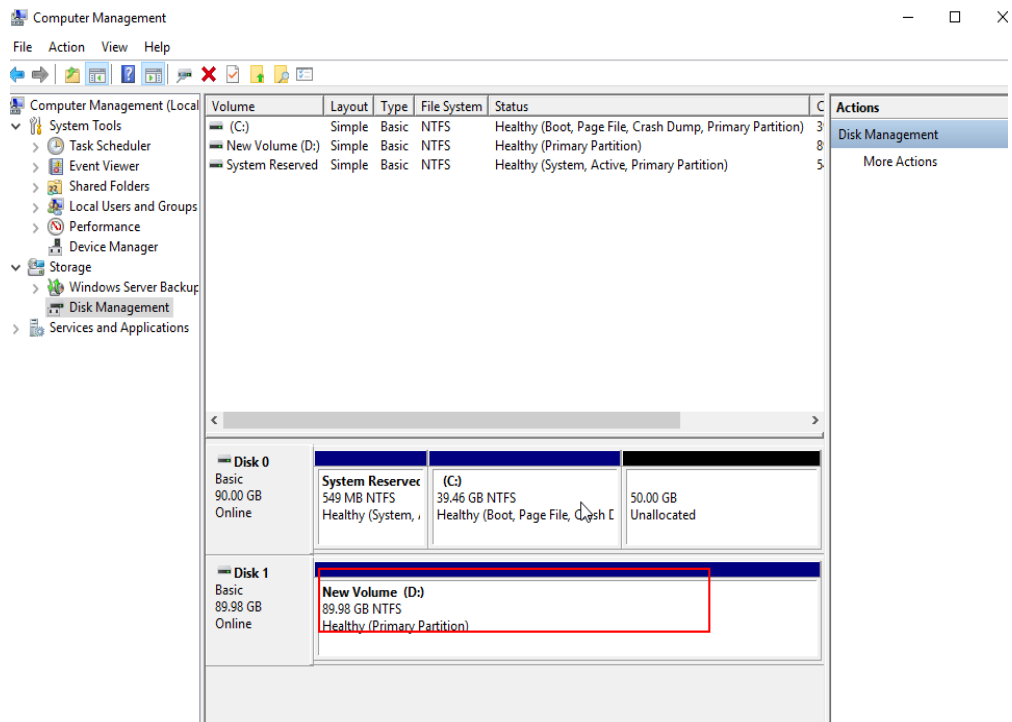
Figure 5-7 Selecting the amount of space you want to add



Step 8 Click **Finish**.

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

Figure 5-8 Extension succeeded



----End

Creating a New Partition

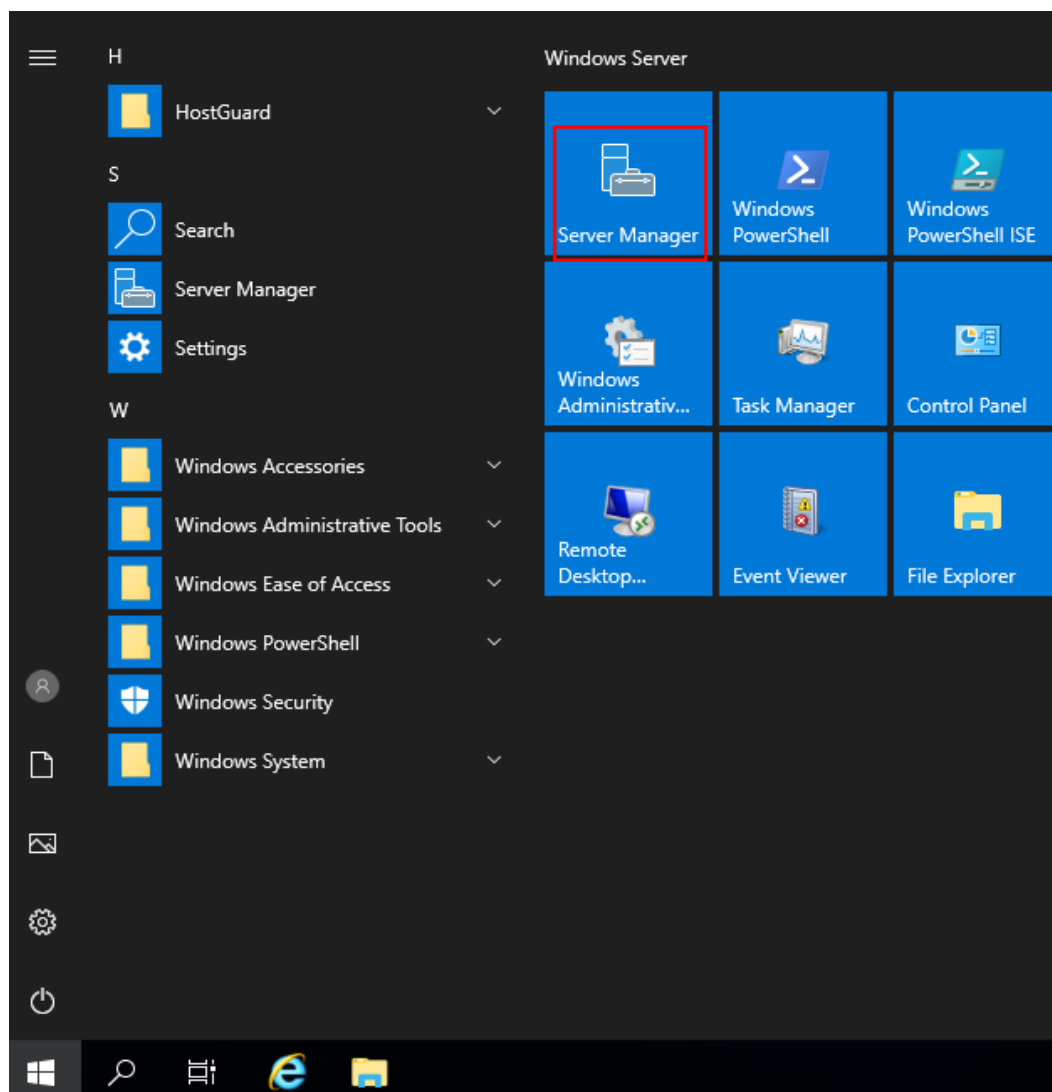
The following example shows you how to create a GPT partition with an NTFS file system on a server running Windows Server 2019.

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

The **Windows Server** window is displayed.

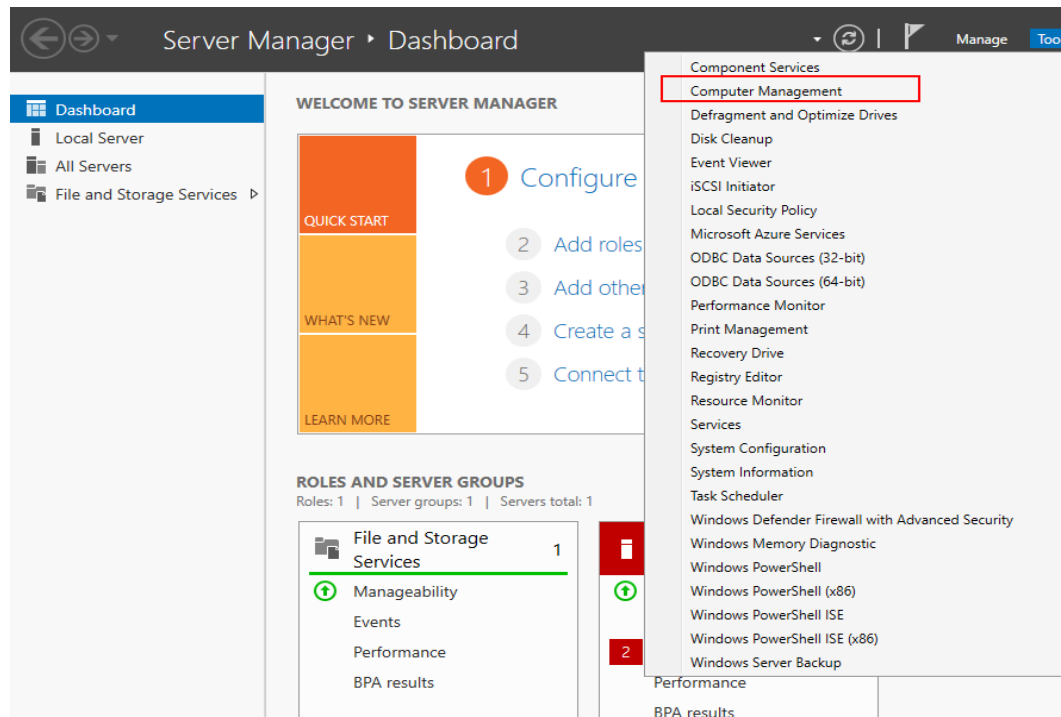
Step 2 Click **Server Manager** to open **Server Manager**.

Figure 5-9 Server Manager



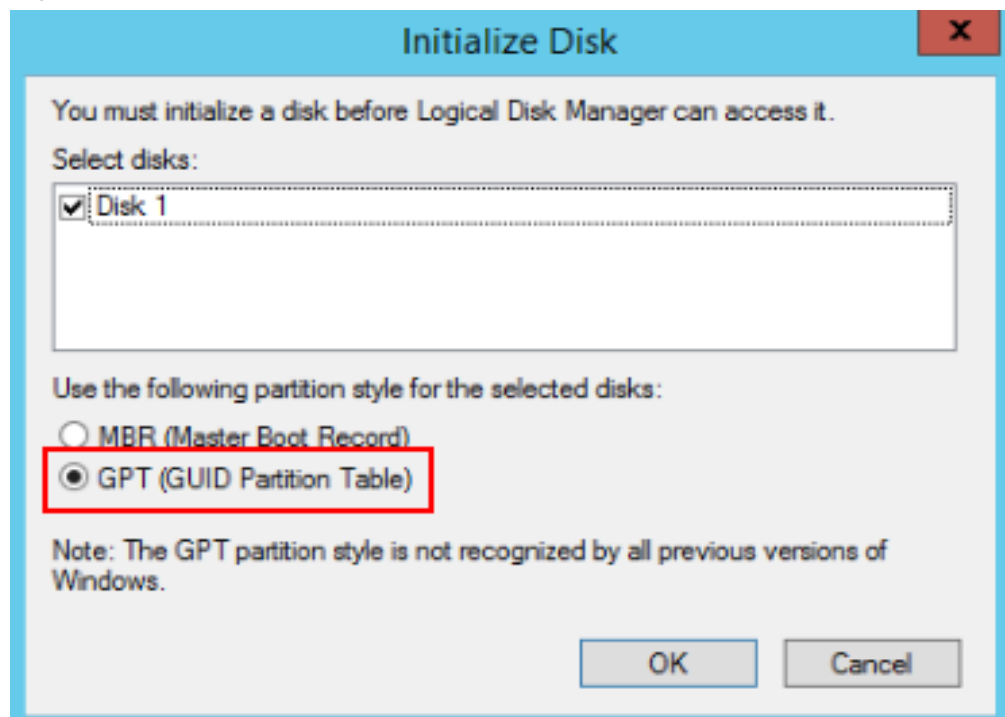
Step 3 In the upper right corner, choose **Tools > Computer Management**.

Figure 5-10 Computer Management



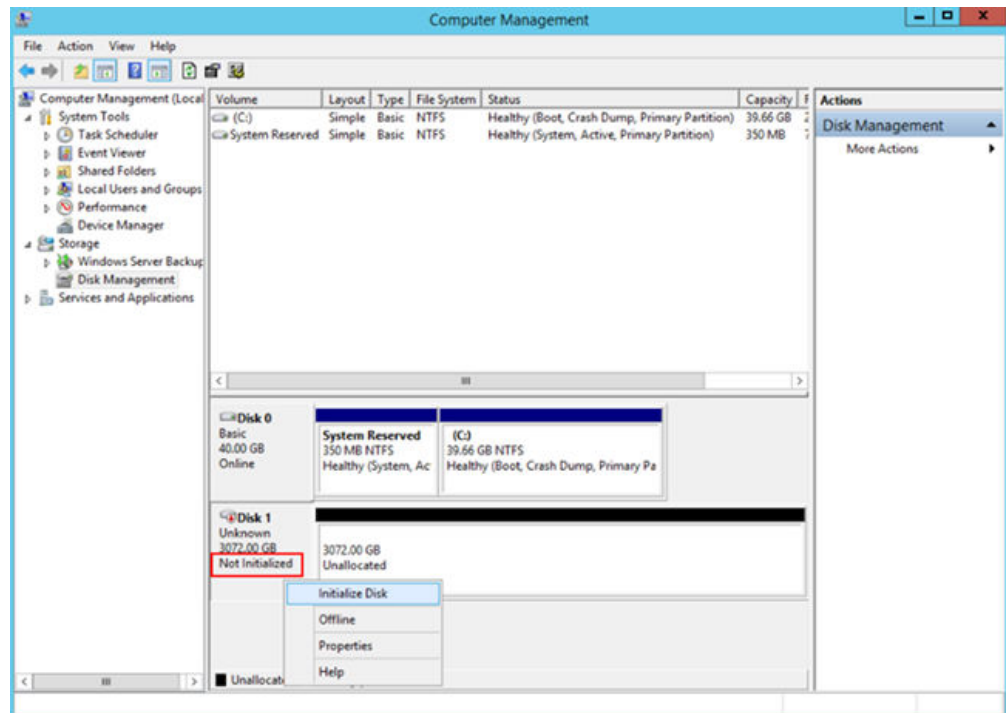
- Step 4** Choose **Storage > Disk Management** to go to the disk list page.
- Disks are displayed in the right pane. If there is a disk that is not initialized, the system will prompt you with the **Initialize Disk** dialog box.

Figure 5-11 Disk list



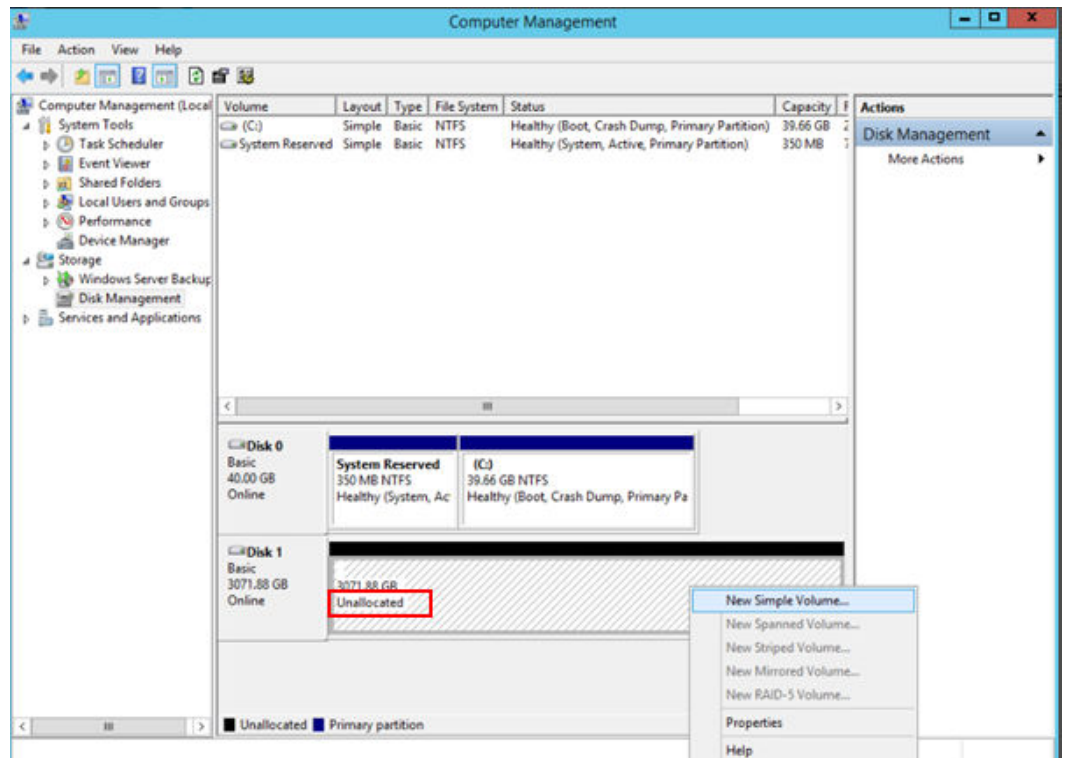
- If the **Initialize Disk** dialog box is not prompted up and the disk has no partitions (entire disk shown as **Unallocated**), right-click the area where the to-be-initialized disk is and choose **Initialize Disk** from the shortcut menu.

Figure 5-12 Initialize Disk



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

Figure 5-13 Disk expanded but additional space not allocated



NOTE


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

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

NOTE


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

Figure 5-14 Selecting GPT



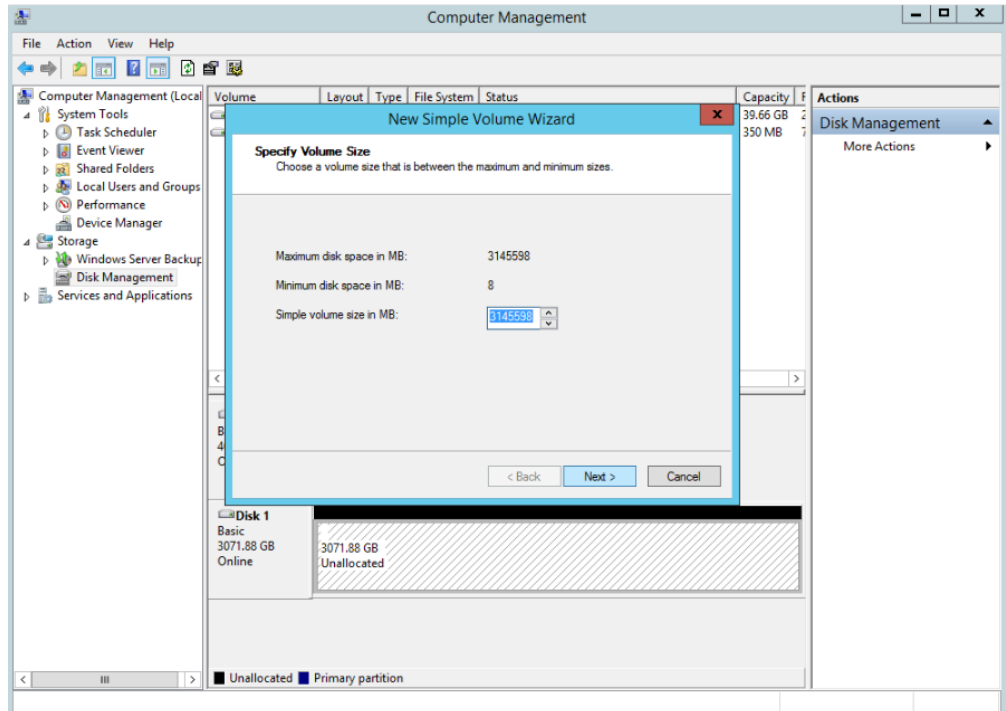
Step 6 In the **Unallocated** area of **Disk 1**, right-click and choose **New Simple Volume** from the shortcut menu and initialize the disk as prompted.

Figure 5-15 New Simple Volume



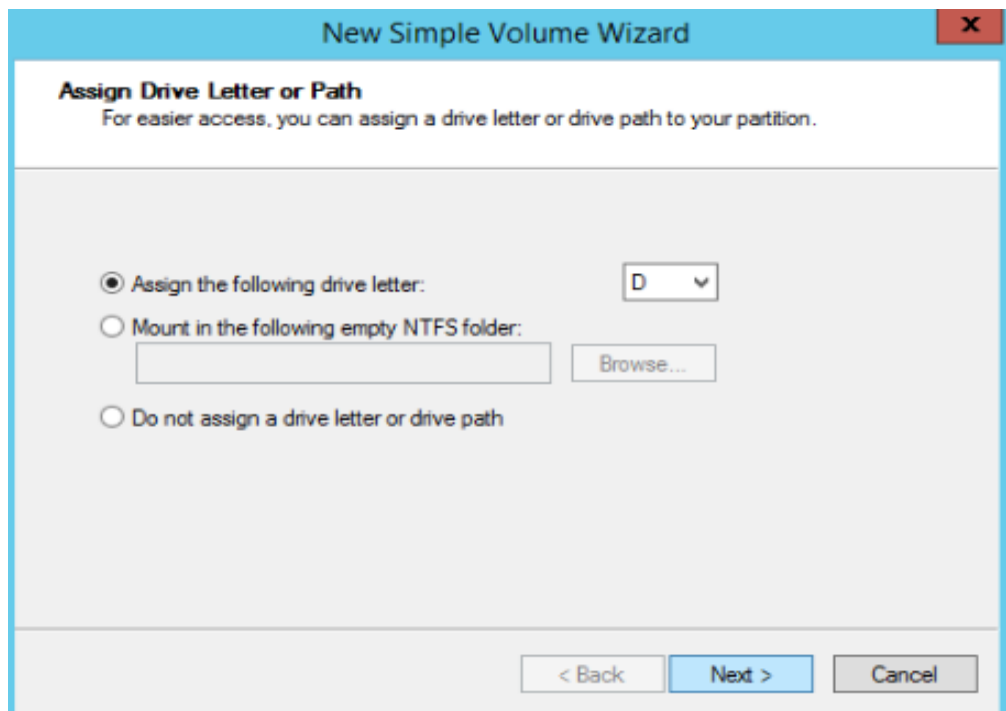
1. On the **Specify Volume Size** page, retain the default settings and click **Next**. The system uses the maximum disk space as the default volume size. You can specify a volume size as needed.

Figure 5-16 Specify Volume Size




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

Figure 5-17 Assign Drive Letter or Path



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

Figure 5-18 Format Partition



NOTICE

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

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

Figure 5-19 Completing the New Simple Volume Wizard

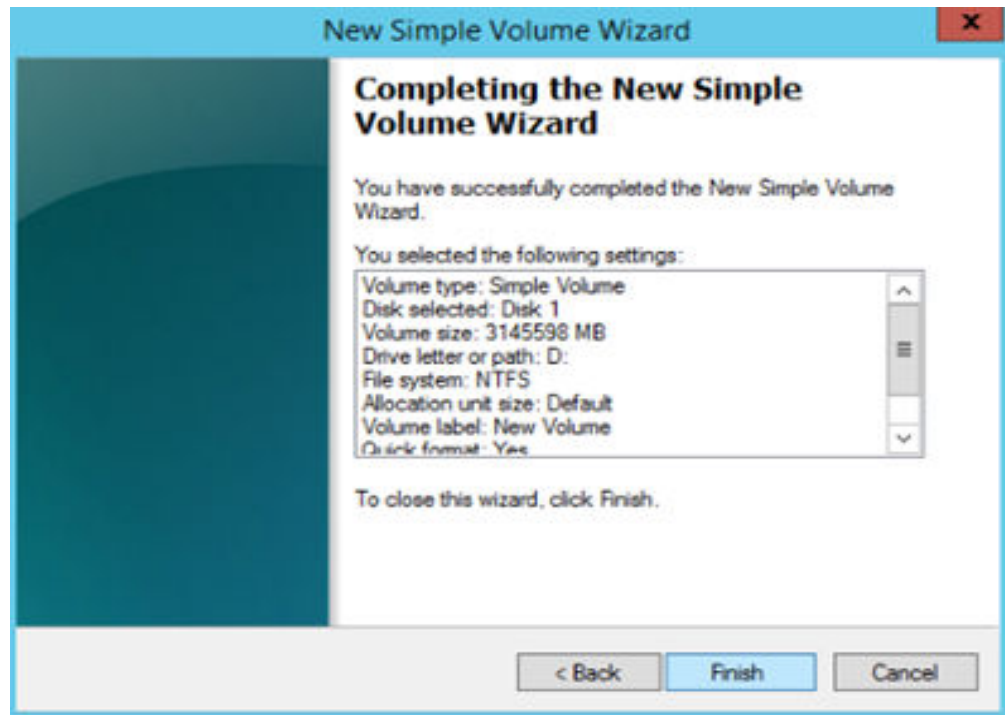




Figure 5-20 Viewing the initialization results

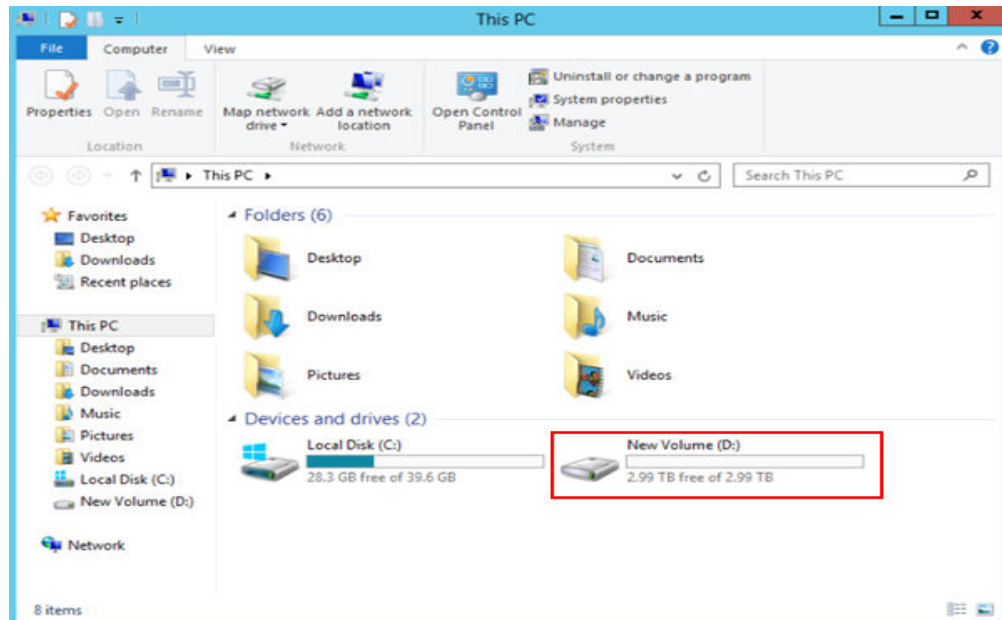


Step 7 (Optional) Alternatively, choose **Server Manager > File and Storage Services > Volumes > Disks** to view the disk status, capacity, and partition style.

Step 8 After the volume is created, click  on the task bar and check whether a new volume appears in the File Explorer. In this example, New Volume (D:) is the new volume.

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

Figure 5-21 File Explorer



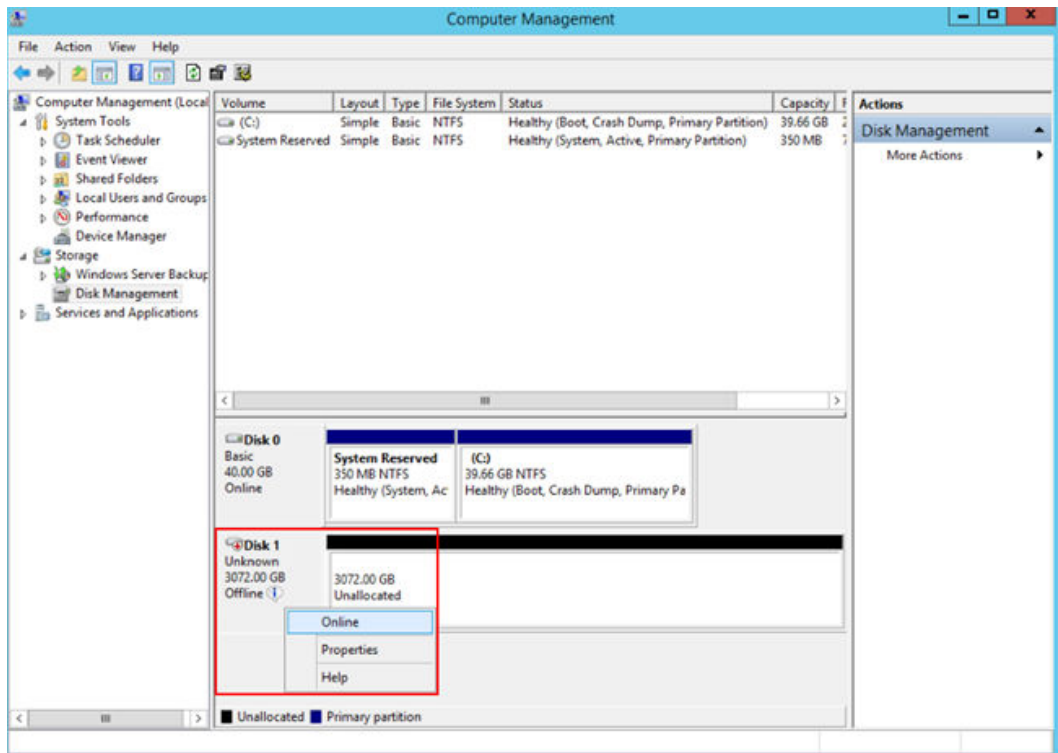
----End

Related Operations

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

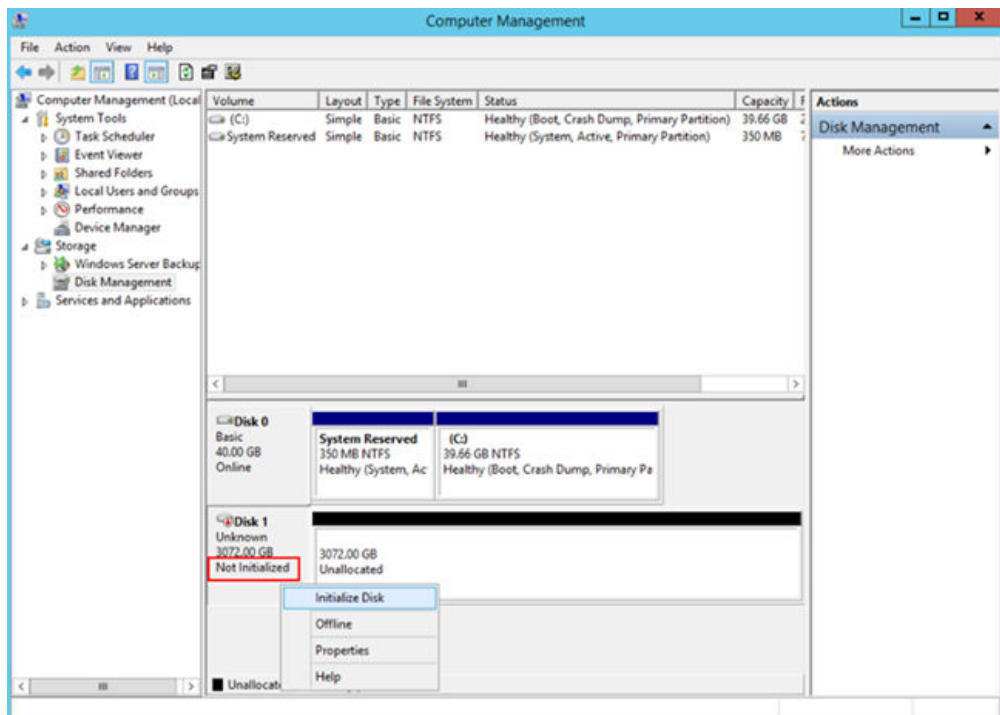
Step 1 In the **Disk 1** area, right-click and choose **Online** from the shortcut menu.

Figure 5-22 Online



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

Figure 5-23 Online succeeded



----End

6 Detaching and Deleting an EVS Disk

6.1 Detaching an EVS Disk

Scenarios

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

 NOTE

- For an attached system disk, the disk function is displayed as **System disk**, and the disk status is displayed as **In-use** in the disk list. After the system disk is detached, the disk function changes to **Bootable disk**, and the status changes to **Available**.
- Bootable disks are the system disks detached from servers. A bootable disk can be re-attached to a server to be used as a system disk or data disk depending on the disk function selected.
- For an attached data disk, the disk function is displayed as **Data disk**, and the disk status is displayed as **In-use** in the disk list. After the data disk is detached, the disk function remains unchanged, and the status changes to **Available**. For a shared disk, the status changes to **Available** only after it is detached from all its servers.
- A detached disk will not be automatically deleted, and it will still be billed. To avoid unintended charges, you can delete or unsubscribe from it if it is no longer needed.

Notes and Constraints

- After a system disk is detached, some operations cannot be performed on the original server and the system disk. The restrictions are as follows:
 - Server: starting the server, remote login, resetting the password, changing server billing mode, changing server specifications, changing the OS, reinstalling the OS, creating images, creating backups, adding disks, changing the security group, and changing the VPC
 - System disk: changing disk billing mode

Prerequisites

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

Detaching a System Disk

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

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

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

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

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

Step 4 Click the name of this server.

The server details page is displayed.

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

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

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

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


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

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

----End

Detaching a Non-Shared Data Disk

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

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

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

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

- If yes, perform the following procedure:
 - a. In the disk list, click the name of the to-be-detached disk.
The disk details page is displayed.
 - b. Click the **Servers** tab to view the server where the disk has been attached.
 - c. Click to select the server and click **Detach Disk**.
The **Detach Disk** dialog box is displayed.
 - d. Click **Yes** to detach the disk.
- If no, perform the following procedure:
 - a. In the disk list, locate the row that contains the target disk and choose **More > Detach** in the **Operation** column.
The **Detach Disk** dialog box is displayed.
 - b. Click **Yes** to detach the disk.

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

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

----End

Helpful Links

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

6.2 Unsubscribing from or Deleting an EVS Disk

Scenarios

If an EVS disk is no longer used, you can delete the disk to release the virtual resources. When a disk is deleted, EVS immediately destroys the metadata to ensure that data can no longer be accessed. In addition, the physical storage space of the disk is reclaimed and cleared before being re-assigned. For any new disk created based on the re-assigned physical space, before data is written to the disk, EVS returns zero for all the read requests to the disk.

When deleting a disk, you can choose not to delete the disk immediately, but move it to the recycle bin to prevent any data loss that may be caused by unintended operations. The recycle bin function is disabled by default. If you need to use this function, enable it on the console. For details, see [Enabling the Recycle Bin](#).

Yearly/Monthly disks cannot be deleted. You can unsubscribe from them if needed. System disks must be unsubscribed from together with their servers. For how to unsubscribe from data disks, see [Table 6-1](#).

Table 6-1 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 a Yearly/Monthly Disk on the EVS Console Unsubscribing from a Yearly/Monthly Disk 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 a Yearly/Monthly Disk on the EVS Console

Unsubscription Scenario	Sub-scenario	Reference
Unsubscribing from yearly/monthly data disks that were purchased on the EVS console	Unsubscribing from data disks separately	Unsubscribing from a Yearly/Monthly Disk on the EVS Console

 NOTE

You will be not billed for a disk after it is deleted or unsubscribed from.

Notes and Constraints

- The disk status is **Available**, **Error**, **Expansion failed**, **Restoration failed**, or **Rollback failed**.
- The disk is not locked by any service.
- The shared disk has been detached from all its servers.
- 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.
- 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**.
- Yearly/Monthly data disks purchased on the EVS console have different expiration times as the server. They can be unsubscribed from separately.


NOTICE

When you delete a disk, all the disk data including the legacy snapshots created for this disk will be deleted.

A deleted disk cannot be recovered.

Deleting Pay-per-Use EVS Disks



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

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



Step 3 In the disk list, locate the 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

Unsubscribing from a Yearly/Monthly Disk on the EVS Console

- Step 1** Log in to the [console](#).
- Step 2** Click  in the upper left corner and choose **Storage > Elastic Volume Service**.
The **Elastic Volume Service** page is displayed.
- Step 3** In the disk list, locate the 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 the information 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

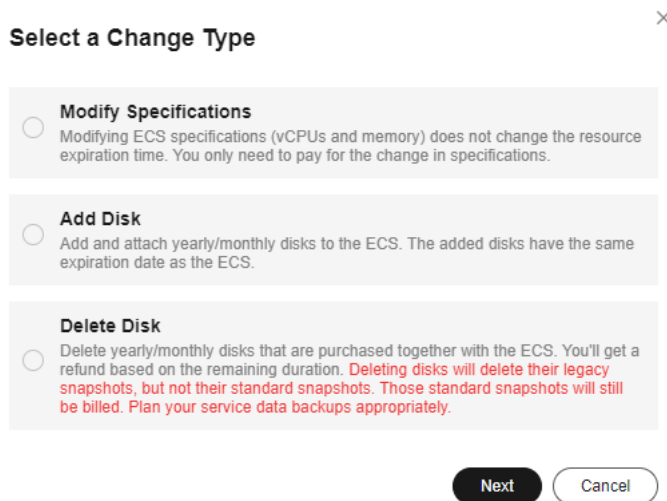
Unsubscribing from a Yearly/Monthly Disk on the ECS Console

-  **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 [console](#).
- Step 2** Click  . Choose **Compute > 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 6-1 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

Helpful Links

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

6.3 Unsubscribing from Yearly/Monthly EVS Disks

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 6-2](#).

Notes and 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**.
- Yearly/Monthly data disks purchased on the EVS console have different expiration times as the server. They can be unsubscribed from separately.

Table 6-2 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 [console](#).

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

Step 3 In the disk list, locate the target disk and choose **More > 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 [console](#).

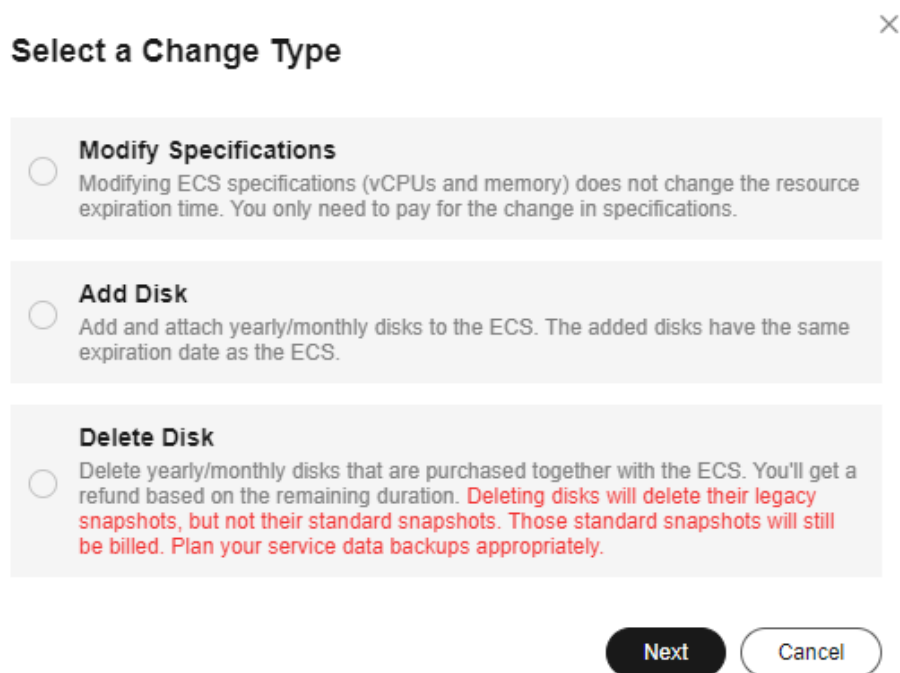
Step 2 Click  . Choose **Compute > Elastic Cloud Server**.

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 6-2 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

7 Managing EVS Recycle Bin

7.1 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

See the EVS recycle bin region availability in [Function Overview](#).

Notes and Constraints

- When you delete a disk, regardless of whether the disk will be moved to the recycle bin or not, legacy snapshots of the disk will always be deleted permanently.
- There are no limits on the capacity and quantity of disks in the recycle bin.
- Separately deleted disks are stored in the recycle bin for a maximum of seven days. During this period, you can recover or permanently delete the disks. After the disks expire, they are permanently deleted and 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.

Billing Description

EVS disks in the recycle bin are billed on a pay-per-use basis. For details, see [Billing for EVS Recycle Bin](#).

If you want to view the bills of the recycle bin, see [How Do I View the EVS Recycle Bin Bill?](#)

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 or unsubscribe from yearly/monthly disks before they expires.
- 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:

- Your account is restricted or frozen.
- The number of days elapses since the disk creation is less than what you specified in the recycle bin policy.
- The pay-per-use disks you deleted or yearly/monthly disks unsubscribed from are already in a grace or retention period.
- The system permanently deletes the pay-per-use disks or yearly/monthly disks whose retention period has expired.

Recycle Bin Policy Configuration Suggestions

When using the recycle bin, you need to configure a recycle bin policy. For details, see [Configuring a Recycle Bin Policy](#).

Both ECS and EVS support the recycle bin function, and their recycle bin policies can be configured separately. If you use both recycle bins, you are advised to configure the same minimum number of days for moving ECSs and EVS disks to the recycle bins to avoid issues brought by different lifecycles.

In special cases, you may need to configure different recycle bin policies for ECS and EVS. For details, see [Table 7-1](#).

Assume that an ECS (including a system disk and a data disk) was created 8 days ago. If you delete or subscribe from the ECS, the results would be as follows.

Recycle Bin Operations

Table 7-1 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

Operation	Description	Reference
Disable the recycle bin.	You can disable the recycle bin if you no longer need it. 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.	You can permanently delete the EVS disks from the recycle bin at any time.	Permanently Deleting Disks from the Recycle Bin

7.2 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.

Notes and 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.
- Separately deleted disks are stored in the recycle bin for a maximum of seven days. During this period, you can recover or permanently delete the disks. After the disks expire, they are permanently deleted and cannot be recovered.

Procedure

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


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

Step 3 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.

Figure 7-1 Enable Recycle Bin



----End

Related Operations

You can **disable** the recycle bin if you no longer need it.

7.3 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.


Notes and Constraints

- The default recycle bin policy is 7 days, which means that disks created a least 7 days ago will be moved to the recycle bin upon deletion or unsubscription. You can customize a recycle bin policy ranging from 7 to 1,000 days as required.
- When the OS of a server 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

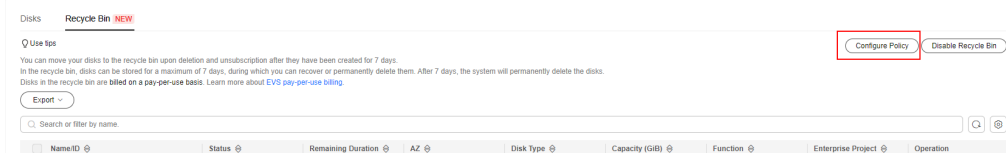
Step 1 Log in to the [console](#).

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

Figure 7-2 Configure Policy



Step 5 Set a 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

If you set 7 days for the recycle bin policy, disks created less than 7 days will not be moved to the recycle bin upon deletion and will be deleted permanently. 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 **OK**. Then, disks will be moved to the recycle bin upon deletion or unsubscription and deleted permanently based on the recycle bin policy.

----End

Related Operations

You can manually [recover](#) or [permanently delete](#) disks in the recycle bin.

7.4 Recovering Disks from the Recycle Bin

Scenarios

Before recycle bin disks expire, you can recover them from the recycle bin if needed.

Notes and 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 [console](#).

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

The **Elastic Volume Service** 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

7.5 Permanently Deleting Disks from the Recycle Bin

Scenarios

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

Notes and Constraints


- Separately deleted disks are stored in the recycle bin for a maximum of seven days. During this period, you can recover or permanently delete the disks. After the disks expire, they are permanently deleted and cannot be recovered.

 CAUTION

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

Permanently Deleting an EVS Disk

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

- Step 2** Click  in the upper left corner and choose **Storage > Elastic Volume Service**.
The **Elastic Volume Service** page is displayed.
- Step 3** 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

7.6 Disabling the Recycle Bin

Scenarios


You can disable the recycle bin as needed.

Notes and 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 [console](#).
- Step 2** Click  in the upper left corner and choose **Storage > Elastic Volume Service**.
The **Elastic Volume Service** 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

8 Managing EVS Snapshots

8.1 EVS Snapshot Overview

Overview

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

NOTE

EVS is now deploying the snapshot function commercially in regions one by one. You may see the snapshot function in OBT (legacy snapshots) or commercial use (standard snapshots) in different regions, and there are differences between them. By default, a snapshot created in a commercially deployed region is a standard snapshot.

- Commercially deployed regions (standard snapshots): CN East2
- OBT regions (legacy snapshots): regions except CN East2

Table 8-1 Snapshot-related operations

Operation	Description	Reference
Creating snapshots	You can create a snapshot to save the disk data at a specified time. NOTE Snapshots are read-only. After snapshots are created, data in the snapshots cannot be modified.	Creating an EVS Snapshot
Rolling back data	If data on a disk is incorrect or damaged, you can roll back data from a snapshot to the source disk.	Rolling Back Disk Data from a Snapshot
Creating disks from a snapshot	You can create disks from a snapshot to quickly copy the snapshot data to disks.	Creating a Disk from a Snapshot

Operation	Description	Reference
Using Instant Snapshot Restore	The more data on an EVS disk, the more time it takes to create a standard snapshot. Instant Snapshot Restore allows you to create disks or roll back disk data from snapshots even if the snapshots are being created, and the restoration speed and creation speed are fast.	Enabling or Disabling Instant Snapshot Restore (for Snapshots in Commercial Use)
Checking snapshot information	You can check the storage used by all snapshots of an EVS disk, the total storage used by all snapshots in a specified period, and the total storage used by all snapshots of your account in a specified region. You can check the snapshot details, including the region and AZ, source disk information, and tags.	Checking the EVS Snapshot Storage Usage (for Snapshots in Commercial Use) Checking EVS Snapshot Details
Deleting snapshots	If you no longer require certain snapshots or the snapshot quantity reaches the maximum allowed, you can delete the snapshots.	Deleting an EVS Snapshot

Billing

Legacy snapshots are free. You can use them free of charge.

For the billing information about standard snapshots, see [Billing for EVS Snapshots](#).

Snapshot Usage Scenarios

The snapshot function helps address your following needs:

- Routine data backup
You can create snapshots for disks on a timely basis and use snapshots to recover your data in case that data loss or data inconsistency occurred due to unintended operations, viruses, or attacks.
- Rapid data restoration
You can create a snapshot or multiple snapshots before an application software upgrade or a service data migration. If an exception occurs during the upgrade or migration, service data can be rapidly restored to the time when the snapshot was created.

For example, a fault occurred on system disk A of server A, and therefore server A cannot be started. As system disk A is already faulty, data on system disk A cannot be restored by rolling back data from snapshots. However, you can create disk B using an existing snapshot of system disk A and attach disk

B to a properly running server, for example server B. In this case, server B obtains the data of system disk A from disk B.

NOTE

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

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

Snapshot Principles

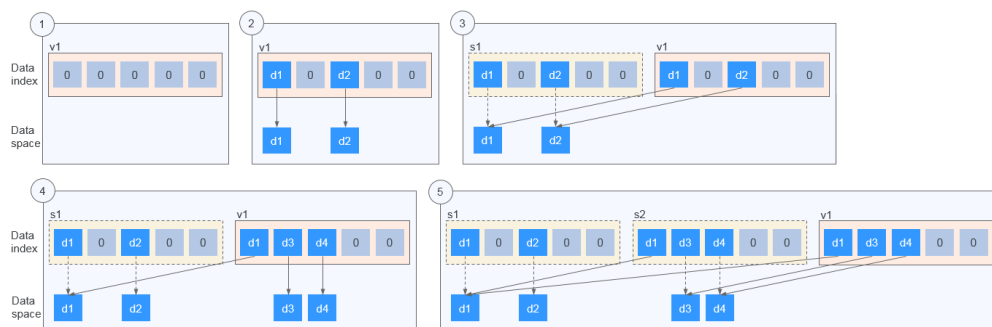
Legacy Snapshot Principles

Legacy snapshots and standard snapshots are different in that a standard snapshot stores data in OBS, while a legacy snapshot establishes a relationship between the snapshot and disk data. For details, see [Differences Between Disk Backups and Disk Snapshots](#)

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

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

Figure 8-1 Snapshot principles



Standard Snapshot Principles

Standard snapshots back up data by data block. They include **full snapshots** and **incremental snapshots**. The first snapshot created for an EVS disk is a full snapshot, which backs up all data blocks on the disk at the time of the snapshot. Subsequent snapshots are incremental snapshots, which back up only changed data blocks since the last snapshot.

Metadata files of full and incremental snapshots record information about all data blocks when the snapshots were created. So, you can use any snapshot to restore your disk data to the state when the snapshot was created.

Figure 8-2 Standard snapshot principles



Based on the source of data blocks, a snapshot's metadata file contains information about three types of data blocks: **inherited data blocks** (inherited from the last snapshot), **modified data blocks** (have modifications compared with the last snapshot), and **new data blocks** (new compared with the last snapshot).

A snapshot's data file stores only the changed data blocks (modified and new data blocks) compared with the last snapshot.

Let's use the preceding figure for illustration. Assume that data is written to an EVS disk at 09:30 and 10:30. Snapshot 1 is created at 09:00, snapshot 2 is created at 10:00, and snapshot 3 is created at 11:00.

- At 09:00, snapshot 1 is created for the disk. This is the first time that a snapshot is created for this disk, so snapshot 1 is a full snapshot and it contains all the data on the disk, including data blocks A, B, and C. The metadata file of snapshot 1 records information about the disk's full data blocks: A, B, and C.
- After snapshot 1 is created, data block A is changed to A1, data block B is changed to B1, and data block D is added. Then, snapshot 2 is created at 10:00. It is an incremental snapshot. Compared with snapshot 1, data blocks A1, B1, and D are changed data blocks. The metadata file of snapshot 2 records information about the disk's full data blocks: A1, B1, C, and D, among which data block C is inherited from snapshot 1.

- After snapshot 2 is created, data block A1 is changed to A2, data block C is changed to C1, and data block E is added. Then, snapshot 3 is created at 11:00. It is an incremental snapshot. Compared with snapshot 2, data blocks A2, C1, and E are changed data blocks. The metadata file of snapshot 3 records information about the disk's full data blocks: A2, B1, C1, D, and E, among which data blocks B1 and D are inherited from snapshot 2.

Calculating the Standard Snapshot Storage Usage

The total snapshot storage usage of an EVS disk is calculated by snapshot chain. A snapshot chain collects the storage space used by data blocks of all the snapshots of a disk.

- **Snapshot chain's storage usage calculation after snapshots are added**

Figure 8-3 Snapshot chain with snapshots added



Take the scenario in [Figure 8-3](#) as an example. Assume that the size of a snapshot's data block is fixed at 2 MiB. The snapshot chain's storage usage is calculated as follows:

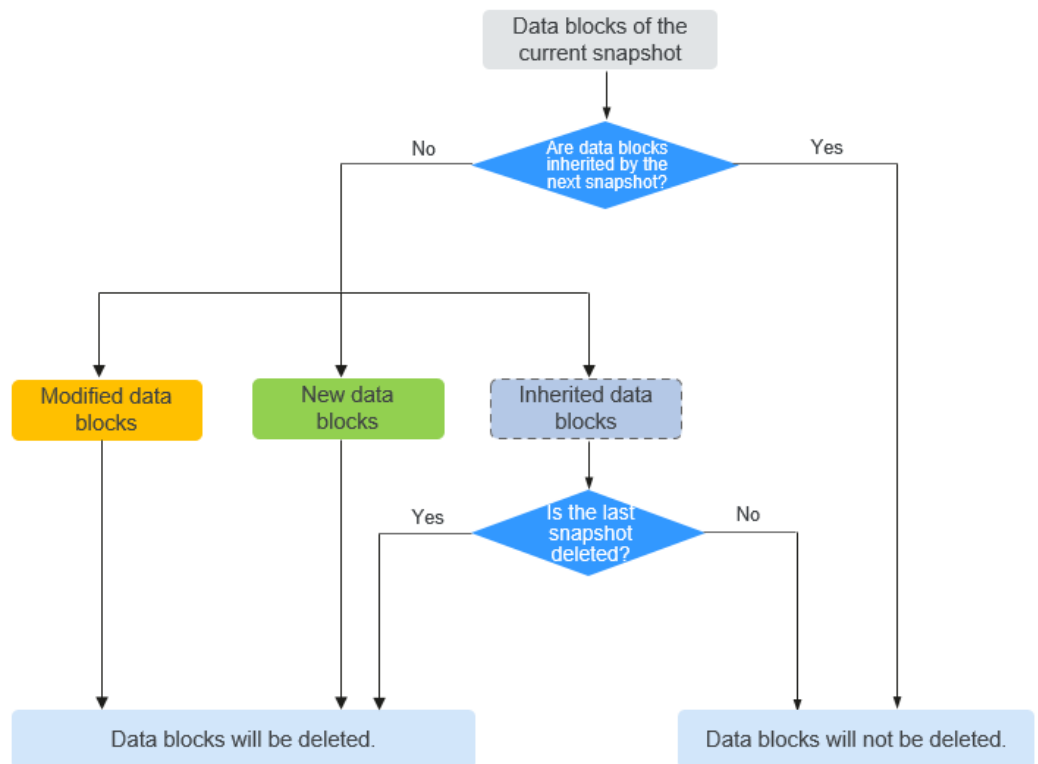
- After snapshot 1 is created, the snapshot chain of the disk contains only one snapshot. Snapshot chain's storage usage = Snapshot 1's storage usage = Size of data block A + Size of data block B + Size of data block C = 6 MiB
- After snapshot 2 is created, the snapshot chain of the disk contains two snapshots: snapshot 1 and snapshot 2. Snapshot chain's storage usage = Snapshot 1's storage usage + Snapshot 2's storage usage = 6 MiB + (Size of data block A1 + Size of data block B1 + Size of data block D) = 12 MiB
- After snapshot 3 is created, the snapshot chain of the disk contains three snapshots: snapshot 1, snapshot 2, and snapshot 3. Snapshot chain's storage usage = Snapshot 1's storage usage + Snapshot 2's storage usage + Snapshot 3's storage usage = 6 MiB + 6 MiB + (Size of data block A2 + Size of data block C1 + Size of data block E) = 18 MiB

- **Snapshot chain's storage usage calculation after snapshots are deleted**

When a snapshot is deleted, all data block information in this snapshot's metadata file is traversed, and the following deletion rules are applied:

- If a data block is inherited by the next snapshot, it will not be deleted.
- If a data block is not inherited by the next snapshot:
 - For an inherited data block, if the previous snapshot that the data block is inherited from is not deleted, the data block will not be deleted. Otherwise, it will be deleted.
 - For a modified data block, it will be deleted.
 - For a new data block, it will be deleted.

Figure 8-4 Snapshot data block deletion rules



The following example describes how to calculate a snapshot chain's storage usage after snapshots are deleted.

Figure 8-5 Snapshot chain with snapshots deleted



Take the scenario in [Figure 8-5](#) as an example. Assume that snapshot 2 is deleted at 14:00 and snapshot 3 is deleted at 15:00. The snapshot chain's storage usage is calculated as follows:

- Before any snapshot is deleted, the snapshot chain's storage usage is 18 MiB (Snapshot 1's storage usage + Snapshot 2's storage usage + Snapshot 3's storage usage).
 - When snapshot 2 is deleted at 14:00, information about all data blocks in the metadata file of snapshot 2 is traversed.
 - Data block A1: It is not inherited by snapshot 3 and is modified from data block A of snapshot 1. So, data block A1 will be deleted.
 - Data block B1: It is inherited by snapshot 3, so it will not be deleted.
 - Data block C: It is not inherited by snapshot 3, but is inherited from snapshot 1 and snapshot 1 is not deleted. So, data block C will not be deleted.
 - Data block D: It is inherited by snapshot 3. So, it will not be deleted.
- After snapshot 2 is deleted, the snapshot chain's storage usage is 16 MiB (18 MiB – Size of data block A1).
- When snapshot 3 is deleted at 15:00, information about all data blocks in the metadata file of snapshot 3 is traversed.
 - Data block A2: It is not inherited by the next snapshot and is modified from data block A1 of snapshot 2. So, data block A2 will be deleted.
 - Data block B1: It is not inherited by the next snapshot, but is inherited from snapshot 2 and snapshot 2 has been deleted. So, data block B1 will be deleted.
 - Data block C1: It is not inherited by the next snapshot and is modified from data block C of snapshot 2. So, data block C1 will be deleted.

- Data block D: It is not inherited by the next snapshot, but is inherited from snapshot 2 and snapshot 2 has been deleted. So, data block D will be deleted.
- Data block E: It is not inherited by the next snapshot and is newly added in snapshot 3. So, data block E will be deleted.

After snapshot 3 is deleted, the snapshot chain's storage usage is 6 MiB (16 MiB – Size of data block A2 – Size of data block B1 – Size of data block C1 – Size of data block D – Size of data block E).

EVS allows you to view the snapshot storage usage on the console. For details, see [Checking the EVS Snapshot Storage Usage \(for Snapshots in Commercial Use\)](#).

Differences Between Disk Backups and Disk Snapshots

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

Table 8-2 Differences between backups and snapshots

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

Item	Storage Solution	Data Synchronization	DR Range	Service Recovery
Legacy Snapshot	Snapshots are stored on the same disk as the source data. NOTE Creating a backup requires a certain amount of time because data needs to be transferred to OBS. Creating a snapshot or rolling back data from a snapshot consumes less time than creating a backup.	A snapshot is the state of a disk at a specific point in time and is stored on the same disk. If the disk is deleted, all its snapshots will also be deleted. For example, if you reinstalled or changed the server OS, snapshots of the system disk were also automatically deleted. Snapshots of the data disks can be used as usual.	A snapshot and its source disk reside in the same AZ.	You can use a snapshot to roll back data to its source disk or create a new disk.
Standard Snapshot	Standard snapshots are stored in OBS, instead of disks. This ensures data restoration upon disk damage or corruption.	A snapshot is the state of a disk at a specific point in time and is stored in OBS. If the disk is deleted, all its snapshots will not be deleted.	A snapshot and its source disk reside in different AZs.	You can use a snapshot to roll back data to its source disk or create a new disk. The data durability is high.

8.2 Using EVS Snapshots

8.2.1 Creating an EVS Snapshot

Scenarios

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

NOTE

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

Prerequisites

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

Notes and Constraints

- Snapshots can be created for both system disks and data disks.
- Snapshots of encrypted disks are stored encrypted, and those of non-encrypted disks are stored non-encrypted.

Snapshot function in OBT ([view supported regions](#))

- You can manually create a maximum of seven snapshots for a disk.
- Huawei Cloud reserves the right to restrict user snapshots created during OBT.
- The enterprise project of a snapshot is the same as that of the snapshot's source disk.

Snapshot function in commercial use ([view supported regions](#))


- You can manually create a maximum of 256 standard snapshots for a disk, of which up to seven can have Instant Snapshot Restore enabled.
- You can create one standard snapshot for a disk at a time. You can only create the next standard snapshot for the same disk after the previous snapshot has been created.
- Standard snapshots cannot be created for the disks in edge AZs. For details about the differences between edge AZs and general AZs, see [CloudPond User Guide](#).
- When standard snapshots are created for Common I/O and High I/O disks, Instant Snapshot Restore cannot be enabled.
- It usually takes several minutes to create a standard snapshot. The time required varies depending on the amounts of data written to the disk. The larger the data volume, the longer the time required. The initial standard snapshot usually takes more time because data of the entire disk is backed up. Subsequent standard snapshots are quicker, but the time required is still determined by the amounts of changed data compared with each last snapshot. The more the changed data, the longer the time required.
- If the data on a disk is rolled back from a snapshot, the next standard snapshot created for this disk will be a full snapshot.

- During the creation of a standard snapshot, any incremental data written to the disk will not be backed up to the snapshot created.
- During the creation of a standard snapshot, deleting the snapshot's source disk does not affect the creation of the snapshot.

Snapshot Function in OBT

Creating a Snapshot on the Disks Page

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

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

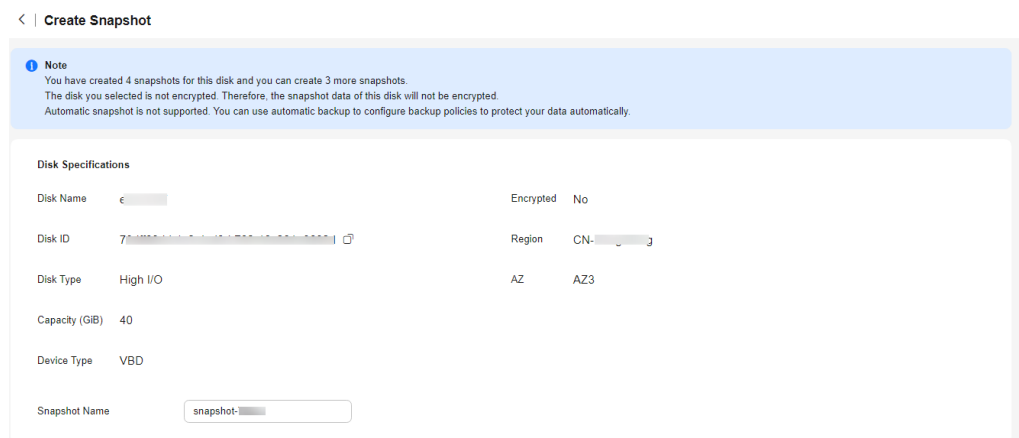
Step 3 In the disk list, locate the target disk and click **Create Snapshot** in the **Operation** column.

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

Table 8-3 Snapshot parameter

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

Figure 8-6 Create Snapshot



Step 4 Click **Create Now**.


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

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

----End

Creating a Snapshot on the Snapshots Page

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

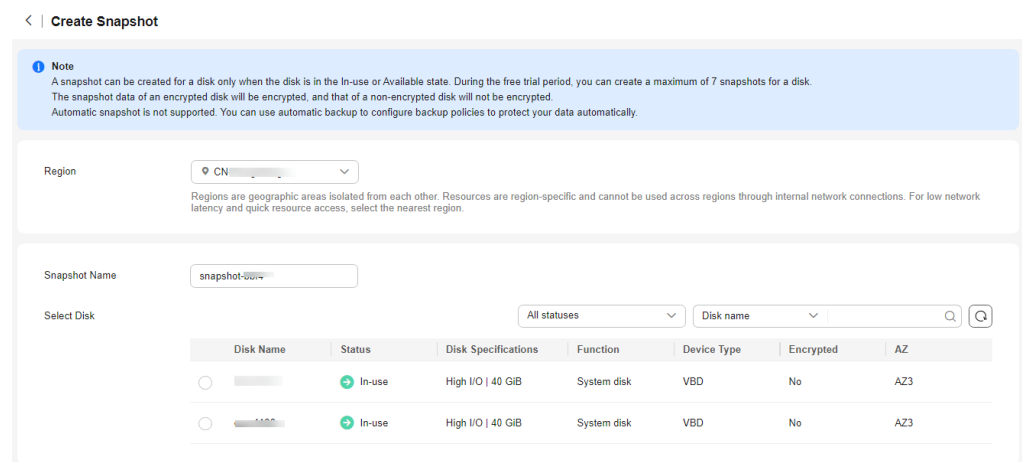
Step 2 Click  in the upper left corner and choose **Storage > Elastic Volume Service**.
The **Elastic Volume Service** page is displayed.

Step 3 In the navigation pane on the left, choose **Elastic Volume Service > Snapshots**.
On the **Snapshots** page, click **Create Snapshot**.
Configure the snapshot parameters according to [Table 8-4](#).

Table 8-4 Snapshot parameters

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

Figure 8-7 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 Function in Commercial Use

Creating a Snapshot on the Disks Page

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

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

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

Configure the snapshot parameters according to [Table 8-5](#).

Table 8-5 Snapshot parameters

Parameter	Description	Example Value
Region	The region to which the snapshot belongs. The snapshot must be in the same region as its source disk.	CN North-Beijing4
Disk Name/ID	The name and ID of the source disk.	-
Snapshot Name	Mandatory The name can contain a maximum of 64 characters.	snapshot-01Created_from_evstest
Snapshot Description	Optional The description can contain up to 255 characters.	-
Snapshot Type	The type of the snapshot. Only standard snapshot is supported currently. The time required for creating a standard snapshot depends on the size of data being backed up, but the initial snapshot usually takes a bit longer.	Standard snapshot

Parameter	Description	Example Value
Instant Snapshot Restore	Snapshots with Instant Snapshot Restore enabled allow you to create disks or roll back disk data even if the snapshots are being created, and the restoration speed and creation speed are fast. For more information, see Enabling or Disabling Instant Snapshot Restore (for Snapshots in Commercial Use) .	-
Advanced Settings	<p>Optional</p> <p>You can add tags when creating standard snapshots. Tags can help you to identify, classify, and search for your snapshots.</p> <p>NOTE</p> <ul style="list-style-type: none">You can add a maximum of 20 tags to a snapshot.Tag keys of the same snapshot must be unique. <p>A tag consists of a tag key and a tag value.</p> <ul style="list-style-type: none">A tag key cannot start or end with a space, or start with _sys_. It can contain a maximum of 128 characters and contain letters, digits, spaces, and the following special characters: <code>._:=-+@</code>A tag value can contain a maximum of 255 characters and contain letters, digits, spaces, and the following special characters: <code>._:/=-+@</code>	-
Enterprise Project	<p>Mandatory</p> <p>When creating a snapshot, you can add the snapshot to an existing enterprise project or a new one.</p> <p>An enterprise project facilitates project-level management and grouping of cloud resources and users. The default project is default.</p>	default

Figure 8-8 Create Snapshot

Note

1. Standard snapshots are now available for commercial use and will be billed. For billing details, see [Billing rules](#).
2. Snapshots can only be created for Available or In-use disks.
3. A maximum of 256 standard snapshots can be manually created for a disk, of which up to 7 can have Instant Snapshot Restore enabled.
4. Automatic snapshot is not supported. You can use automatic backup to configure backup policies to protect your data automatically.

Region

Regions are geographic areas isolated from each other. Resources are region-specific and cannot be used across regions through internal network connections. For low network latency and quick resource access, select the nearest region.

Disk Name/ID

Snapshot Name

Snapshot Description

Snapshot Type **Standard snapshot**

Snapshots are created in minutes, depending on the size of data being backed up, but the initial snapshot usually takes a bit longer. Additional pricing applies. [Billing rules](#). Standard snapshots already created for the selected disk: 1 Standard snapshots you can still create for this disk: 255

Instant Snapshot Restore Enable

Snapshots with Instant Snapshot Restore enabled allow you to create disks or roll back disk data even if the snapshots are being created.

Advanced Settings

Enterprise Project [Create Enterprise Project](#)

Step 4 Click **Create Now**.


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

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

----End

Creating a Snapshot on the Snapshots Page

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

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

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

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

Configure the snapshot parameters according to [Table 8-6](#).

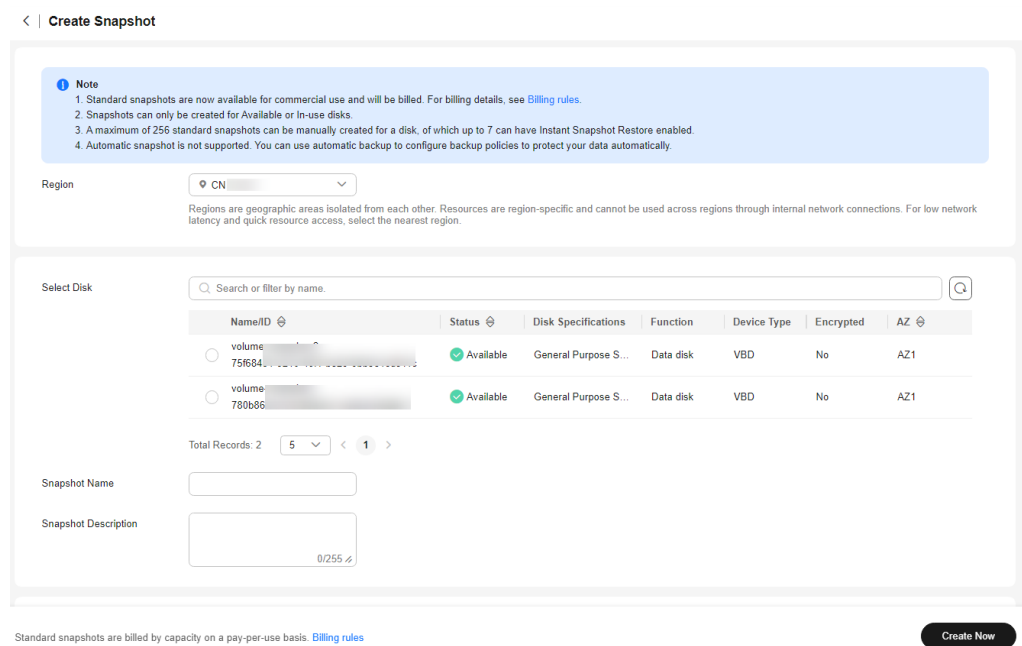
Table 8-6 Snapshot parameters

Parameter	Description	Example Value
Region	Mandatory After you select a region, disks in the selected region will be displayed for you to choose from.	CN North-Beijing4
Select Disk	Mandatory Select a disk for which you want to create a snapshot.	-

Parameter	Description	Example Value
Snapshot Name	Mandatory The name can contain a maximum of 64 characters.	snapshot-01Created_from_evstest
Snapshot Description	Optional The description can contain up to 255 characters.	-
Snapshot Type	The type of the snapshot. Only standard snapshot is supported currently.	Standard snapshot
Instant Snapshot Restore	Snapshots with Instant Snapshot Restore enabled allow you to create disks or roll back disk data even if the snapshots are being created, and the restoration speed and creation speed are fast. For more information, see Enabling or Disabling Instant Snapshot Restore (for Snapshots in Commercial Use) .	Enable
Advanced Settings > Tag	Optional You can add tags when creating standard snapshots. Tags can help you to identify, classify, and search for your snapshots. NOTE <ul style="list-style-type: none">You can add a maximum of 20 tags to a snapshot.Tag keys of the same snapshot must be unique. A tag consists of a tag key and a tag value. <ul style="list-style-type: none">A tag key cannot start or end with a space, or start with _sys_. It can contain a maximum of 128 characters and contain letters, digits, spaces, and the following special characters: <code>_:+=-@</code>A tag value can contain a maximum of 255 characters and contain letters, digits, spaces, and the following special characters: <code>_:./+=-@</code>	-

Parameter	Description	Example Value
Enterprise Project	<p>Mandatory</p> <p>When creating a snapshot, you can add the snapshot to an existing enterprise project or a new one.</p> <p>An enterprise project facilitates project-level management and grouping of cloud resources and users. The default project is default.</p>	default
Select Disk	<p>Mandatory</p> <p>Select a disk based on which the snapshot will be created.</p>	volume-01

Figure 8-9 Create Snapshot



Step 4 Click **Create Now**.

Step 5 Go back to the **Snapshots** page to view the snapshot creation information. After the snapshot status changes to **Available**, the snapshot has been created.

----End

8.2.2 Rolling Back Disk Data from a Snapshot

Scenarios


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

Notes and Constraints

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

Rolling Back Disk Data from a Snapshot

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

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

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

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

The **Snapshots** page is displayed.

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

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

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

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

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

----End

8.2.3 Creating a Disk from a Snapshot

Scenarios


You can create new disks from snapshots by either locating the target snapshot in the snapshot list and create a disk or specifying parameter **Create from snapshot** when creating a new disk.

Notes and Constraints

Snapshot Type	Notes and Constraints
Snapshot function in OBT (view supported regions)	<ul style="list-style-type: none"> Batch disk creation from a snapshot is not supported. A disk created from a snapshot has the same device type (SCSI or VBD), encryption attribute, AZ, region, and disk type as the snapshot's source disk. A snapshot whose name starts with autobk_snapshot_vbs_, manualbk_snapshot_vbs_, autobk_snapshot_csbs_, or manualbk_snapshot_csbs_ is automatically generated during backup. Such a snapshot can only be viewed. It cannot be used to create new disks.
Snapshot function in commercial use (view supported regions)	<ul style="list-style-type: none"> A standard snapshot with Instant Snapshot Restore not enabled can only be used to create disks when the snapshot status is Available. If Instant Snapshot Restore is enabled for a standard snapshot, when its upload is in progress, you can use it to create a disk but cannot change the device type (SCSI or VBD), encryption attribute, AZ, and type of the new disk. They are kept the same as those of the snapshot's source disk. A standard snapshot can only be used to batch create disks when its upload is complete. After a standard snapshot has been uploaded, you can change the device type (SCS or VBD), encryption attribute, AZ, or type of the disks when using this snapshot to create disks on the console. <p>NOTE You can view the snapshot upload progress in the status column. If there is a progress bar, the upload is still in progress. After the progress bar disappears, the upload is complete.</p>

Creating an EVS Disk from a Snapshot

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

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

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

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

The **Snapshots** page is displayed.

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

Step 5 Configure the disk parameters.

 **NOTE**

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

- In the condition that you do not specify a disk capacity, if the snapshot size is smaller than 10 GiB, the default capacity 10 GiB will be used as the disk capacity; if the snapshot size is greater than 10 GiB, the snapshot size will be used as the disk capacity.
- 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 Confirm the configuration and click **Submit**.

Step 8 Make the payment and click **OK**.

The disk list page is displayed.

Step 9 In the disk list, view the disk status.

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

----End

8.2.4 Enabling or Disabling Instant Snapshot Restore (for Snapshots in Commercial Use)

Scenarios

The more data on an EVS disk, the more time it takes to create a standard snapshot. Instant Snapshot Restore allows you to create disks or roll back disk data from snapshots even if the snapshots are being created, and the restoration speed and creation speed are fast.

Notes and Constraints


- Instant Snapshot Restore is supported for the following types of disks: Extreme SSD V2, Extreme SSD, General Purpose SSD V2, General Purpose SSD, and High I/O.
- You can only enable Instant Snapshot Restore when creating standard snapshots. It cannot be enabled later.
- You can enable Instant Snapshot Restore for up to seven snapshots for a disk.
- When Instant Snapshot Restore is enabled and snapshots are being created, you cannot disable Instant Snapshot Restore.
- When you delete a disk whose standard snapshots have Instant Snapshot Restore enabled, the snapshots will be not deleted, but Instant Snapshot Restore will be disabled automatically.

Enabling Instant Snapshot Restore

You can only enable Instant Snapshot Restore when creating standard snapshots. It cannot be enabled later. For details, see [Creating an EVS Snapshot](#).

Disabling Instant Snapshot Restore

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

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

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

Step 4 In the snapshot list, locate the target snapshot and click **Disable Instant Snapshot Restore** in the **Operation** column.

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

----End

8.2.5 Checking the EVS Snapshot Storage Usage (for Snapshots in Commercial Use)

Scenarios


You can check the storage used by all snapshots of an EVS disk, the total storage used by all snapshots in a specified period, and the total storage used by all snapshots of your account in a specified region.


Notes and Constraints

- The size of a single snapshot is smaller than the capacity of its source disk.
- A snapshot chain's storage usage may be greater than the capacity of the corresponding disk, because one disk may have multiple snapshots.

Checking the Total Storage Usage of All Snapshots of a Disk by Snapshot Chain


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

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

Step 3 Locate the disk that you want to check their total snapshot usage and click  to copy the disk ID.

Step 4 In the navigation pane on the left, choose **Elastic Volume Service > Snapshots**.
The **Snapshots** page is displayed.

Step 5 Click the **Snapshot Chains** tab.


Step 6 In the search box above the list, select Disk ID, paste the copied disk ID, and click .

- Step 7** View the capacity displayed in the **Snapshot Storage Usage** column.
 - Step 8** (Optional) Click the number displayed in the **Snapshots** column to view all the snapshots in the snapshot chain.
- End

Querying the Total Snapshot Storage Usage in a Specified Period

- Step 1** Log in to the [console](#).
- Step 2** Choose **Storage > Elastic Volume Service**.
- Step 3** In the navigation pane on the left, choose **Elastic Volume Service > Snapshots**.
The **Snapshots** page is displayed.
- Step 4** Click the **Snapshot Storage Usage** tab.
- Step 5** View the snapshot storage usage and snapshot quantity above the tabs.
- Step 6** Specify a time range for the query (minimum interval: 1 hour). You can also query the snapshot storage usage by the last 1 day, last 7 days, last 15 days, or last 30 days.

Figure 8-10 Querying the total snapshot storage usage in a specified period



----End

8.2.6 Checking EVS Snapshot Details

Scenarios

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

Checking Snapshot Details


- Step 1** Log in to the [console](#).
- Step 2** Choose **Storage > Elastic Volume Service**.
- Step 3** In the navigation pane on the left, choose **Elastic Volume Service > Snapshots**.
The **Snapshots** page is displayed.

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

If the snapshot function is in commercial use: Click the snapshot name to view the snapshot details on the **Basic Information** tab page.

If the snapshot function is in OBT: Select the snapshot and check the snapshot details at the bottom of the page.

Figure 8-11 OBT Snapshot details at the bottom of the page



----End

Snapshot Statuses

An EVS snapshot has several statuses. [Table 8-7](#) lists the EVS snapshot statuses, the meaning of each status, and the operations a snapshot in each status allows.

Table 8-7 Snapshot status details

Snapshot Status	Description	Allowed Operation
Creating	The snapshot is being created.	No operations are allowed.
Available	The snapshot is successfully created.	<ul style="list-style-type: none"> Creating EVS disks using snapshots Deleting snapshots Rolling back data to EVS disks using snapshots
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.

Snapshot Status	Description	Allowed Operation
Rolling back	The snapshot is rolling back data. NOTE <ul style="list-style-type: none">When you roll back from a snapshot, you can only roll back data to the source EVS disk. Rollback to a specified disk is not supported.A snapshot can only be used for rollback when its source disk is in the Available or Rollback failed state.	No operations are allowed.
Backing up	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. NOTE Temporary snapshots are created through the CBR service. Do not perform any operation on these snapshots.	No operations are allowed.

8.2.7 Deleting an EVS Snapshot

Scenarios

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

Prerequisites

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

Notes and Constraints

- If a snapshot is deleted, disks rolled back or created from this snapshot are not affected.

Snapshot function in OBT ([view supported regions](#))


- If a snapshot's source disk is deleted, all legacy snapshots of this disk are also deleted.
- If you reinstall or change the server OS, snapshots of the system disk are automatically deleted. Those of the data disks can be used as usual.
- A snapshot whose name starts with **autobk_snapshot_vbs_**, **manualbk_snapshot_vbs_**, **autobk_snapshot_csbs_**, or **manualbk_snapshot_csbs_** is automatically generated during backup. You can only check details of such snapshots and cannot delete them.

Snapshot function in commercial use ([view supported regions](#))

- Standard snapshots are not deleted even if their source disks are deleted.
- When you delete a disk whose standard snapshots have Instant Snapshot Restore enabled, the standard snapshots will be not deleted, but Instant Snapshot Restore will be disabled automatically.
- If you reinstall or change the server OS, standard snapshots will be not deleted, but Instant Snapshot Restore will be disabled automatically if it has been enabled for the standard snapshots of the system disk.

Procedure

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

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

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

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

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

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

----End

9 Managing Encrypted EVS Disks

What Is EVS Disk Encryption?

EVS enables you to encrypt data on newly created disks as required.

It uses the industry-standard XTS-AES-256 cryptographic algorithm and keys to encrypt EVS disks. Keys used to encrypt EVS disks are provided by the Key Management Service (KMS) of Data Encryption Workshop (DEW), which is secure and convenient. You do not need to establish and maintain the key management infrastructure. KMS uses the Hardware Security Module (HSM) that complies with FIPS 140-2 level 3 requirements to protect keys. All user keys are protected by the root key in HSM to prevent key exposure.

NOTICE

The encryption attribute of a disk cannot be changed after the disk is purchased. For details about how to create an encrypted disk, see [Purchasing an EVS Disk](#).

Keys Used for EVS Encryption

Keys provided by KMS include a Default Key and Custom Keys.

- **Default Key:** A key that is automatically created by EVS through KMS and named **evs/default**.
It cannot be disabled and does not support scheduled deletion.
- **Custom keys:** Keys created by users. You can use existing keys or create new ones to encrypt disks. For details, see "Key Management Service" > "Creating a CMK" in the *Data Encryption Workshop User Guide*.
- **Shared keys:** You can use DEW to create grants to share keys with other accounts. For details, see [Creating a Grant](#).

When an encrypted disk is attached, EVS accesses KMS, and KMS sends the data key (DK) to the host memory for use. The disk uses the DK plaintext to encrypt and decrypt disk I/Os. The DK plaintext is only stored in the memory of the host housing the ECS and is not stored persistently on the media. If a custom key is disabled or deleted in KMS, the disk encrypted using this custom key can still use

the DK plaintext stored in the host memory. If this disk is later detached, the DK plaintext will be deleted from the memory, and data can no longer be read from or written to the disk. Before you re-attach this encrypted disk, ensure that the custom key is enabled.

If you use a custom key to encrypt disks and this custom key is then disabled or scheduled for deletion, data cannot be read from or written to these disks or may never be restored. See [Table 9-1](#) for more information.

Table 9-1 Impact of custom key unavailability

Custom Key Status	Impact	How to Restore
Disabled	<ul style="list-style-type: none">For an encrypted disk already attached: Reads and writes to the disk are normal. If the disk is detached, it cannot be attached again.For an encrypted disk not attached: The disk cannot be attached anymore.	Enable the custom key. For details, see Enabling One or More Custom Keys .
Scheduled deletion		Cancel the scheduled deletion for the custom key. For details, see Canceling the Scheduled Deletion of One or More Custom Keys .
Deleted		Data on the disks can never be restored.

NOTICE

You will be billed for the custom keys you use. If pay-per-use keys are used, ensure that you have sufficient account balance. If yearly/monthly keys are used, renew your order timely. Or, your services may be interrupted and data may never be restored if encrypted disks become inaccessible.

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 you select when creating the server.

Table 9-2 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 9-3 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.

Purchased On	Method of Purchase	Whether Data Disk Will Be Encrypted	Description
	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.
	Creating from an image (The image's source disk is encrypted.)	Yes	-
	Creating from an image (The image's source disk is not encrypted.)	No	-

Notes and Constraints

Table 9-4 Constraints on disk encryption

Item	Description
Types of disks supporting encryption	All disk types
Constraints on encrypted disks	The encryption attribute of a disk cannot be changed after the disk is created, meaning that: <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	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. <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 EVS KMSAccess permissions to EVS. Then, the user can create and obtain keys to encrypt and decrypt disks.• If the user is not the first user, the user can use encryption directly.
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.

For details about how to create an encrypted disk, see [Purchasing an EVS 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 CMK status. For details, see [Keys Used for EVS Encryption](#).

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

10 Managing Shared EVS Disks

What Is Disk Sharing?

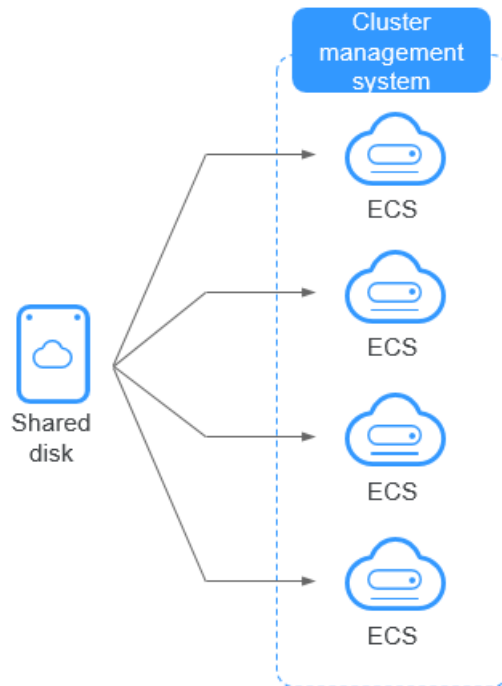
Disk sharing allows you to create shared EVS disks. Shared EVS disks are block storage devices that support concurrent read/write operations and can be attached to multiple servers. Shared EVS disks provide multiple attachments, high concurrency, high performance, and high reliability. They are usually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

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

NOTICE

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

Figure 10-1 Application scenario of shared EVS disks



Advantages

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

Specifications and Performance

Shared EVS disks have the same specifications and performance as non-shared EVS disks.

How Do I Use Shared VBD and SCSI Disks?

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

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

NOTICE

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

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

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

Constraints on Shared Disks

- A shared disk can be attached to a maximum of 16 servers.
- The sharing attribute of a disk cannot be changed after the disk is created.
- Shared disks can only be used as data disks, not system disks.
- A shared file system or cluster management system must be set up before you can properly use a shared disk. If you simply attach a shared disk to multiple servers, data cannot be shared between those servers and may be overwritten.
- When a shared disk is attached to multiple servers, the total performance of the disk on all servers cannot exceed the maximum allowed on a single disk.

Attaching a Shared EVS Disk

A non-shared EVS disk can only be attached to one server, whereas a shared EVS disk can be attached to up to 16 servers.

For details, see [Attaching a Shared Disk](#).

Deleting a Shared EVS Disk

Because a shared EVS disk can be attached to multiple servers, ensure that the shared EVS disk is detached from all the servers before deletion.

For details, see [Unsubscribing from or Deleting an EVS Disk](#).

Expanding a Shared EVS Disk

Shared EVS disks must be expanded when they are in the **Available** state. For details, see [Step 1: Expand Disk Capacity](#).

Data Sharing Principles and Common Usage Mistakes

A shared EVS disk is essentially the disk that can be attached to multiple servers for use. It is similar to a physical disk in that the disk can be attached to multiple physical servers, and each server can read data from and write data to any space on the disk. If no data read/write rules, such as the read/write sequence and meaning, between these servers are defined, data reads and writes between these servers may conflict, or other unpredictable errors may occur.

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

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

- Data inconsistency caused by read/write conflicts

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

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

- Data inconsistency caused by data caching

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

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

Before you buy a shared EVS disk, determine its device type (VBD or SCSI) based on the applications that will use the shared disk. Shared SCSI EVS disks support SCSI reservations. Before using SCSI reservations, you need to install a driver in the server OS and ensure that the OS image is included in the compatibility list.

NOTICE

If you simply attach a shared disk to multiple servers, data or files cannot be shared between the servers, because the shared disk does not have the cluster management capability. To share files between servers, build a shared file system or deploy a cluster management system.

Helpful Links

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

11 Managing EVS Disk Backups

11.1 CBR Overview

What Is CBR?

Cloud Backup and Recovery (CBR) enables you to easily back up cloud servers and cloud disks. In case of a virus attack, accidental deletion, or software or hardware fault, you can restore data to any point when the data was backed up.

CBR protects your workloads by ensuring the security and consistency of your data.

CBR Architecture

CBR involves backups, vaults, and policies.

- **Backup**

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

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

- **Vault**

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

Vaults can be either backup vaults or replication vaults. Backup vaults store resource backups, and replication vaults store backup replicas.

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

- **Policy**

There are backup policies and replication policies.

- A backup policy defines when you want to take a backup and for how long you would retain each backup.
- A replication policy defines when you want to replicate from backup vaults and for how long you would retain each replica. Backup replicas are stored in replication vaults.

Backup Mechanism

The first backup is a full backup and backs up all used data blocks.

For example, if a disk size is 100 GB and 40 GB has been used, only the 40 GB of data is backed up.

Subsequent backups are incremental backups. An incremental backup backs up only the data changed since the last backup to save the storage space and backup time.

When a backup is deleted, data blocks will not be deleted if they are depended on by other backups, ensuring that other backups can still be used for restoration. Both a full backup and an incremental backup can be used to restore data to a given backup point in time.

When creating a backup of a disk, CBR also creates a snapshot for it. CBR keeps only the latest snapshot. Every time it creates a new snapshot during backup, it deletes the old snapshot.

CBR stores backups in OBS to ensure data security.

Backup Options

CBR supports one-off backup and periodic backup. A one-off backup task is manually created by users and is executed only once. Periodic backup tasks are automatically executed based on a user-defined backup policy.

Table 11-1 One-off backup and periodic backup

Item	One-Off Backup	Periodic Backup
Backup policy	Not required	Required
Number of backup tasks	One manual backup task	Periodic tasks driven by a backup policy
Backup name	User-defined backup name, which is manualbk_XXXX by default	System-assigned backup name, which is autobk_XXXX by default
Backup mode	Full backup for the first time and incremental backup subsequently, by default	Full backup for the first time and incremental backup subsequently, by default

Item	One-Off Backup	Periodic Backup
Application scenario	Executed before patching or upgrading the OS or upgrading an application on a resource. A one-off backup can be used to restore the resource to the original state if the patching or upgrading fails.	Executed for routine maintenance of a resource. The latest backup can be used for restoration if an unexpected failure or data loss occurs.

11.2 Backing Up EVS Disks

Scenarios

EVS disk backups are created using the CBR service.

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

Notes and Constraints

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

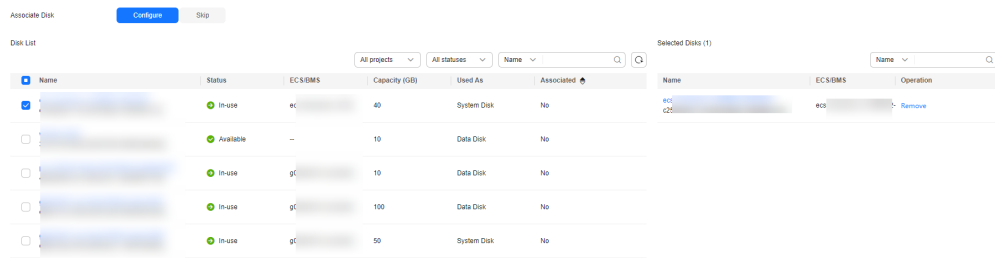
Step 2 In the upper right corner, click **Buy Disk Backup Vault**.

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 and you can buy or delete vaults at any time.

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 11-1 Selecting disks



NOTE

- Only **Available** and **In-use** disks can be selected.
- You can also associate disks with the vault you are creating later if you skip this step.

Step 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 policy. After the vault is created, CBR will apply the policy to this vault, and all disks associated with this vault will be automatically backed up based on this policy.
- If you select **Skip**, disks associated with this vault are not automatically backed up.

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

An enterprise project makes it easy to manage projects and groups of cloud resources and users. Use the **default** enterprise project or create one.

Step 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 11-2 describes the tag parameters.

Table 11-2 Tag parameters

Parameter	Description	Example Value
Key	<p>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.</p> <p>A tag key:</p> <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: <code>=*<>\ /</code> 	Key_0001

Parameter	Description	Example Value
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 and view the created vault in the vault list.

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

----End

12 Managing EVS Transfers

Scenarios

EVS transfer allows you to transfer disks from one account to another. After a transfer succeeds, the ownership of the disk belongs to the target account only.

Users can use disk transfer via API only. For more information, see [EVS Transfer](#).

Notes and Constraints

- Encrypted EVS disks cannot be transferred.
- EVS disks with backups and snapshots available cannot be transferred.
- EVS disks associated with backup policies cannot be transferred.

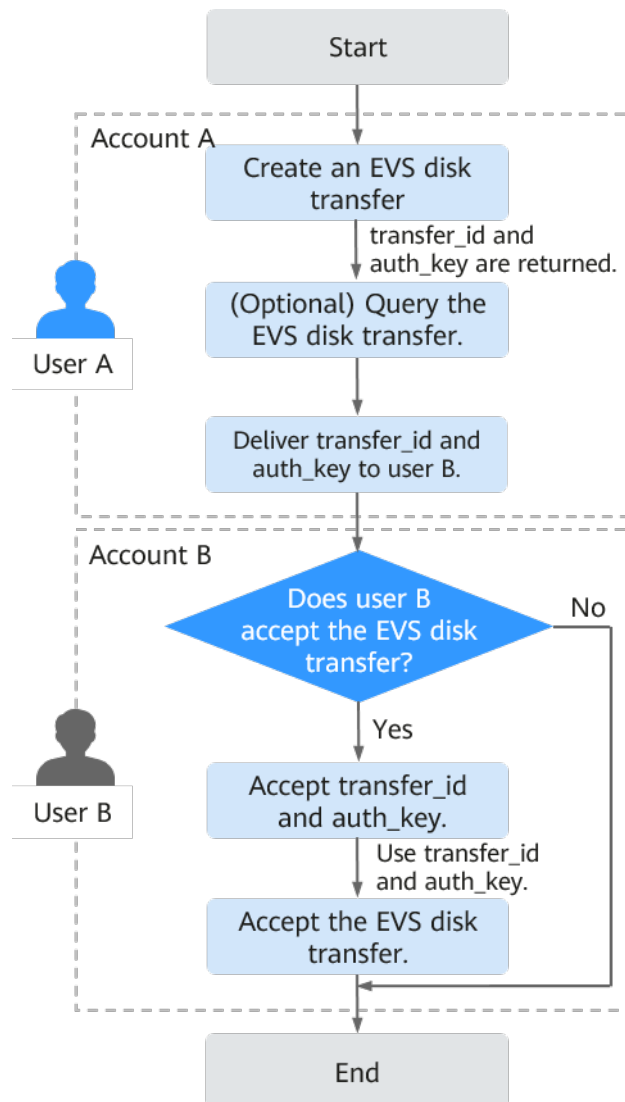
Procedure

The following example shows you how to transfer an EVS disk from account A to account B. User A belongs to account A, and user B belongs to account B. User A creates the transfer. User B accepts the transfer using the transfer ID (**transfer_id**) and authentication key (**auth_key**). After the transfer has been accepted, the transfer is complete. [Figure 12-1](#) shows the basic transfer process.

NOTE

- **transfer_id** specifies the disk transfer ID. Each EVS disk transfer has a transfer ID, and user B uses this ID to accept the disk transfer. The transfer ID expires after user B accepts the transfer.
- **auth_key** specifies the identity authentication key of the disk transfer. Each EVS disk transfer has an authentication key, and user B uses this key for authentication when accepting the disk transfer.

Figure 12-1 EVS disk transfer process



Step 1 User A creates 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

13 Managing EVS Tags

13.1 Tag Overview

Tags identify EVS resources for purposes of easy categorization and quick search.

If your organization has enabled the tag policy type for EVS and has a tag policy attached, you must comply with the tag policy rules when creating disks, otherwise disks may fail to be created. Contact the organization administrator to learn more about tag policies.

Table 13-1 Tag overview

Operation	Scenario
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.

13.2 Adding a Tag

Scenarios

You can add a tag for an existing EVS disk. You can also add tags when creating a disk.

Tag Rules

A tag consists of a tag key and a tag value. Tag rules are described as follows: (Tag rules vary depending on regions. See the rules displayed on the console.)


- First set of rules:
 - A tag key can contain a maximum of 36 characters. It can contain only letters, digits, special characters (.-_), and Unicode characters.
 - A tag value can contain a maximum of 43 characters. It can contain only letters, digits, special characters (.-_), and Unicode characters.
- Second set of rules:
 - A tag key can contain a maximum of 36 characters. It cannot contain special characters (=*<>\\,|/) or start or end with spaces.
 - A tag value can contain a maximum of 43 characters. It cannot contain special characters (=*<>\\,|/) or start or end with spaces.
- Third set of rules:
 - A tag key can contain a maximum of 128 characters. It cannot contain special characters (*<>\\,|), start with **_sys_**, or start or end with spaces.
 - A tag value can contain a maximum of 255 characters. It cannot contain special characters (*<>\\,|) or start or end with spaces.

Notes and Constraints

- A maximum of 20 tags can be added for an EVS disk.
- Tag keys of the same EVS disk must be unique.

Procedure

Step 1 Log in to the management console.

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

Step 3 In the disk list, locate the desired disk and click the disk name.
The disk details page is displayed.

Step 4 Click the **Tags** tab.

Step 5 Click **Add Tag**.
The **Add Tag** page is displayed.

Step 6 Enter a key and a value for a tag and click **OK**.

- **Key:** This parameter is mandatory.
- **Value:** This parameter is optional.

The **Tags** tab is displayed, and you can view the newly added tag.

----End

13.3 Modifying a Tag

Scenarios

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

Tag Rules

A tag consists of a tag key and a tag value. Tag rules are described as follows: (Tag rules vary depending on regions. See the rules displayed on the console.)

- First set of rules:
 - A tag key can contain a maximum of 36 characters. It can contain only letters, digits, special characters (.-_), and Unicode characters.
 - A tag value can contain a maximum of 43 characters. It can contain only letters, digits, special characters (.-_), and Unicode characters.
- Second set of rules:
 - A tag key can contain a maximum of 36 characters. It cannot contain special characters (=*<>\\,|/) or start or end with spaces.
 - A tag value can contain a maximum of 43 characters. It cannot contain special characters (=*<>\\,|/) or start or end with spaces.
- Third set of rules:
 - A tag key can contain a maximum of 128 characters. It cannot contain special characters (*<>\\,|), start with **_sys_**, or start or end with spaces.
 - A tag value can contain a maximum of 255 characters. It cannot contain special characters (*<>\\,|) or start or end with spaces.

Notes and Constraints

- A maximum of 20 tags can be added for an EVS disk.
- Tag keys of the same EVS disk must be unique.

Procedure

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

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

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

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

The disk details page is displayed.

Step 4 Click the **Tags** tab.

Step 5 Locate the target tag and click **Edit** in the **Operation** column.

The **Edit Tag** page is displayed.

Step 6 Change the value of the tag and click **OK**.

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

----End


13.4 Deleting a Tag

Scenarios

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

Procedure

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

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

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

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

The disk details page is displayed.

Step 4 Click the **Tags** tab.

Step 5 Locate the target tag and click **Delete** in the **Operation** column.

The **Delete Tag** page is displayed.

Step 6 Confirm the information and click **Yes**.

The tag is deleted if it disappears from the tag list.

----End


13.5 Searching for Disks by Tag

Scenarios

Tags can be used to categorize EVS disks, and users can quickly search for their desired EVS disks by tags. This section is used to guide users to search for EVS disk by existing tags.

Procedure

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

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

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

Step 3 In the search box, select a tag key under **Resource Tag** and then a tag value to trigger auto search.

You can search for disks by multiple tags and they are automatically joined with AND.

----End

14 Managing EVS Quotas

14.1 Querying EVS Resource Quotas

Scenarios

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

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

Procedure


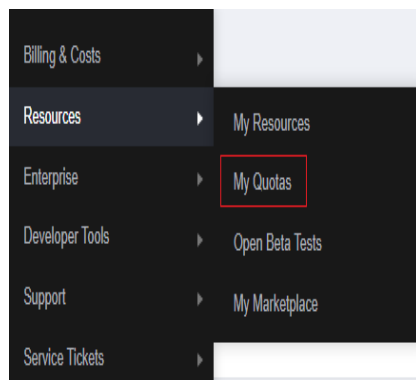
- Step 1** Log in to the [console](#).
- Step 2** Click  in the upper left corner and select the desired region and project.
- Step 3** In the upper right corner of the page, choose **Resources > My Quotas**.
The **Service Quota** page is displayed.

Figure 14-1 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

14.2 Increasing EVS Resource Quotas

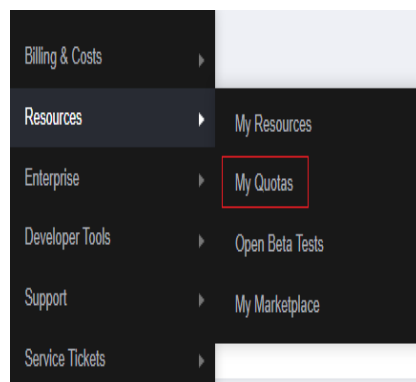
Scenarios

If any resource quota no longer meets your service requirements, you can apply for a higher quota.

How Do I Apply for a Higher Quota?

1. Log in to the management console.
2. In the upper right corner of the page, choose **Resources > My Quotas**.
The **Service Quota** page is displayed.

Figure 14-2 My Quotas



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

Figure 14-3 Increasing quota

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

15 Cloud Eye Monitoring

15.1 Viewing Basic EVS Monitoring Data

Description

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

Namespace

SYS.EVS

Metrics

Table 15-1 EVS metrics

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period
disk_device_read_bytes_rate	Disk Read Bandwidth	Number of bytes read from the monitored disk per second Unit: Bytes/s	≥ 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

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period
disk_device_read_requests_rate	Disk Read IOPS	Number of read requests sent to the monitored disk per second Unit: Requests/s	≥ 0 Requests/s	EVS disk	5 minutes in average
disk_device_write_requests_rate	Disk Write IOPS	Number of write requests sent to the monitored disk per second Unit: Requests/s	≥ 0 Requests/s	EVS disk	5 minutes in average
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

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period
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
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_upper_limit_reached_count	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_upper_limit_reached_count	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	<p><i>Server ID-drive letter</i>, for example, 6f3c6f91-4b24-4e1b-b7d1-a94ac1cb011d-vda (vda is the drive letter)</p> <p><i>Server ID-volume- Volume ID</i>, for example, 6f3c6f91-4b24-4e1b-b7d1-a94ac1cb011d-volume-31f45764-38b3-44ad-aaca-4015c83371e6</p>

Viewing Monitoring Data

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

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

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

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

The disk details page is displayed.

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

The **Monitoring Metrics** page is displayed.

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

----End

15.2 Viewing EVS Monitoring Data Included in OS Metrics (with Agent Installed)

Description

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

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

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

Monitoring Metrics

Table 15-2 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 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 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. <p>The path of the device name prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</p> Windows does not support this metric. 	0-100%	ECS	5 minutes in average

Dimensions

Key	Value
instance_id	Specifies the ECS ID.

16 Recording EVS Operations Using CTS

Scenarios

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

Prerequisites

CTS has been enabled.

Key EVS Operations Recorded by CTS

Table 16-1 EVS operations that can be recorded by CTS

Operation	Resource	Trace
Create disk	evs	createVolume
Update disk	evs	updateVolume
Expand disk capacity	evs	extendVolume
Delete disk	evs	deleteVolume
Create disk tag	evs	createVolumeTag

Viewing Traces

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