

# Elastic Volume Service(EVS)

## User Guide

**Issue** 01  
**Date** 2026-01-13



**Copyright © Huawei Technologies Co., Ltd. 2026. 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

---

<b>1 Using IAM to Grant Access to EVS.....</b>	<b>1</b>
1.1 Using IAM Roles or Policies to Grant Access to EVS.....	1
1.2 Using IAM Identity Policies to Grant Access to EVS.....	4
<b>2 Purchasing and Using an EVS Disk.....</b>	<b>8</b>
2.1 Overview.....	8
2.2 Purchasing an EVS Disk.....	9
2.3 Attaching an EVS Disk.....	18
2.3.1 Attaching a Non-Shared Disk.....	18
2.3.2 Attaching a Shared Disk.....	21
2.4 Initializing EVS Data Disks.....	24
2.4.1 Initialization Overview.....	24
2.4.2 Initializing a Linux Data Disk (Less Than or Equal to 2 TiB).....	27
2.4.3 Initializing a Linux Data Disk (Greater Than 2 TiB).....	34
2.4.4 Initializing a Windows Data Disk.....	40
<b>3 Expanding the EVS Disk Capacity.....</b>	<b>48</b>
3.1 Expansion Overview.....	48
3.2 Step 1: Expand Disk Capacity.....	51
3.3 Step 2: Extend Disk Partitions and File Systems.....	55
3.3.1 Extending Disk Partitions and File Systems (Linux).....	55
3.3.2 Extending Disk Partitions and File Systems (Windows).....	74
<b>4 Managing EVS Snapshots.....</b>	<b>89</b>
4.1 EVS Snapshot Overview.....	89
4.2 Using EVS Snapshots.....	101
4.2.1 Creating an EVS Snapshot.....	101
4.2.2 Rolling Back Disk Data from a Snapshot.....	110
4.2.3 Creating a Disk from a Snapshot.....	112
4.2.4 Enabling or Disabling Instant Snapshot Restore (for Standard Snapshots).....	114
4.2.5 Checking the EVS Snapshot Storage Usage (for Standard Snapshots).....	115
4.2.6 Checking EVS Snapshot Details.....	117
4.2.7 Deleting an EVS Snapshot.....	118
4.3 Using EVS Snapshot Consistency Groups.....	120
4.3.1 Creating a Snapshot Consistency Group.....	120

4.3.2 Rolling Back Disk Data Using a Snapshot Consistency Group.....	122
4.3.3 Deleting a Snapshot Consistency Group.....	123
<b>5 Changing the EVS Disk Type (OBT).....</b>	<b>125</b>
<b>6 Viewing EVS Disk Details.....</b>	<b>129</b>
<b>7 Detaching and Deleting an EVS Disk.....</b>	<b>133</b>
7.1 Detaching an EVS Disk.....	133
7.2 Unsubscribing from or Deleting an EVS Disk.....	136
<b>8 Managing EVS Recycle Bin.....</b>	<b>141</b>
8.1 Recycle Bin Overview.....	141
8.2 Enabling the Recycle Bin.....	144
8.3 Configuring a Recycle Bin Policy.....	146
8.4 Recovering Disks from the Recycle Bin.....	148
8.5 Permanently Deleting Disks from the Recycle Bin.....	149
8.6 Disabling the Recycle Bin.....	150
<b>9 Managing Encrypted EVS Disks.....</b>	<b>152</b>
9.1 EVS Encryption Overview.....	152
9.2 Creating an Encrypted EVS Disk.....	155
9.3 Configuring Default Encryption for EVS Disks (OBT).....	160
<b>10 Managing Shared EVS Disks.....</b>	<b>164</b>
<b>11 Managing EVS Disk Backups.....</b>	<b>169</b>
11.1 CBR Overview.....	169
11.2 Backing Up EVS Disks.....	171
<b>12 Managing EVS Transfers.....</b>	<b>173</b>
<b>13 Managing EVS Tags.....</b>	<b>176</b>
13.1 Tag Overview.....	176
13.2 Adding a Tag.....	176
13.3 Modifying a Tag.....	178
13.4 Deleting a Tag.....	179
13.5 Searching for Disks by Tag.....	179
<b>14 Managing EVS Quotas.....</b>	<b>181</b>
14.1 Querying EVS Resource Quotas.....	181
14.2 Increasing EVS Resource Quotas.....	182
<b>15 Cloud Eye Monitoring.....</b>	<b>184</b>
15.1 Basic EVS Monitoring Data.....	184
15.2 EVS Monitoring Data Included in OS Metrics (with Agent Installed).....	189
15.3 EVS Events Supported by Event Monitoring.....	210
15.4 Setting an Alarm Rule.....	210

---

**16 Recording EVS Operations Using CTS..... 212**

# 1 Using IAM to Grant Access to EVS

---

## 1.1 Using IAM Roles or Policies to Grant Access to EVS

System-defined permissions in [role/policy-based authorization](#) provided by [Identity and Access Management \(IAM\)](#) let you control access to 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 users only the permissions required to perform a given task based on their job responsibilities.
- Entrust a Huawei Cloud account or a cloud service to perform efficient O&M on your EVS resources.

If your Huawei Cloud account meets your permissions requirements, you can skip this section.

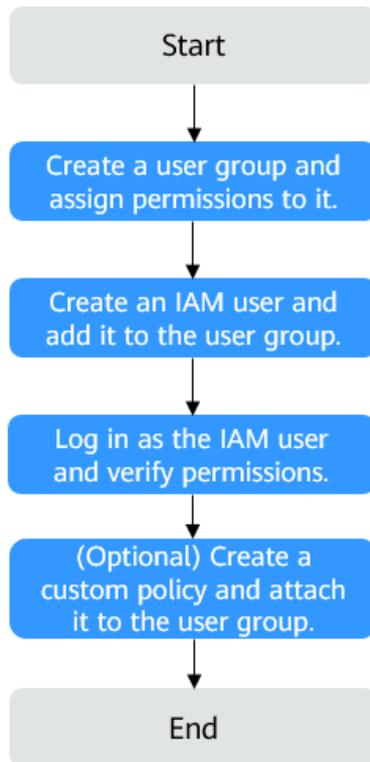
[Figure 1-1](#) shows the process flow of role/policy-based authorization.

### Prerequisites

Before granting permissions to user groups, learn about [system-defined permissions in role/policy-based authorization](#) for EVS. To grant permissions for other services, learn about all [system-defined permissions](#) supported by IAM.

## Process Flow

Figure 1-1 Process of 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.**
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.

## Example Custom Policies

You can create custom policies to supplement the system-defined policies of EVS.

To create a custom policy, choose either visual editor or JSON.

- Visual editor: Select cloud services, actions, resources, and request conditions. This does not require knowledge of policy syntax.
- JSON: Create a JSON policy or edit an existing one.

For details, see [Creating a Custom Policy](#). The following lists examples of common EVS custom policies.

- Example 1: Grant permission to create disks.

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Action": [
        "evs:volumes:create",
        "evs:volumes:list",
        "evs:volumes:get",
        "evs:types:get",
        "evs:quotas:get",
        "ecs:cloudServerFlavors:get",
        "ecs:cloudServers:list"
      ],
      "Effect": "Allow"
    }
  ]
}
```

- Example 2: Grant permission to deny disk deletion.

A policy with only "Deny" permissions must be used together with other policies. If the permissions granted to an IAM user contain both "Allow" and "Deny", the "Deny" permissions take precedence over the "Allow" permissions.

Assume that you want to grant the permissions of the **EVS FullAccess** policy to a user but want to prevent them from deleting EVS disks. You can create a custom policy for denying EVS disk deletion, and attach this policy together with the **EVS FullAccess** policy to the user. As an explicit deny in any policy overrides any allows, the user can perform all operations on EVS disks excepting deleting them.

Example policy denying disk deletion:

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

- Example 3: Create a custom policy containing multiple actions.

A custom policy can contain the actions of one or multiple services that are of the same type (global or project-level).

Example policy containing multiple actions:

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Action": [
        "evs:volumes:create",
        "evs:volumes:list",
        "evs:volumes:get",
        "evs:types:get",
        "evs:quotas:get",
        "evs:volumes:use"
      ],
      "Effect": "Allow"
    }
  ],
}
```

```
{
  "Action": [
    "ecs:cloudServerFlavors:get",
    "ecs:cloudServers:list",
    "ecs:cloudServers:get",
    "ecs:cloudServers:attach",
    "ecs:cloudServers:detachVolume"
  ],
  "Effect": "Allow"
}
```

## 1.2 Using IAM Identity Policies to Grant Access to EVS

System-defined permissions in [identity policy-based authorization](#) provided by [Identity and Access Management \(IAM\)](#) let you control access to EVS resources. With IAM, you can:

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

If your Huawei Cloud account meets your permissions requirements, you can skip this section.

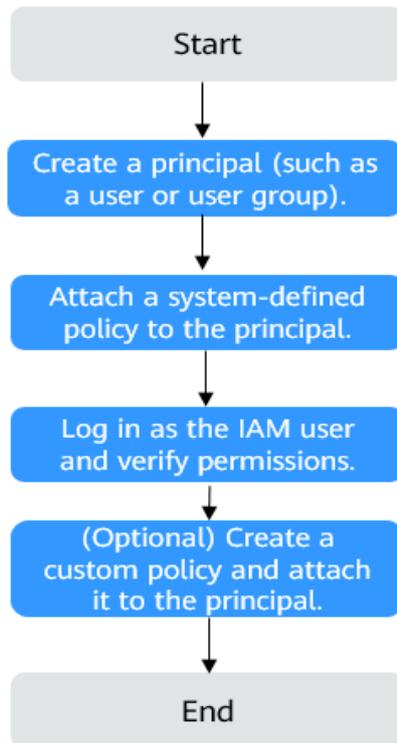
[Figure 1-2](#) shows the process flow of identity policy-based authorization.

### Prerequisites

Before granting permissions, learn about system-defined permissions in [identity policy-based authorization](#) for EVS. To grant permissions for other services, learn about all [system-defined permissions](#) supported by IAM.

## Process Flow

Figure 1-2 Process of granting EVS permissions



1. On the IAM console, **create an IAM user** or **create a user group**.
2. **Attach a system-defined identity policy** (**EVSReadOnlyPolicy** as an example) to the user or user group.
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 **EVSReadOnlyPolicy** 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 **EVSReadOnlyPolicy** policy is in effect.

## Example Custom Policies

- Example 1: Grant permissions to create disks.

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

```

```
        "evs:quotas:get",  
        "ecs:cloudServerFlavors:get",  
        "ecs:cloudServers:listServersDetails"  
    ]  
  }  
]  
}
```

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

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

```
{  
  "Version": "5.0",  
  "Statement": [  
    {  
      "Effect": "Allow",  
      "Action": [  
        "evs:volumes:create",  
        "evs:volumes:list"  
      ]  
    },  
    {  
      "Effect": "Allow",  
      "Action": [  
        "evs:volumes:create",  
        "evs:volumes:list",  
        "evs:volumes:get",  
        "evs:types:get",  
        "evs:quotas:get",  
        "evs:volumes:use"  
      ]  
    },  
    {  
      "Effect": "Allow",  
      "Action": [  
        "ecs:cloudServerFlavors:get",  
        "ecs:cloudServers:listServersDetails",  
        "ecs:cloudServers:showServer",  
        "ecs:cloudServers:attach",  
        "ecs:cloudServers:detachVolume"  
      ]  
    }  
  ]  
}
```

- 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

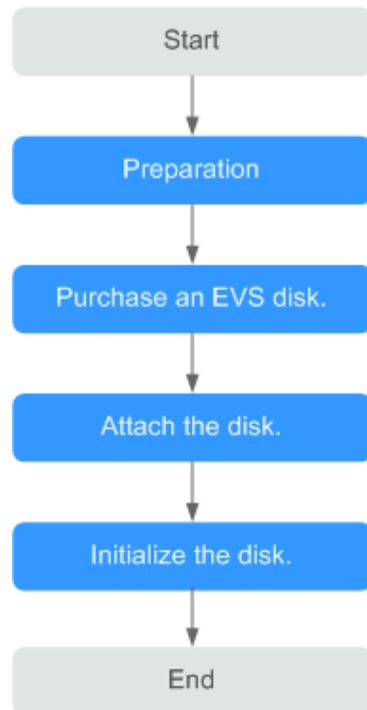
You can attach EVS disks to cloud servers to be used as system disks or data disks. For details, see [Table 2-1](#).

**Table 2-1** Method for purchasing disks

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

[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, capacity, 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 \(Greater Than 2 TiB\)](#)
  - [Initializing a Windows Data Disk](#)

## 2.2 Purchasing an EVS Disk

### Scenarios

You can use EVS disks as system disks or data disks for servers. If the storage space on your servers is insufficient, you can buy EVS disks to add more storage space. This section describes how to buy new EVS disks.

**Table 2-2** Scenarios

Buy On	Description	Reference
EVS console	You can only buy data disks on the EVS console. You can select to attach the disk to an ECS during the purchase or <b>manually attach</b> it after the purchase.	<a href="#">Buying Disks On the EVS Console</a>
Cloud server console	<ul style="list-style-type: none"><li>The system automatically attaches the data disks purchased together with servers or those added after the server purchase.</li><li>You can only buy system disks together with servers. After that, the system automatically attaches the system disks.</li></ul>	<a href="#">Purchasing and Using a Linux ECS</a>

## Constraints

**Table 2-3** Constraints

Category	Description
General constraints	<ul style="list-style-type: none"><li>Capacities of multiple disks cannot be combined, and the capacity of a single disk cannot be split.</li><li>You can only attach disks to servers in the same region and AZ. Once a disk is created, its region and AZ cannot be changed.</li><li>There are quantity and capacity quotas on EVS disks, so properly plan the number of disks and total disk capacity your workloads require. For details, see <a href="#">Managing EVS Quotas</a>.</li><li>After a disk is created, its device type, sharing attribute, and encryption attribute cannot be changed.</li><li>You can create a maximum of 100 disks in a batch.</li></ul>

Category	Description
Create from backup	<ul style="list-style-type: none"><li>• One backup cannot be used for concurrent disk creation operations at the same time. For example, if you are creating disk A from a backup, this backup can only be used to create another disk after disk A has been created.</li><li>• If a disk is created from a backup of a system disk, the new disk can only be used as a data disk.</li><li>• When you create a disk from a backup, 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 backup size will be used as the disk capacity.</li></ul>
Create from image	<ul style="list-style-type: none"><li>• The device type of the new disk is the same as that of the image's source disk.</li><li>• The encryption attribute of the new disk is the same as that of the image's source disk.</li></ul>
Create from snapshot	When you create a disk from a snapshot, the disk capacity must be greater than or equal to the snapshot size. In the condition that you do not specify a disk capacity, if the snapshot size is smaller than 10 GiB, the default capacity 10 GiB will be used as the disk capacity; if the snapshot size is greater than 10 GiB, the snapshot size will be used as the disk capacity.

## Billing

- EVS disks are billed based on the disk type, capacity, and usage duration. For details, see [Billing for EVS Disks](#).
- If you create disks together with a server, the disks will have the same billing mode as their server.
- Yearly/Monthly disks that you add to a cloud server after the server purchase have a different expiration time as the server. Such disks cannot be automatically renewed or unsubscribed from together with the server. If they expire earlier, services on the server may be interrupted.
- If KMS encryption is used, what you use beyond the free quota given by KMS will be billed. For details, see [DEW Billing](#).

## Buying Disks On the EVS Console

**Step 1** Sign in to the [EVS console](#).

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

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

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

**Table 2-4** Parameter description

Parameter	Example Value	Description
Attach to Server	Now	<ul style="list-style-type: none"> <li>● <b>Now:</b> 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.</li> <li>● <b>Later:</b> When no server is available, you can select this option to create the disk first and attach the disk after the purchase.</li> </ul> <p><b>NOTE</b> This parameter is available only in some regions. Whether it is displayed depends on the region where you use EVS.</p>
Region	-	Resources in different regions cannot communicate with each other over an internal network. For low network latency and quick resource access, select the nearest region.
AZ	AZ1	The availability zone (AZ) where you want to create the disk.
Billing Mode	Pay-per-use	<p>You can pay for EVS disks in two ways:</p> <ul style="list-style-type: none"> <li>● Yearly/Monthly</li> <li>● Pay-per-use</li> </ul> <p><b>NOTICE</b></p> <ul style="list-style-type: none"> <li>● Selecting <b>Now</b> for <b>Attach to Server</b>: <ul style="list-style-type: none"> <li>- 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 <b>Later</b> for <b>Attach to Server</b>, buy a pay-per-use disk, and attach it to the yearly/monthly server after the purchase.</li> <li>- 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 <b>Later</b> for <b>Attach to Server</b>, buy a yearly/monthly disk, and attach it to the pay-per-use server after the purchase.</li> </ul> </li> <li>● Selecting <b>Later</b> for <b>Attach to Server</b>: 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.</li> </ul>

Parameter	Example Value	Description
Data Source	-	<p>Optional</p> <p>Supported data source types include backup, snapshot, and image.</p> <ul style="list-style-type: none"> <li>• <b>Backup:</b> The disk will be created from a backup. Choose <b>Backup</b> &gt; <b>Select Data Source</b>. On the displayed page, select a backup and click <b>OK</b>.</li> <li>• <b>Snapshot:</b> The disk will be created from a snapshot. Choose <b>Snapshot</b> &gt; <b>Select Data Source</b>. On the displayed page, select a snapshot and click <b>OK</b>.</li> </ul> <p><b>NOTE</b> For more information about creating disks from snapshots, see <a href="#">Creating a Disk from a Snapshot</a>.</p> <ul style="list-style-type: none"> <li>• <b>Image:</b> The disk will be created from an image. Choose <b>Image</b> &gt; <b>Select Data Source</b>. On the displayed page, select an image and click <b>OK</b>.</li> </ul>
Disk Type	Ultra-high I/O	<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 <a href="#">Disk Types and Performance</a>.</p> <p><b>NOTE</b> General Purpose SSD V2 disks allow you to specify the disk IOPS and throughput.</p> <p>You can change the disk type after a disk is created. For details, see <a href="#">Changing the EVS Disk Type (OBT)</a>.</p>
Capacity (GiB)	100 GiB	<p>Set a disk capacity. You can only create data disks on the current page. The disk capacity ranges from 10 GiB to 32,768 GiB.</p> <p><b>NOTE</b> The system shows you the remaining disk space you can purchase. If the disk capacity you need exceeds the upper limit, click <b>Increase Quota</b> to obtain a higher quota. You can purchase the disk capacity you need after the request is approved.</p>

Parameter	Example Value	Description
Disk Encryption	-	<p>You can only encrypt data disks here, and you need to create an agency to authorize EVS to access KMS. After the authorization is successful, configure the following parameters on the <b>Encryption Setting</b> page displayed:</p> <ul style="list-style-type: none"> <li> <p>Select an existing key If you select <b>Select an existing key</b>, select a key from the drop-down list. You can select one of the following keys:</p> <p>Default keys: After KMS access permissions are granted to EVS, the system automatically creates a default key and names it <b>evs/default</b>.</p> <p>Custom keys: You can choose an existing key or create a new one. For details about how to create a key, see <a href="#">Creating a Custom Key</a>.</p> </li> <li> <p>Enter a key ID If you select <b>Enter a key ID</b>, 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 how to share keys, see <a href="#">Creating a Grant for a Custom Key</a>. Keys can only be shared with accounts, not users.</p> </li> <li> <p>Data Encryption Algorithm Supported encryption algorithms are AES256 and SM4.</p> </li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>System disk encryption depends on images. For details, see <a href="#">Encrypting Images</a>.</li> <li>Before using encryption, you need to create an agency to grant KMS access permissions to EVS. If you have the right to grant the permission, grant KMS access permissions to EVS directly. After KMS access permissions are granted, follow-up operations do not require the permissions to be granted again. If you do not have this permission, contact a user with the security administrator permissions to grant KMS access permissions to EVS, then repeat the preceding operations.</li> </ul>

Parameter	Example Value	Description
Advanced Settings <ul style="list-style-type: none"> <li>• Share</li> <li>• SCSI</li> </ul>	-	<ul style="list-style-type: none"> <li>• <b>Share</b> If you select <b>Share</b>, a shared disk will be created. A shared disk can be attached to up to 16 servers. If you do not select <b>Share</b>, a non-shared disk will be created, and the disk can be attached to one server only.  If you select both <b>Share</b> and <b>SCSI</b>, a shared SCSI disk will be created.</li> <li><b>NOTE</b> For details about shared EVS disks, see <a href="#">Managing Shared EVS Disks</a>.</li> <li>• <b>SCSI</b> If you select <b>SCSI</b>, a SCSI disk will be created. Such disks allow the server OS to directly access the underlying storage media and send SCSI commands to the disks. If you do not select <b>SCSI</b>, a VBD disk will be created. That said, the disk device type is VBD, the default device type.</li> <li><b>NOTE</b> To learn more about the ECS types and OSs supported by SCSI disks and the requirements on the ECS software, see <a href="#">Device Types</a>.</li> </ul>

Parameter	Example Value	Description
Automatic Backup	-	<p>Optional</p> <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 a backup vault and apply the selected policy to the vault to periodically back up the disk.</p> <ul style="list-style-type: none"> <li>• Skip this configuration if backup is not required. If you need backup protection after a disk has been purchased, go to the CBR console, locate the desired vault, and associate the disk with the vault.</li> <li>• <b>Use existing:</b> <ol style="list-style-type: none"> <li>1. <b>Vault:</b> Select an existing vault from the drop-down list.</li> <li>2. <b>Backup Policy:</b> Select a backup policy from the drop-down list, or go to the CBR console and configure a desired one.</li> </ol> </li> <li>• <b>Buy new:</b> <ol style="list-style-type: none"> <li>1. Enter a vault name, which can contain a maximum of 64 characters, including letters, digits, underscores (_), and hyphens (-), for example, <b>vault-f61e</b>. The default naming rule is <b>vault_XXXX</b>.</li> <li>2. Enter a vault capacity for storing disk backups. The vault capacity cannot be less than the size of the disk to be backed up. The value ranges from the disk size to 10,485,760 in the unit of GiB.</li> <li>3. Select a backup policy from the drop-down list, or go to the CBR console and configure a desired one.</li> </ol> </li> </ul>
Disk Name	<p>Assume that you create two disks and enter <b>volume for Disk Name</b>, the EVS disk names will be <b>volume-0001</b> and <b>volume-0002</b>.</p>	<ul style="list-style-type: none"> <li>• If you create a single disk, the name you entered will be used as the disk name. The value can contain a maximum of 64 characters.</li> <li>• If you create multiple disks in a batch, the name you entered will be used as the prefix of disk names. Each disk name will be composed of the name you entered and a four-digit number. The value can contain a maximum of 59 characters.</li> </ul>

Parameter	Example Value	Description
Tag	-	<p>Optional</p> <p>You can add tags when creating disks. Tags can help you identify, classify, and search for your disks. For details about tag rules, see <a href="#">Adding a Tag</a>.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• Tag rules vary depending on regions. See the rules displayed on the console.</li> <li>• Except for tagging the disk during disk creation, you can also add, modify, or delete tags for existing disks. For more information about tags, see <a href="#">Managing EVS Tags</a>.</li> </ul>
Enterprise Project	default	<p>When creating EVS disks, you can add the disks to an existing enterprise project.</p> <p>An enterprise project facilitates project-level management and grouping of cloud resources and users. The default project is <b>default</b>.</p>
Quantity	<p>Usage Duration: 1 year</p> <p>Select <b>Auto-renew</b>.</p> <p>Quantity: 1</p>	<ul style="list-style-type: none"> <li>• <b>Usage Duration:</b> This parameter is mandatory if you select <b>Yearly/Monthly</b> for <b>Billing Mode</b>. You can choose from 1 month to 3 years.</li> <li>• <b>Auto-renew:</b> This parameter is optional. Your order will automatically renew on a monthly or yearly basis, depending on if you purchased by month or by year.</li> <li>• <b>Quantity:</b> This parameter is optional. The preset disk quantity is <b>1</b>. It means that only one disk will be created. You can create a maximum of 100 disks at a time.</li> </ul> <p><b>NOTE</b></p> <p>The system shows you the remaining number of disks you can purchase. If the number of disks you need exceeds the upper limit, click <b>Increase Quota</b> to obtain a higher quota. You can purchase the disks you need after the request is approved.</p>

**Step 5** Check your configurations on the configuration summary page. If the information is correct, click **Submit**.

- If you select **Yearly/Monthly** for **Billing Mode**:
  - a. On the displayed page, select a desired payment method and confirm the payment. The system displays a message indicating payment processed successfully.
  - b. Click **Go to Elastic Volume Service Console** to return to the **Elastic Volume Service** page.
- If you select **Pay-per-use** for **Billing Mode**:
 

On the page with message "Task submitted successfully" displayed, click **Go to Disk List** to return to the **Elastic Volume Service** page.

**Step 6** In the disk list, view the disk status.

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

**Step 7** Perform follow-up operations. After new disks are purchased, attach and initialize them based on different scenarios.

**Table 2-5** Follow-up operations

Step	Description
Step 1: Attach the disk.	If you choose not to attach the disk when purchasing the disk, you need to manually <b>attach</b> it later.
Step 2: Initialize the disk.	The procedure for initializing a newly created empty data disk differs from that for a data disk with data on it. For details, see <b>Initialization Overview</b> .

----End

## Related Links

- To store data on EVS disks, upload files to ECSs by referring to [How Do I Upload Files to My ECS?](#)
- You can also create EVS disks through API calls. For details, see [Creating EVS Disks](#).
- To create EVS disks from snapshots, see [Creating a Disk from a Snapshot](#).

## 2.3 Attaching an EVS Disk

### 2.3.1 Attaching a Non-Shared Disk

#### Scenarios

You need to attach EVS disks to servers to store data. A non-shared disk can only be attached to one server. This section describes how to attach a non-shared 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, the disk function changes to **Bootable disk**, and the status changes to **Available**. You can attach a bootable disk to a server to be used as a system disk or data disk depending on the disk function selected.

#### Prerequisites

- The account is not in arrears.

- At least **one disk has been purchased or created**.
- The status of the non-shared disk is **Available**.
- To attach a disk as a data disk, the status of the server must be **Running** or **Stopped**.
- To attach a disk as a system disk, the status of the server must be **Stopped**.

## 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 non-shared disk and the server must be in the same region and AZ.
- A frozen disk cannot be attached.
- A disk can only be successfully attached when its device type is supported by the desired ECS. For details, see **Device Types Supported by ECS**.
- To re-attach an existing disk, ensure that the following conditions are met.

How Disks Are Purchased	Re-attaching Requirements
Non-shared data disks purchased together with yearly/monthly servers	They can only be re-attached to the original servers as data disks.
System disks purchased together with yearly/monthly servers	<ul style="list-style-type: none"> <li>• To re-attach and use them as system disks, you can only attach them to the original servers.</li> <li>• To re-attach and use them as data disks, you can attach them to any desired servers.</li> </ul>
System disk purchased with pay-per-use servers	<ul style="list-style-type: none"> <li>• To re-attach and use them as system disks, you can only attach them to servers that use the same images as the original servers.</li> <li>• To re-attach and use them as data disks, you can attach them to any desired servers.</li> </ul>

## Attaching the Disk on the EVS Console

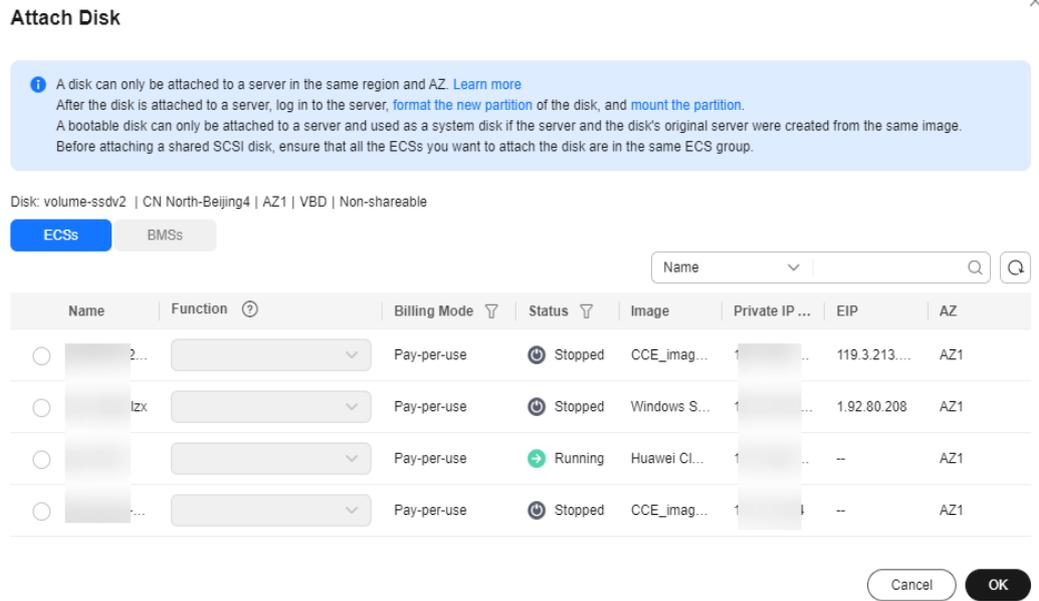
**Step 1** Sign in to the **EVS console**.

**Step 2** In the disk list, locate the disk and click **Attach**.

**Step 3** Select a 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 "Elastic Cloud Server FAQs" > "How Do I Obtain My Disk Name in the ECS OS Using the Device Identifier Provided on the Console?"

**Figure 2-2 Attach Disk**



**Step 4** 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 5** 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. Sign in to the [ECS console](#).
2. In the search box above the ECS list, enter the ECS name, IP address, or ID for search.
3. Click the name of the target ECS.  
The ECS details page is displayed.
4. Click the **Disks** tab. Then, click **Attach Disk**.  
The **Attach Disk** dialog box is displayed.

**Figure 2-3 Attach Disk (KVM)**

5. 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**

- For the restrictions on attaching disks, see [What Are the Requirements for Attaching an EVS Disk to an ECS?](#)
6. Click **OK**. A dialog box is displayed, showing "The disk has been attached but still needs to be initialized before it can be used".  
After the disk is attached, you can view information about it on the **Disks** tab.

## Follow-Up Operations

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

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

```
mount <disk-partition> <mount-point>
```

## Related 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?](#)
- If you no longer need an EVS disk or want to attach it to a different server in the same AZ, detach the disk by referring to [Detaching an EVS Disk](#).

## 2.3.2 Attaching a Shared Disk

### Scenarios

In cloud computing environments, enterprise servers often need to concurrently read or write the same disk for efficient data sharing and fast failover. You can

attach a shared EVS disk to multiple servers to achieve this goal. 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.
- At least **one disk has been purchased or created**.

## 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 region and AZ.
  - A shared, **In-use** disk can only be attached to servers when the maximum number of allowed servers has not been reached.
  - A shared disk can only be attached to servers running the same type of OS (either Windows or Linux).  
For example, if you attach a shared disk to multiple Windows servers and then detach it, the shared disk cannot be attached to Linux servers later. This is because Windows and Linux support different file systems. Improper operations may damage the original file system.
  - A shared disk can only be used as a data disk. It cannot be used as a system disk.
  - Cloud servers created from ISO images are only used for OS installation. They have limited functions and cannot have EVS disks attached.
  - A frozen disk cannot be attached.

## Attaching the Disk on the EVS Console

**Step 1** Sign in to the [EVS console](#).

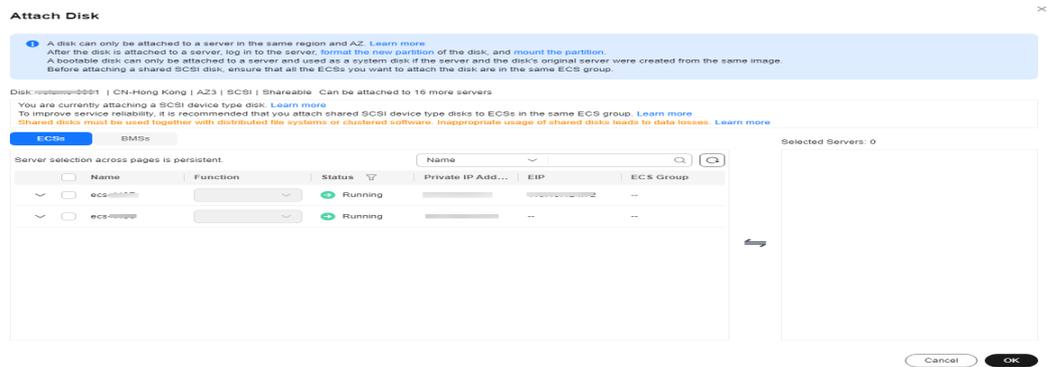
**Step 2** In the disk list, locate the disk and click **Attach**.

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

**Step 3** Select the target servers to attach the shared disk. Ensure that the disk and servers are in the same AZ. After you select servers, the system automatically inputs **Data disk** as the disk function.

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 4** 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 5** 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. Sign in to the [ECS console](#).
2. In the search box above the upper right corner of the ECS list, enter the ECS name, IP address, or ID for search.
3. Click the name of the target ECS.  
The ECS details page is displayed.
4. Click the **Disks** tab. Then, click **Attach Disk**.  
The **Attach Disk** page is displayed.
5. Select the target disk and specify it as a data disk.
  - For Xen ECSs, you can specify a specific device name, such as **/dev/sdb**.
  - For KVM ECSs, you cannot specify a specific device name.

### NOTE

- For the restrictions on attaching disks, see [What Are the Requirements for Attaching an EVS Disk to an ECS?](#)
6. Click **OK**. A dialog box is displayed, showing "The disk has been attached but still needs to be initialized before it can be used".

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

## Follow-Up Operations

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

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

```
mount <disk-partition> <mount-point>
```

## Related 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?](#)
- If you no longer need an EVS disk or want to attach it to a different server in the same AZ, detach the disk by referring to [Detaching an EVS Disk](#).

# 2.4 Initializing EVS Data Disks

## 2.4.1 Initialization Overview

### Scenarios

- **System disk**  
When a server is created, a system disk is automatically initialized with Master Boot Record (MBR).
- **Newly created empty data disk**  
After you attach a newly created empty data disk to a server, you must initialize the disk before using it. The initialization includes creating partitions, creating file systems, and mounting partitions.
  - 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 separately, you need to first attach it to a server and then initialize it.For detailed operation instructions, see [Table 2-6](#).
- **Existing data disk**  
An existing data disk is a disk created from a data source (snapshot, backup, or image) or one detached from another server.
  - You can choose not to re-partition the disk, but use the disk existing partitions.

- In Linux, **create new mount points and mount the partitions**, 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 2-6](#).

## Impact on the System

- An initialization operation includes partitioning, which deletes all the data on the disk.
- If you change the partition style of a disk, data on the disk will be erased. Select an appropriate partition style when initializing disks.
- 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

Major initialization steps include **creating partitions and file systems, mounting partitions, and configuring auto mount at system startup**. In Linux, you can choose different partition styles based on your disk capacity. For details, see [Table 2-7](#).

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 plan to expand it to over 2 TiB later, use GPT when initializing disks.

**Table 2-6** Disk initialization instructions

Disk Capacity	Partition Style	OS	Reference
Capacity ≤ 2 TiB	Guid Partition Table (GPT) or MBR	Linux	<a href="#">Initializing a Linux Data Disk (Less Than or Equal to 2 TiB)</a>
		Windows	<a href="#">Initializing a Windows Data Disk</a>
Capacity > 2 TiB	GPT	Linux	<a href="#">Initializing a Linux Data Disk (Greater Than 2 TiB)</a>

Disk Capacity	Partition Style	OS	Reference
		Windows	<a href="#">Initializing a Windows Data Disk</a>

## Partition Styles

Common disk partition styles include MBR and GPT. In Linux, you can choose different partition styles based on your disk capacity, as described in [Partition Styles](#).

**Table 2-7** Partition styles

Partition Style	Max. Disk Size Supported	Max. Number of Partitions Supported	Linux Partitioning Tool	Reference
MBR	2 TiB	<p>MBR partitions include primary partitions and extended partitions. A maximum of four primary partitions are supported. If you need more partitions, create one extended partition. The following is an example:</p> <ul style="list-style-type: none"> <li>• Four primary partitions</li> <li>• Three primary partitions and one extended partition</li> </ul> <p>The extended partition must be divided into logical partitions before use. For example, to create six partitions, you can create them in the following two ways:</p> <ul style="list-style-type: none"> <li>• Three primary partitions and one extended partition, with the extended partition divided into three logical partitions</li> <li>• One primary partition and one extended partition, with the extended partition divided into five logical partitions</li> </ul> <p>For how to create logical partitions, see <a href="#">Creating a Logical Volume Using LVM</a>.</p>	<ul style="list-style-type: none"> <li>• fdisk</li> <li>• parted</li> </ul>	<p><a href="#">Initializing a Linux Data Disk (Less Than or Equal to 2 TiB)</a></p> <p><a href="#">Initializing a Windows Data Disk</a></p>

Partit ion Style	Max. Disk Size Supported	Max. Number of Partitions Supported	Linux Partitio ning Tool	Reference
GPT	18 EiB 1 EiB = 1048576 TiB	Not limited GPT partitions are not categorized.	parted	<a href="#">Initializing a Linux Data Disk (Greater Than 2 TiB)</a> <a href="#">Initializing a Windows Data Disk</a>

## Common EVS Device Names in Linux

The following lists some common EVS device names in Linux:

- System disk: /dev/vda, /dev/sda, and /dev/xvda
- Data disk: /dev/vd[b-z], /dev/sd[b-z], and /dev/xvd[b-z]  
Examples: /dev/vdb, /dev/vdc, /dev/sdb, /dev/sdc, /dev/xvdb, and /dev/xvdc
- Partition: /dev/vd[a-z][digit], /dev/sd[a-z][digit], /dev/xvd[a-z][digit]  
Examples: /dev/vda1, /dev/vda2, /dev/vdb1, /dev/vdb2

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

### Scenarios

After a newly purchased data disk is attached to a server, you must log in to the server and initialize the disk before you can use the disk. This section describes how to initialize a Linux data disk. The operations may vary depending on the server OS.

The maximum disk size that MBR supports is 2 TiB, and that GPT supports is 18 EiB. If your disk capacity is less than 2 TiB, you can use either the MBR or GPT partition style. However, if your disk is greater than 2 TiB or you plan to expand it to over 2 TiB later, use **GPT** when initializing disks.

**Table 2-8** Initialization instructions

Operation	OS Requirements	Partition Style	Common File Systems	Example Configuration
<a href="#">Initializing a Data Disk Using fdisk</a>	None	MBR	ext* (such as ext2, ext3, and ext4), xfs, and btrfs	<ul style="list-style-type: none"> <li>● OS: CentOS 7.4 64-bit</li> <li>● Partitioning tool: fdisk</li> <li>● Device name: /dev/vdb</li> <li>● File system format: ext4</li> <li>● Mount points: /mnt/sdc and /mnt/sdd</li> <li>● Partition 1: /dev/vdb1                             <ul style="list-style-type: none"> <li>– Size: 40 GiB</li> <li>– Partition style: MBR</li> </ul> </li> <li>● Partition 2: /dev/vdb2                             <ul style="list-style-type: none"> <li>– Size: 60 GiB</li> <li>– Partition style: MBR</li> </ul> </li> </ul>
<a href="#">Initializing a Data Disk Using parted</a>	None	<ul style="list-style-type: none"> <li>● GPT</li> <li>● MBR</li> </ul>		See <a href="#">Initializing a Data Disk Using parted</a> .

## Prerequisites

The disk has been attached to a server. For how to attach disks, see [Attaching an EVS Disk](#).

## 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. For how to back up data using CBR, see [Backing Up EVS Disks](#). For how to back up data using snapshots, see [Creating an EVS Snapshot](#).
- 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 Using fdisk

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

**Step 1** Log in to the server as user **root**.

For how to log in to an ECS, see [How Do I Log In to My ECS?](#)

For how to log in to a BMS, see [Linux BMS Login Methods](#).

**Step 2** Create partitions. In this example, 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 40G 0 part /
vdb 253:16 0 100G 0 disk
```

The command output shows that there are two EVS disks. **/dev/vda** is the system disk, and **/dev/vdb** is the new data disk, whose 100 GiB is not partitioned.

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

#### fdisk /dev/vdb

**n**

**p**

**1**

```
[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**

- Entering **p** for **Partition type** creates a primary partition, and entering **e** creates an extended partition. The default value is **p**.
- **Partition number** indicates the partition serial number. Enter a value ranging from **1** to **4**.

#### NOTE

- Enter **n** and press **Enter** to create a new partition.
- Enter **p** and press **Enter** to create a primary partition.
- Enter **1** and press **Enter** to set a partition number. Partition number **1** is used in this example.

3. Set **First sector** to **2048** and **Last sector** to **83886079** for partition **/dev/vdb1** (40 GiB). For how to calculate the values, see [Table 2-9](#).

```
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
```

**Table 2-9** First and last sectors in this example are calculated as follows

Sector	/dev/vdb1 (40 GiB)	/dev/vdb2 (60 GiB)	Formula for Calculating the Value of sectors
First sector	2048 (The first sector of the /dev/vdb data disk is used.)	<b>Last sector of /dev/vdb1 + 1</b> = 83886079 + 1 = 83886080	<b>Value of sectors</b> = Capacity × 1073741824/512
Last sector	<b>Value of sectors - 1</b> = (40 × 1073741824/512) - 1 = 83886079	<b>First sector + Value of sectors - 1</b> = 83886080 + (60 × 1073741824/512) - 1 = 209715199	

4. Create the second primary partition /dev/vdb2 for /dev/vdb.

**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
```

**NOTE**

- Enter **n** and press **Enter** to create a new partition.
- Enter **p** and press **Enter** to create a primary partition.
- Enter **2** and press **Enter** to set a partition number. Partition number 2 is used in this example.

5. 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
```

6. Check the sizes and partition styles of the new partitions.

**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):
```

**Disk label type:** `dos` indicates the MBR partition style.

 **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 and step 2.

7. Write the changes to the partition table.

**w**

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

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

The partition is created.

 **NOTE**

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.

8. Synchronize the new partition table to the OS.

**partprobe**

**Step 3** Create file systems.

1. 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**

- `mkfs -t <file-system-format> <disk-partition-name>`: To create an xfs file system, the command is `mkfs -t xfs <disk-partition-name>`. To create a btrfs file system, the command is `mkfs -t btrfs <disk-partition-name>`.
- 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
```

2. Check whether the file system format is ext4.

**blkid /dev/vdb**

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

**Step 4** Create directories (mount points) and mount the new partitions on the created mount points.

1. Mount **/dev/vdb1** on **/mnt/sdc**.

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

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

2. Mount **/dev/vdb2** on **/mnt/sdd**.

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

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

3. View the mount results.

**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 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 5** (Optional) Use the partition UUIDs to configure auto mount at startup.

UUIDs are the unique character strings for identifying partitions in Linux. Mounts become invalid after a system reboot. You can configure auto mount at startup by adding information of the new partitions into the **/etc/fstab** file.

**NOTE**

- 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 server stops or starts. This can even prevent your server from booting up.
- 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/vdb
/dev/vdb1: UUID="0b3040e2-1367-4abb-841d-ddb0b92693df" TYPE="ext4"
/dev/vdb2: UUID="0d6769k2-1745-9dsf-453d-hgd0b34267dj" TYPE="ext4"
```

Take note of the partition UUIDs, which will be used in the next step. In this example, the partition UUIDs 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 2-10** Content description

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

**Step 6** (Optional) Verify that auto mount takes effect.

You can restart the server to check whether auto mount takes effect. Alternatively, you can perform the following steps to simulate auto mount.

1. To verify auto mount, unmount the partitions first.

```
umount /dev/vdb1
```

```
umount /dev/vdb2
```

2. Reload all the content in the **/etc/fstab** file. **/etc/fstab** is a static file system table that contains the list of file systems that need to be automatically mounted during system startup.

```
mount -a
```

3. Query the file system mount 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

## Initializing a Data Disk Using parted

You can use parted to create either MBR or GPT partitions. The only difference is that the command used to set the partition style is different. All other operations are the same. For the initialization instructions with parted, see [Initializing a Data Disk Using parted](#).

Commands for setting partition styles:

MBR:

```
mklabel msdos
unit s
p
```

GPT:

```
mklabel gpt
unit s
p
```

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

### Scenarios

After a newly purchased data disk is attached to a server, you must log in to the server and initialize the disk before you can use the disk.

When your disk is greater than 2 TiB, you can only use the parted tool and create GPT partitions. The initialization operations may vary depending on the server OS. Perform initialization operations based on your server OS. To learn about the differences between MBR and GPT, see [Partition Styles](#).

Operation	Partition Style	OS Requirements	Common File Systems	Partitioning Tool	Example Configuration
<a href="#">Initializing a Data Disk Using parted</a>	GPT	None	ext* (such as ext2, ext3, and ext4), xfs, and btrfs	parted	<ul style="list-style-type: none"><li>• OS: CentOS 7.4 64-bit</li><li>• Device name: /dev/vdb</li><li>• File system format: ext4</li><li>• Mount point: /mnt/sdc</li><li>• Partition name: /dev/vdb1</li><li>• Partition style: GPT</li><li>• Size: 3 TiB</li></ul>

## Prerequisites

The disk has been attached to a server. For how to attach disks, see [Attaching an EVS Disk](#).

## 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. For how to back up data using CBR, see [Backing Up EVS Disks](#). For how to back up data using snapshots, see [Creating an EVS Snapshot](#).
- 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 Using parted

Major initialization steps include creating partitions and file systems, mounting partitions, and configuring auto mount at system startup.

**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 as user **root**.

For how to log in to an ECS, see [How Do I Log In to My ECS?](#)

For how to log in to a BMS, see [Linux BMS Login Methods](#).

**Step 2** Create partitions. In this example, 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
```

The command output shows that there are two EVS disks. **/dev/vda** is the system disk, and **/dev/vdb** is the new data disk, whose 3 TiB is not partitioned.

2. Create a partition. Replace **/dev/vdb** in the command with your desired disk name.

**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 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

- **mklabel** *<disk-partition-style>*: If your disk capacity is less than or equal to 2 TiB and you want to use parted to create an MBR partition, run the **mklabel msdos** command.
  - **unit s**: This command sets the measurement unit of the disk to sector.
  - 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 partition name to **/dev/vdb1** and allocate all the storage space to this partition.

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

```
p
```

 NOTE

- **mkpart** *<partition-name>* *<first-sector-value>* *<last-sector-value>*: In the example command, **2048s** is the start sector value, and **100%** indicates to allocate 100% of the disk space to the **/dev/vdb1** partition.
- If you want to allocate the data disk capacity to 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	

5. 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 file systems.

1. In this example, create an ext4 file system on the **/dev/vdb1** partition.

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

 NOTE

- **mkfs -t** *<file-system-format>* *<disk-partition-name>*: To create an xfs file system, the command is **mkfs -t xfs** *<disk-partition-name>*. To create a btrfs file system, the command is **mkfs -t btrfs** *<disk-partition-name>*.
- 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
```

```
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 ~]#
```

2. After the file system is created, run the following commands to check the file system format:

```
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 3299GB 3299GB  ext4        /dev/vdb1

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

If **ext4** is displayed under **File system**, the creation is successful.

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** (Optional) Use the partition UUID to configure auto mount at startup.

UUIDs are the unique character strings for identifying partitions in Linux. Mounts become invalid after a system reboot. You can configure auto mount at startup by adding information of the new partitions into the **/etc/fstab** file.

 **NOTE**

- 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 server stops or starts. This can even prevent your server from booting up.
- 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"
```

Take note of the partition UUID, which will be used in the next step. In this example, 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 2-11** Content description

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

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

**Step 6** (Optional) Verify that auto mount takes effect.

You can restart the server to check whether auto mount takes effect. Alternatively, you can perform the following steps to simulate auto mount.

1. To verify auto mount, unmount the partition first.

```
umount /dev/vdb1
```

2. Reload all the content in the **/etc/fstab** file. **/etc/fstab** is a static file system table that contains the list of file systems that need to be automatically mounted during system startup.

```
mount -a
```

3. Query the file system mount 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.

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

## Prerequisites

The disk has been attached to a server. For how to attach disks, see [Attaching an EVS Disk](#).

## 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. For how to back up data using CBR, see [Backing Up EVS Disks](#). For how to back up data using snapshots, see [Creating an EVS Snapshot](#).
- 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

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 [How Do I Log In to My ECS?](#)

For how to log in to a BMS, see [Windows BMS Login Methods](#).

**Step 2** On the server desktop, right-click  and choose **Disk Management** from the shortcut menu.

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.

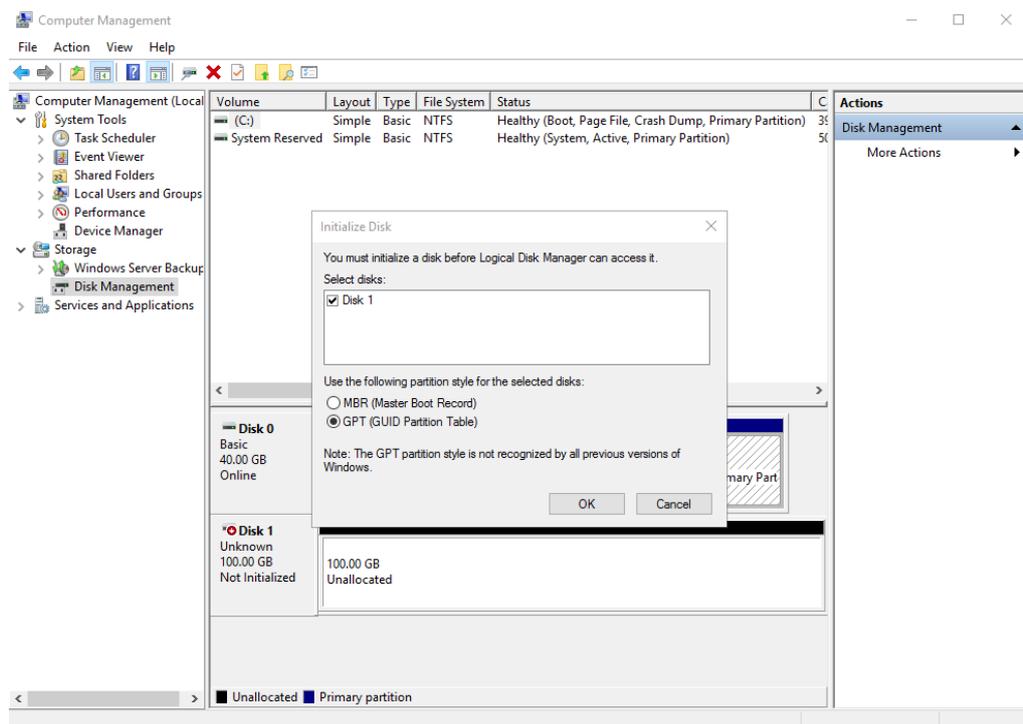
In the **Initialize Disk** dialog box, the to-be-initialized disk is selected. Select a partition style and click **OK**. In this example, **GPT (GUID Partition Table)** is selected.

**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.

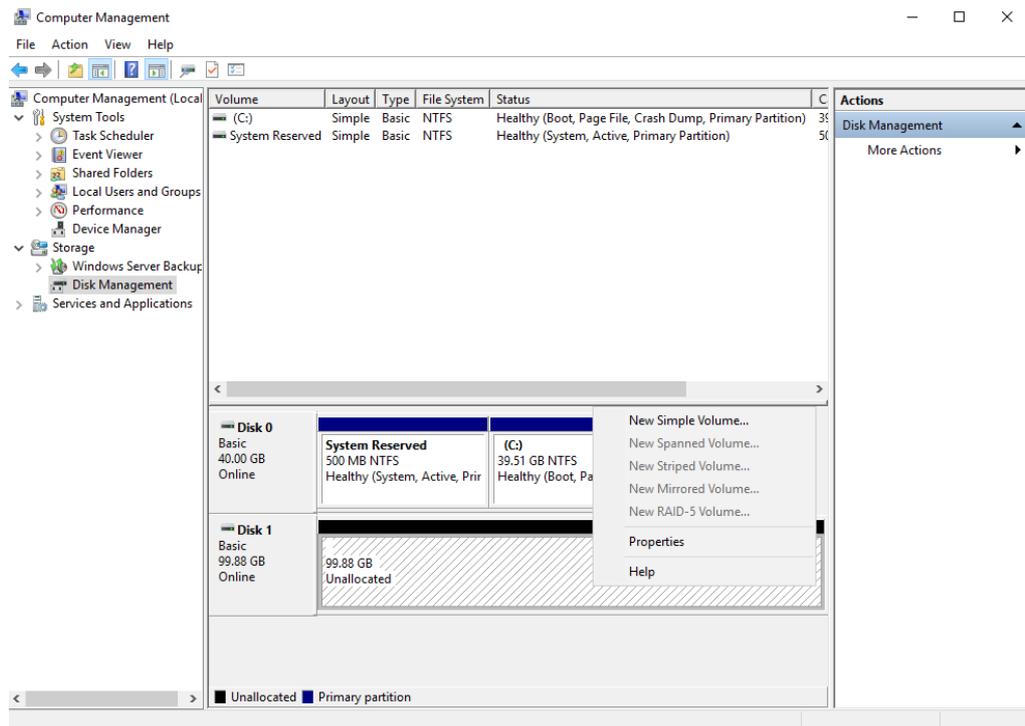
If the partition style of an in-use disk is changed, all data on the disk will be lost, so take care to select an appropriate partition style when initializing the disk. If you must change the partition style to GPT, it is recommended that you back up the disk data before the change.

**Figure 2-5** Disk list



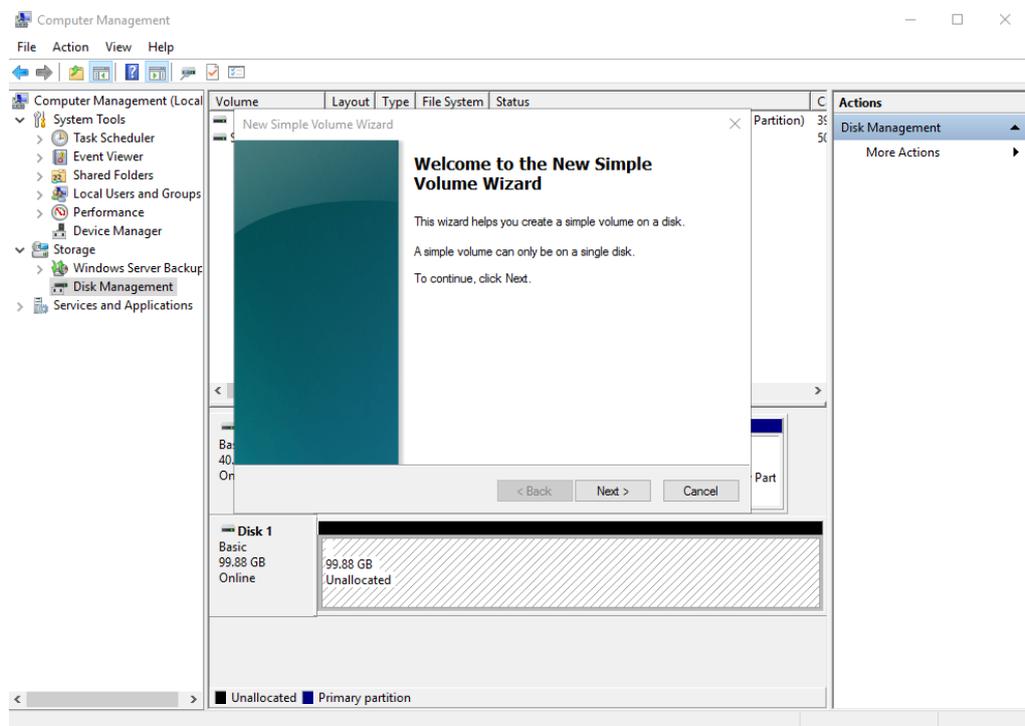
**Step 3** In the **Unallocated** area of **Disk 1**, right-click the blank area and choose **New Simple Volume**.

Figure 2-6 Computer Management



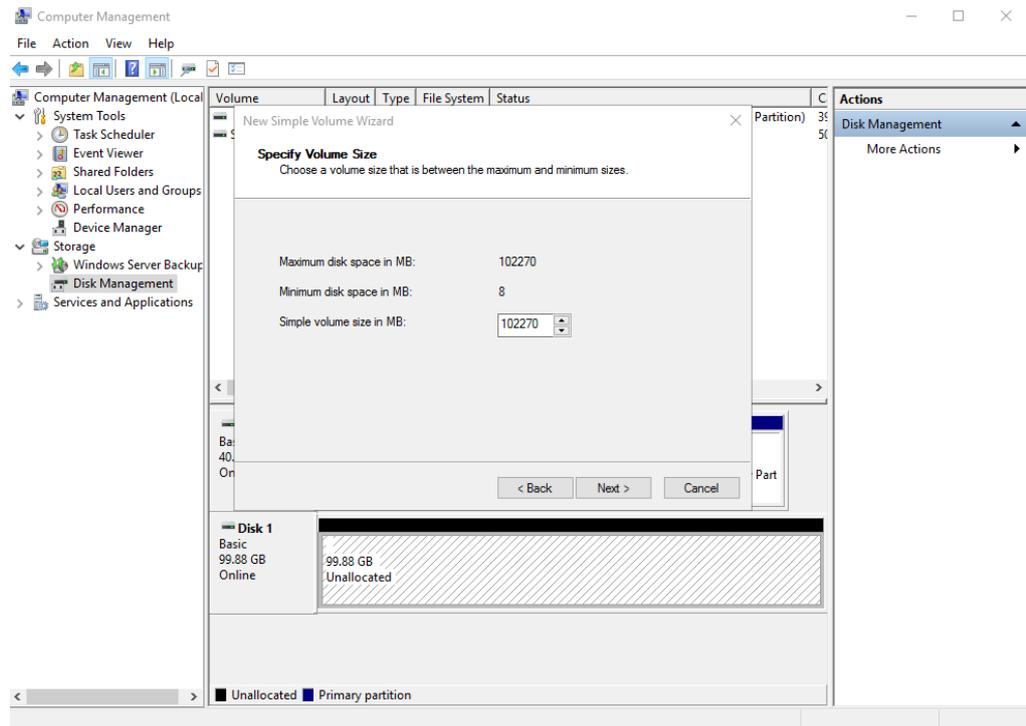
The **New Simple Volume Wizard** window is displayed.

Figure 2-7 New Simple Volume Wizard



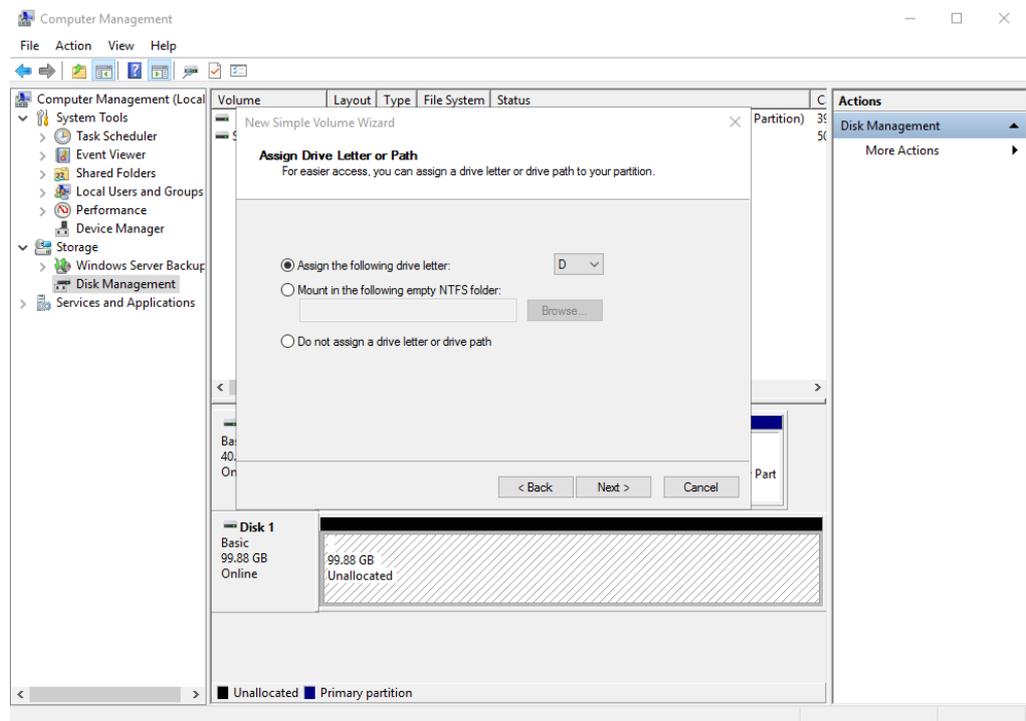
**Step 4** Click **Next** to go to the **Specify Volume Size** page.

Figure 2-8 Specify Volume Size



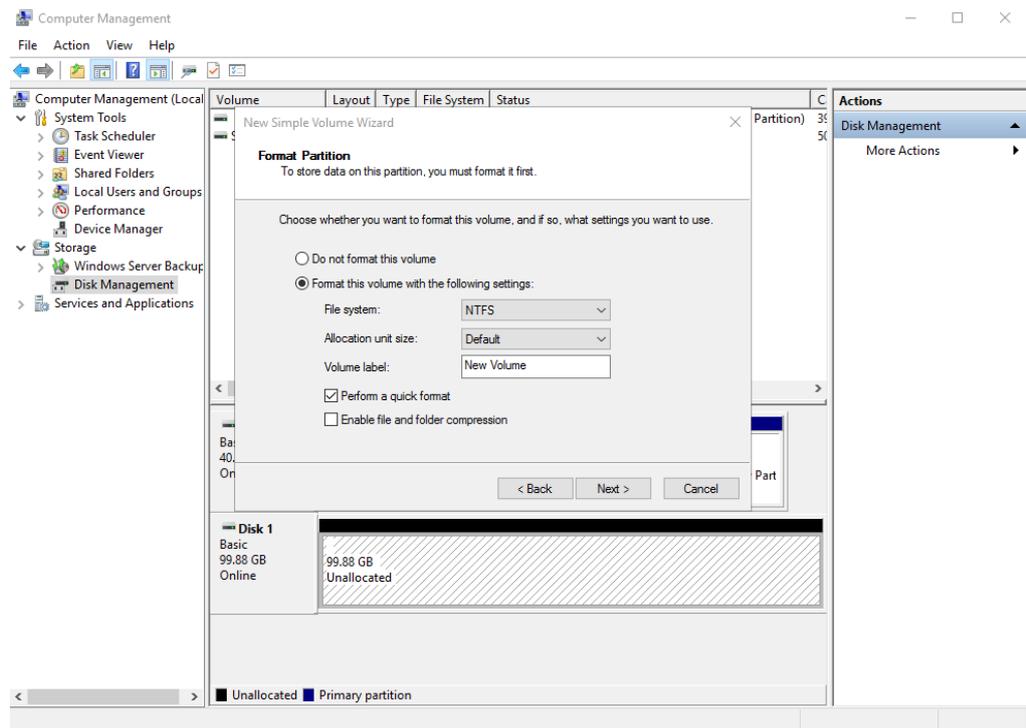
**Step 5** Specify the volume size and click **Next**. The system selects the maximum volume size by default. You can specify the volume size as required. In this example, the default setting is used.

Figure 2-9 Assign Drive Letter or Path



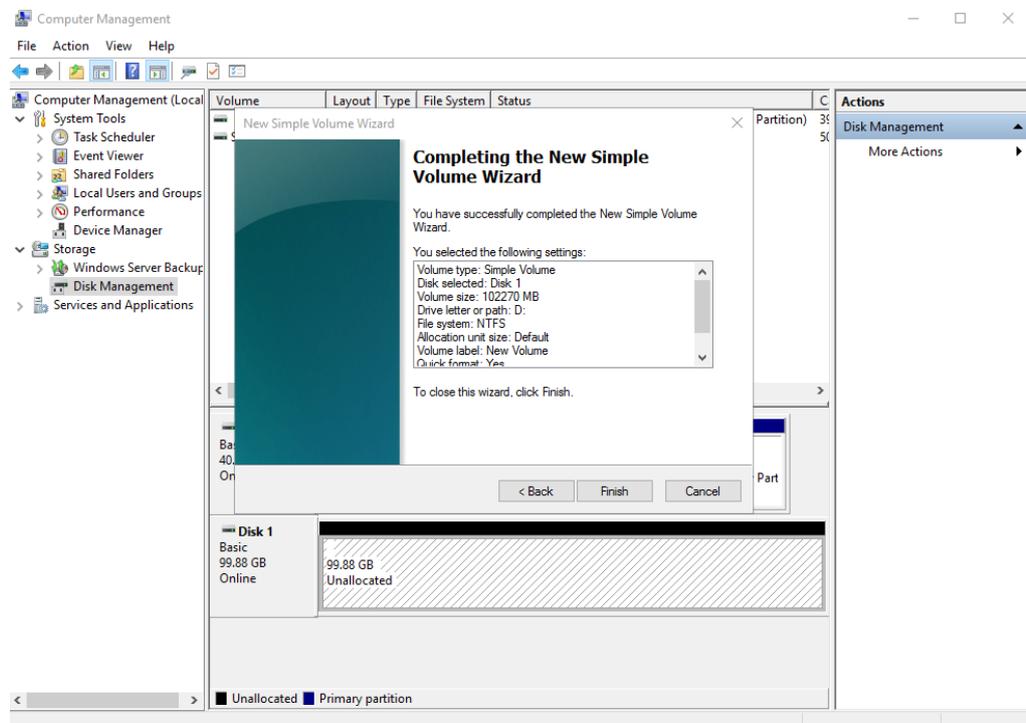
**Step 6** Assign a drive letter or path to your partition and click **Next**. The system assigns drive letter D by default. In this example, the default setting is used.

Figure 2-10 Format Partition



**Step 7** Specify format settings and click **Next**. The system selects the NTFS file system by default. You can specify a file system format as required. In this example, the default setting is used.

Figure 2-11 Completing the New Simple Volume Wizard



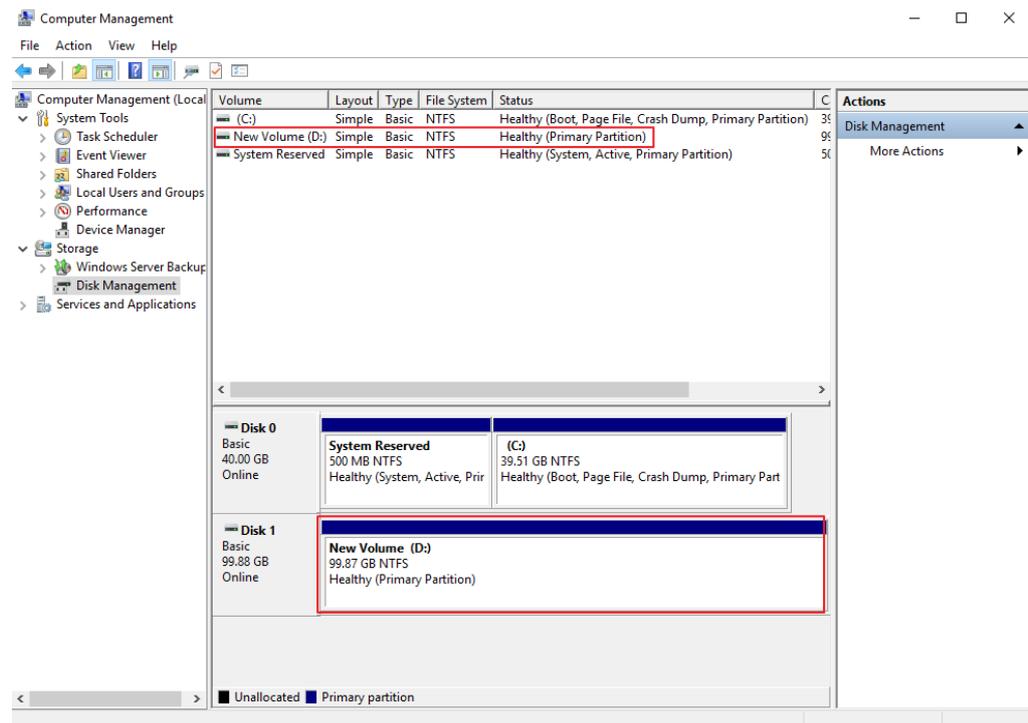
**NOTICE**

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

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

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

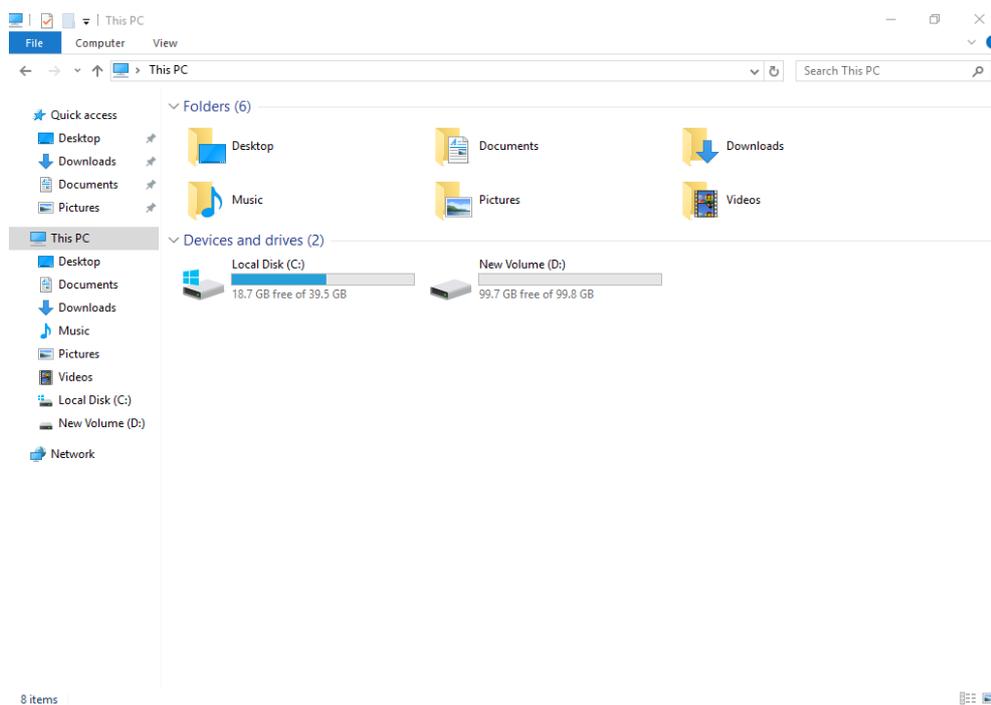
**Figure 2-12** Disk initialized



**Step 9** 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 2-13 File Explorer



-----End

# 3 Expanding the EVS Disk Capacity

---

## 3.1 Expansion Overview

As your business develops or data volume grows, existing EVS disk capacities may no longer meet your storage needs. EVS capacity expansion is the solution to this issue. You can increase the storage space of existing EVS disks without interrupting services while ensuring continuous data storage and service stability.

### Upper Limits on Disk Capacity

The maximum disk capacity is as follows:

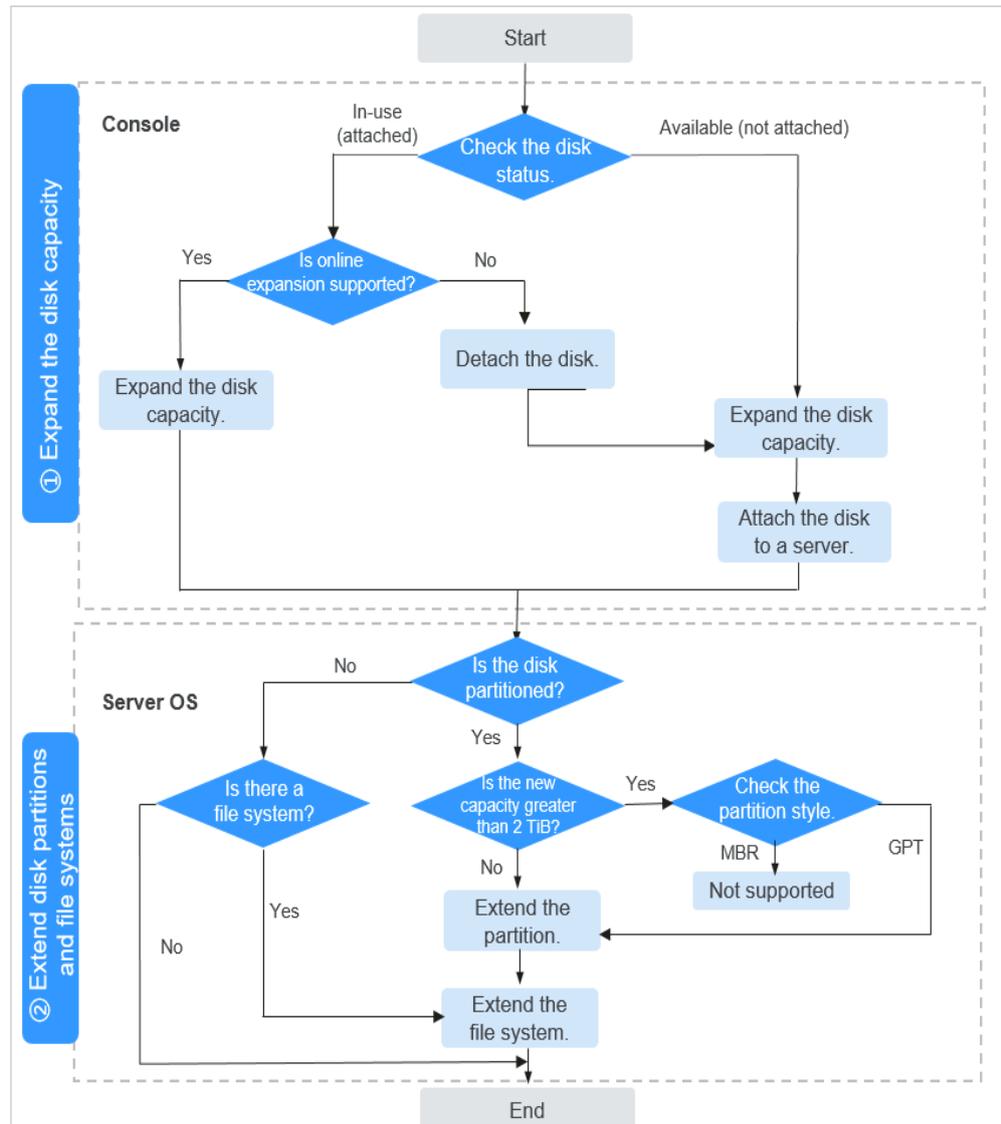
- System disk: 1 TiB
- Data disk: 32 TiB

 **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?

Figure 3-1 Capacity expansion process



- Step 1: **Expand the disk capacity on the console.**  
Choose a proper expansion method based on the disk status. To check the disk status, see [Viewing EVS Disk Details](#). To learn about the expansion conditions under different statuses, see [Prerequisites](#).
- Step 2: **Extend disk partitions and file systems on the server.**  
Before extending the partitions and file systems, you need to select an appropriate extension method based on the disk partitions. For details, see [Table 3-1](#). For how to check the disk partition style in Linux, see [How Do I Check the Disk Partition Style in Linux?](#)

**Table 3-1** Capacity expansion instructions

Operation Category		Operation Instruction
The disk is not partitioned.		On Linux ECSs, you can directly extend the file systems. For details, see <a href="#">How Do I Extend the File System of an Unpartitioned Data Disk in Linux?</a>
The disk is partitioned.	<ul style="list-style-type: none"> <li>• MBR partitions with no more than 2 TiB disk capacity after expansion</li> <li>• GPT partitions with no more than 18 EiB disk capacity after expansion</li> </ul>	<ul style="list-style-type: none"> <li>• Expansion instruction for Linux: <a href="#">Extending Disk Partitions and File Systems (Linux)</a></li> <li>• Expansion instruction for Windows: <a href="#">Extending Disk Partitions and File Systems (Windows)</a></li> </ul>
	MBR partitions with more than 2 TiB disk capacity after expansion	<p>Not supported. The maximum disk capacity that MBR supports is 2 TiB, and the disk space in excess of 2 TiB cannot be used.</p> <ul style="list-style-type: none"> <li>• Solution 1: If your disk uses MBR and you need to expand the disk capacity to over 2 TiB, change the partition style from MBR to GPT. Ensure that the disk data has been backed up before changing the partition style because services will be interrupted and data on the disk will be cleared during this change.</li> <li>• Solution 2: <a href="#">Buy a new disk</a> and partition the disk using GPT.</li> </ul>

 **NOTE**

If a server is stopped during the expansion, the additional space added to 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 instructions shown in the preceding table.

## Billing for Capacity Expansion

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](#).

## 3.2 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.

- [Expanding the Capacity of a Single Disk](#)
- [Expanding Capacities of Multiple Disks in a Batch](#)

### Prerequisites

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

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

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

- The server status must be **Running** or **Stopped**.
- The server OS must meet the requirements described in [Checking Whether the OS Supports the Capacity Expansion of In-use Disks](#).

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

#### NOTE

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

### Constraints

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

### Expanding the Capacity of a Single Disk

**Step 1** Sign in to the [EVS console](#).

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

**Step 3** Choose an entry to expand the capacity.

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

**Step 4** On the **Expand Capacity** page, set **New Capacity** and click **Next**.

**Step 5** In the displayed **Note** dialog box, read the note, and click **Expand Capacity**.

**Step 6** On the **Expand Capacity** page, check the disk configuration.

- Click **Submit** to start expanding a pay-per-use disk. For a yearly/monthly disk, make the payment before you can continue.
- Click **Previous** to change the settings, if required.

**Step 7** 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 disk has been expanded on the console.

 **NOTE**

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

**Step 8** (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 9** After the disk has been expanded on the console, log in to the server and extend the partition and file system, because the previous steps only enlarge the disk space.

The operations vary depending on the server OS.

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

----End

## Expanding Capacities of Multiple Disks in a Batch

**Step 1** Sign in to the [EVS console](#).

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

- Step 3** Expand the service list and click  in the upper left corner to select a region.
- Step 4** Choose **Storage > Elastic Volume Service**.  
The **Elastic Volume Service** page is displayed.
- Step 5** In the disk list, select the disks you want to expand their capacities.
- Step 6** Click **Expand Capacity** above the disk list.
- Step 7** In the **Expand Capacity** dialog box, click **Confirm**.
- Step 8** On the displayed page, set a new capacity for all target disks.
- Step 9** Click **Next**.
- Step 10** In the displayed **Note** dialog box, read the note, select the checkbox, and click **Expand Capacity**.
- Step 11** On the displayed page, check the expansion information.
- Click **Expand Capacity** to start the expansion.
  - Click **Previous** to change the settings, if required.
- Step 12** If there are yearly/monthly disks, make the payment before you can continue.
- Step 13** In the disk list, view the capacity of the target disks.

When the disk statuses change from **Expanding** to **In-use** or **Available**, and the disk capacities increase, 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.

 **NOTE**

If the expansion fails, technical support personnel will contact you and help you handle this error. Do not perform any operations on the disk before the technical support personnel contact you. If you require that the error be handled as soon as possible, contact our technical support personnel. A disk will no longer be billed if its status changes to **Expansion failed**.

- Step 14** (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 15** After the disk has been expanded on the console, log in to the server and extend the partition and file system, because the previous steps only enlarge the disk space.

The operations vary depending on the server OS.

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

----End

## Checking Whether the OS Supports the Capacity Expansion of In-use Disks

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

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

To view such images, click  in the upper left corner, 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 3-2](#).

If it is included in [Table 3-2](#), you can expand capacity while the disk is in use and you do not need to stop and start the server after the expansion. Otherwise, you must detach the disk and then expand its capacity, or stop and start the server after the expansion.

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

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

OS	Version
Windows Server 2012	Windows Server 2012 R2 Standard 64-bit
Windows Server 2008	Windows Server 2008 R2 Enterprise 64-bit

## Related Links

- After expanding disk capacities on the console, you can use LVM to manage disk partitions with more flexibility. For details, see [Using LVM to Manage EVS Disks](#).
- For more expansion FAQs, see [Capacity Expansion](#).
- You can also expand disk capacities through API calls. For details, see [Expanding the Capacity of an EVS Disk](#).
- If the capacity expansion fails, you can roll back data using a snapshot. For details, see [Rolling Back Disk Data from a Snapshot](#).

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

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

#### Scenarios

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

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

**Table 3-3** Operation instructions of extending partitions and file systems in Linux

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

Scenario	Partition Style	Disk Function	OS Requirements	File System Format	Tool	Example Configuration
<b>Extending an Existing MBR Partition (for System Disks Whose Kernel Version Is Earlier Than 3.6.0)</b>	MBR	System disk	The kernel version is earlier than 3.6.0.	ext* (such as ext2, ext3, and ext4), xfs, and btrfs	dracut-module s-growroot	<ul style="list-style-type: none"> <li>• Device name: /dev/vda</li> <li>• File system format: ext4</li> <li>• Mount point: /mnt/sda</li> <li>• Partition name: /dev/vda1</li> <li>• Space added: 60 GiB</li> <li>• Partition style: MBR</li> </ul>
<b>Creating a New MBR Partition</b>	MBR	System disk Data disk	None	ext* (such as ext2, ext3, and ext4), xfs, and btrfs	<ul style="list-style-type: none"> <li>• fdisk</li> <li>• parted</li> </ul>	<ul style="list-style-type: none"> <li>• Device name: /dev/vdb</li> <li>• File system format: ext4</li> <li>• Mount points: /mnt/sdc and /mnt/sdd</li> <li>• Partition 1: /dev/vdb1 <ul style="list-style-type: none"> <li>– Size: 100 GiB</li> <li>– Partition style: MBR</li> </ul> </li> <li>• Partition 2: /dev/vdb2 <ul style="list-style-type: none"> <li>– Size: 50 GiB</li> <li>– Partition style: MBR</li> </ul> </li> </ul>

Scenario	Partition Style	Disk Function	OS Requirements	File System Format	Tool	Example Configuration
<a href="#">Creating a New GPT Partition</a>	GPT	Data disk	None	ext* (such as ext2, ext3, and ext4), xfs, and btrfs	parted	<ul style="list-style-type: none"> <li>• Device name: /dev/vdb</li> <li>• File system format: ext4</li> <li>• Mount points: /mnt/sdc and /mnt/sdd</li> <li>• Partition 1: /dev/vdb1                             <ul style="list-style-type: none"> <li>- Size: 2 TiB</li> <li>- Partition style: GPT</li> </ul> </li> <li>• Partition 2: /dev/vdb2                             <ul style="list-style-type: none"> <li>- Size: 1 TiB</li> <li>- Partition style: GPT</li> </ul> </li> </ul>
<a href="#">Extending a Logical Volume</a>	If you use Logical Volume Manager (LVM) to manage EVS disks, you can extend logical volumes when the capacity fails to meet your requirements.					

 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 That Is Quickly Provisioned?](#)

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

## Constraints

- The additional space of a data disk cannot be added to the root partition. To extend the root partition, expand the system disk instead.
- During an expansion, the additional space is added to the end of the disk. If the disk has multiple partitions, the additional space can only be allocated to the last partition of the disk.
- If the target partition is an extended MBR partition (whose partition number is usually greater than or equal to 5), you need to first expand the extended partition and then the logical partition. Assume that you have three partitions, **/dev/vdb1** (primary partition), **/dev/vdb2** (extended partition),

and `/dev/vdb5` (logical partition), you need to run `growpart /dev/vdb2` and then `growpart /dev/vdb5` to extend the partitions.

- The maximum disk capacity that MBR supports is 2 TiB, and the disk space in excess of 2 TiB cannot be used. 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 [Creating an EVS Snapshot](#) respectively.

## Extending the Disk Partition and File System

When extending a partition, you can choose to extend an existing partition or create a new partition. To check the current disk partition style, run the `parted <disk-name>` command.

## Extending an Existing Partition

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

**Step 1** Log in to the server as user `root`.

For how to log in to an ECS, see [How Do I Log In to My ECS?](#)

For how to log in to a BMS, see [Linux BMS Login Methods](#).

**Step 2** Check the capacity expansion tool and current disk information.

1. Check whether the `growpart` expansion 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 (default: 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 ~]#
```

#### NOTE

- If error message "Read-only file system" is returned, run **mount -o remount,rw /** to change the file system permissions to read/write.
  - If your OS does not allow using online resizing tools (such as growpart and resize2fs) to extend mounted partitions, you can unmount the partitions first and then extend them offline. For details, see [Extending Partitions and File Systems Offline \(Linux Data Disk\)](#).
- 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.

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

#### **lsblk**

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

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

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

- Step 3** Add the additional space to the **/dev/vdb1** partition that already exists.

#### **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

- Command line structure: **growpart** <disk-partition> <partition-number>. Note that there is a space in front of the partition number.
- If the following command output is displayed:  
no tools available to resize disk with 'gpt'  
FAILED: failed to get a resizer for id "  
The disk uses the GPT partition style, and the **gdisk** tool is required when you use **growpart** to add the additional space. In this case, run **yum install gdisk**, enter **y** to install **gdisk**, and then run the preceding **growpart** command.
- If the following command output is displayed:  
growpart /dev/vda 1 unexpected output in sfdisk --version [sfdisk is from util-linux 2.23.2]  
Check whether the system character set (language environment) is **en\_US.UTF-8**. If not, run **export LC\_ALL=en\_US.UTF-8**.
- If error message "NOCHANGE:partition 1 is size 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 (at 100% usage). Back up the disk data and clear unnecessary files or programs.

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

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

**parted /dev/vdb****P**

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

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

(parted)
```

**Partition Table** shows the partition style, which is GPT in this example. **File system** shows the file system format, which is ext4 in this example.

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.
2. As the file system format of **/dev/vdb1** is ext4, run the following command to extend the file system:

**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

- Command line structure: **resize2fs <disk-partition>**
- If the error message "open: No such file or directory while opening /dev/vdb1" is returned, an incorrect partition is specified. Run **parted** to view disk partitions.
- If the file system format is xfs, run the following command (**/mnt/sdc** is the mount point of **/dev/vdb1**. Change it based on your actual conditions):

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

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

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

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

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

**Step 6** (Optional) If you are expanding a data disk whose OS kernel version is earlier than 3.6.0, after the extension operation is complete, 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 [Creating an EVS Snapshot](#).

----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 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.
- 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** Log in to the server as user **root**.

For how to log in to an ECS, see [How Do I Log In to My ECS?](#)

For how to log in to a BMS, see [Linux BMS Login Methods](#).

**Step 2** (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 3** Regenerate the **initramfs** file.

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

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

**Step 4** 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 5** Restart the server.

**reboot**

Reconnect to the server after it is restarted.

**Step 6** 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/vda1** 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** Log in to the server as user **root**.

For how to log in to an ECS, see [How Do I Log In to My ECS?](#)

For how to log in to a BMS, see [Linux BMS Login Methods](#).

**Step 2** 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 3** 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 partitions include primary partitions and extended partitions. A maximum of four primary partitions are supported. If you need more partitions, create 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). For how to calculate the sector values, see [Table 3-4](#).

**Table 3-4** First and last sectors in this example are calculated as follows

Sector	/dev/vdb1 (40 GiB)	/dev/vdb2 (60 GiB)	Formula for Calculating the Value of sectors
First sector	2048 (The first sector of the <b>/dev/vdb</b> data disk is used.)	<b>Last sector of /dev/vdb1 + 1</b> = 83886079 + 1 = 83886080	<b>Value of sectors</b> = Capacity × 1073741824/512
Last sector	<b>Value of sectors - 1</b> = (40 × 1073741824/512) - 1 = 83886079	<b>First sector + Value of sectors - 1</b> = 83886080 + (60 × 1073741824/512) - 1 = 209715199	

**Step 4** 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**

**Step 5** Create an ext4 file system on the `/dev/vdb2` partition.

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

 **NOTE**

- **mkfs -t** *<file-system-format>* *<disk-partition-name>*: To create an xfs file system, the command is **mkfs -t xfs** *<disk-partition-name>*. To create a btrfs file system, the command is **mkfs -t btrfs** *<disk-partition-name>*.
- 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/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 6** 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 7** 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 server 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"
```

Take note of the partition UUID, which will be used in the next step. In this example, the UUID of the **/dev/vdb2** partition 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 3-5** Content description

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

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

**Step 8** Verify that auto mount takes effect.

You can restart the server to check whether auto mount takes effect. Alternatively, you can perform the following steps to simulate auto mount.

1. To verify auto mount, unmount the partition first.

```
umount /dev/vdb2
```

2. Reload all the content in the **/etc/fstab** file. **/etc/fstab** is a static file system table that contains the list of file systems that need to be automatically mounted during system startup.

```
mount -a
```

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

3. Query the file system mount 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** Log in to the server as user **root**.

For how to log in to an ECS, see [How Do I Log In to My ECS?](#)

For how to log in to a BMS, see [Linux BMS Login Methods](#).

**Step 2** 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 3** 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, which will be used in the next step. In this example, the last sector of the `/dev/vdb1` partition is **4294965247s**.

**unit s** means that the display and operation unit is sectors.

**p** prints the partition table information.

#### 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 end of the new disk.

- 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

- **mkpart** *<partition-name>* *<first-sector-value>* *<last-sector-value>*: In the example command, the first sector is **4294965248s**, which is the last sector of `/dev/vdb1` plus one. **100%** indicates to allocate 100% of the disk space to the `/dev/vdb2` partition.
- If you want to allocate the additional space to two or more partitions, calculate the first and last sectors of the partitions based on the method provided in [Table 3-4](#).

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

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

 NOTE

- **mkfs -t** *<file-system-format>* *<disk-partition-name>*: To create an xfs file system, the command is **mkfs -t xfs** *<disk-partition-name>*. To create a btrfs file system, the command is **mkfs -t btrfs** *<disk-partition-name>*.
- 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/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 5** 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 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 server 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"
```

Take note of the partition UUID, which will be used in the next step. In this example, the UUID of the `/dev/vdb2` partition 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 3-6** Content description

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

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

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

You can restart the server to check whether auto mount takes effect. Alternatively, you can perform the following steps to simulate a system restart and auto mount.

1. To verify auto mount, unmount the partition first.

```
umount /dev/vdb2
```

2. Reload all the content in the `/etc/fstab` file. `/etc/fstab` is a static file system table that contains the list of file systems that need to be automatically mounted during system startup.

```
mount -a
```

3. Query the file system mount 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
```

## Related Links

- For more expansion FAQs, see [Capacity Expansion](#).
- If the capacity expansion fails, you can roll back data using a snapshot. For details, see [Rolling Back Disk Data from a Snapshot](#).

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

### Scenarios

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

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

- Extending an Existing Partition
- Creating a New Partition

## Constraints

- The additional space of a data disk cannot be added to the root partition. To extend the root partition, expand the system disk instead.
- During an expansion, the additional space is added to the end of the disk. If the disk has multiple partitions, the additional space can only be allocated to the last partition of the disk.
- If the target partition is an extended MBR partition (whose partition number is usually greater than or equal to 5), you need to first expand the extended partition and then the logical partition. Assume that you have three partitions, `/dev/vdb1` (primary partition), `/dev/vdb2` (extended partition), and `/dev/vdb5` (logical partition), you need to run `growpart /dev/vdb2` and then `growpart /dev/vdb5` to extend the partitions.
- The maximum disk capacity that MBR supports is 2 TiB, and the disk space in excess of 2 TiB cannot be used. 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 [Creating an EVS Snapshot](#) respectively.

## Extending the Disk Partition and File System

When extending a partition, you can choose to extend an existing partition or create a new partition.

### Extending an Existing Partition

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

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

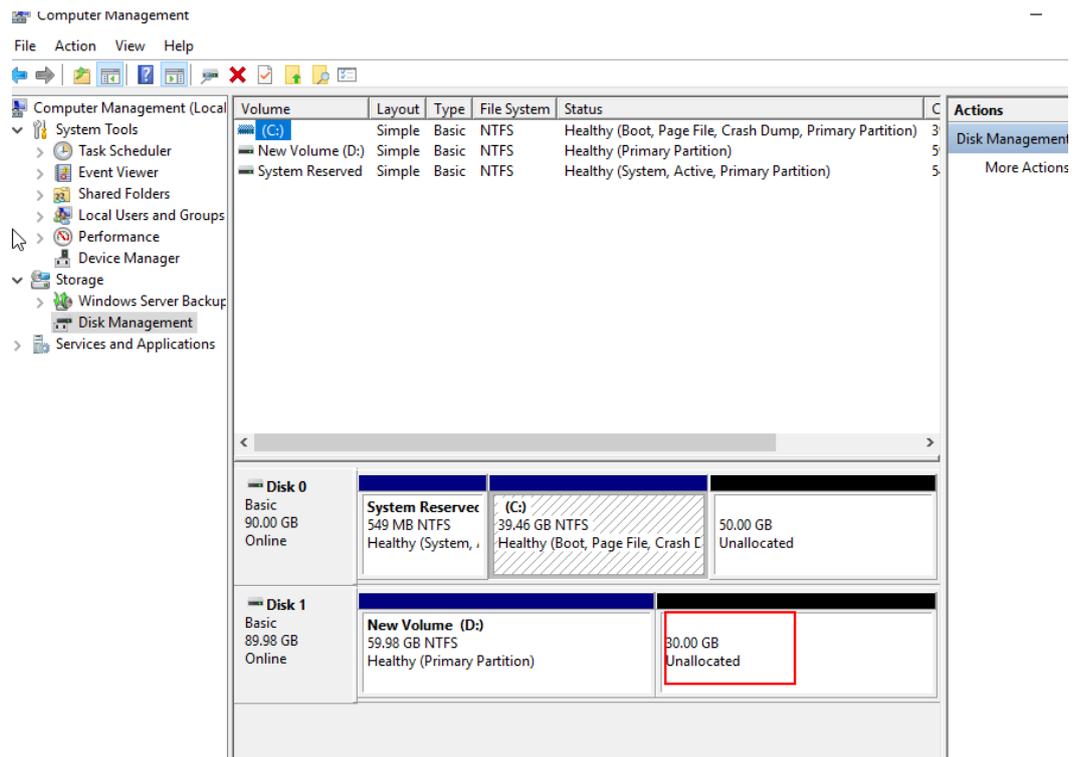
For how to log in to an ECS, see [How Do I Log In to My ECS?](#)

For how to log in to a BMS, see [Windows BMS Login Methods](#).

**Step 2** Log in to the server. On the server desktop, right-click  and choose **Disk Management**.

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

**Figure 3-2** Disk expanded but additional space not allocated



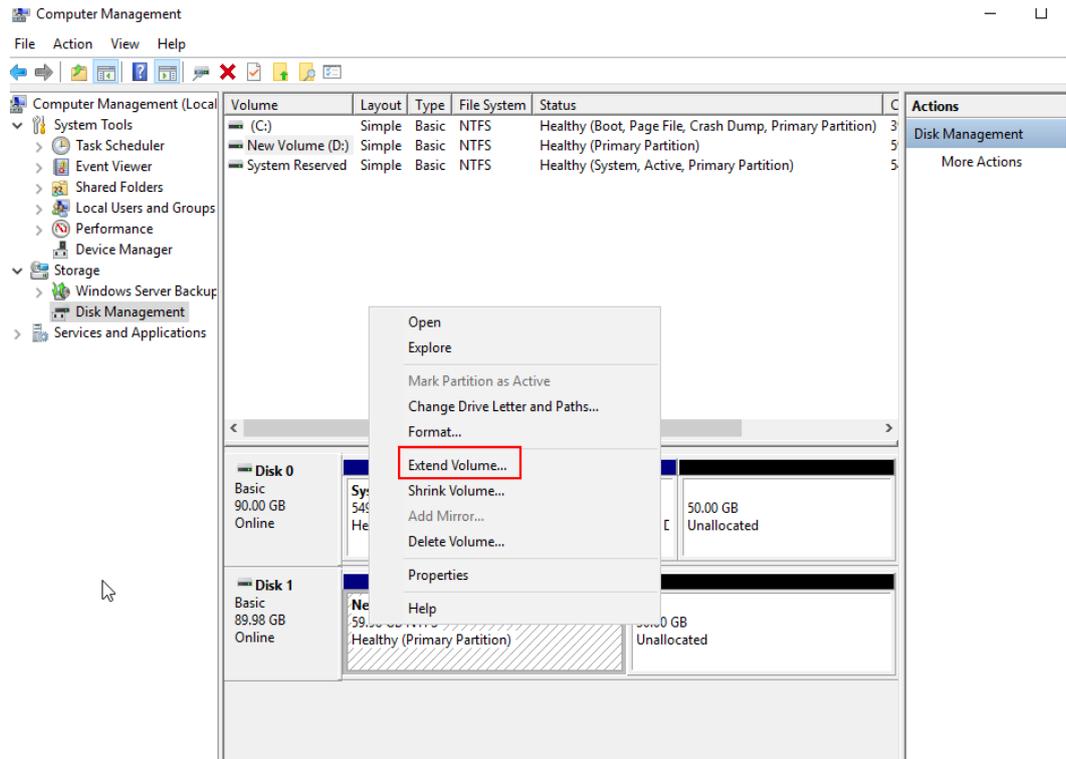
**NOTE**

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

**Step 3** On the **Disk Management** page, find the disk and volume that you want to extend. Check the size and unallocated space.

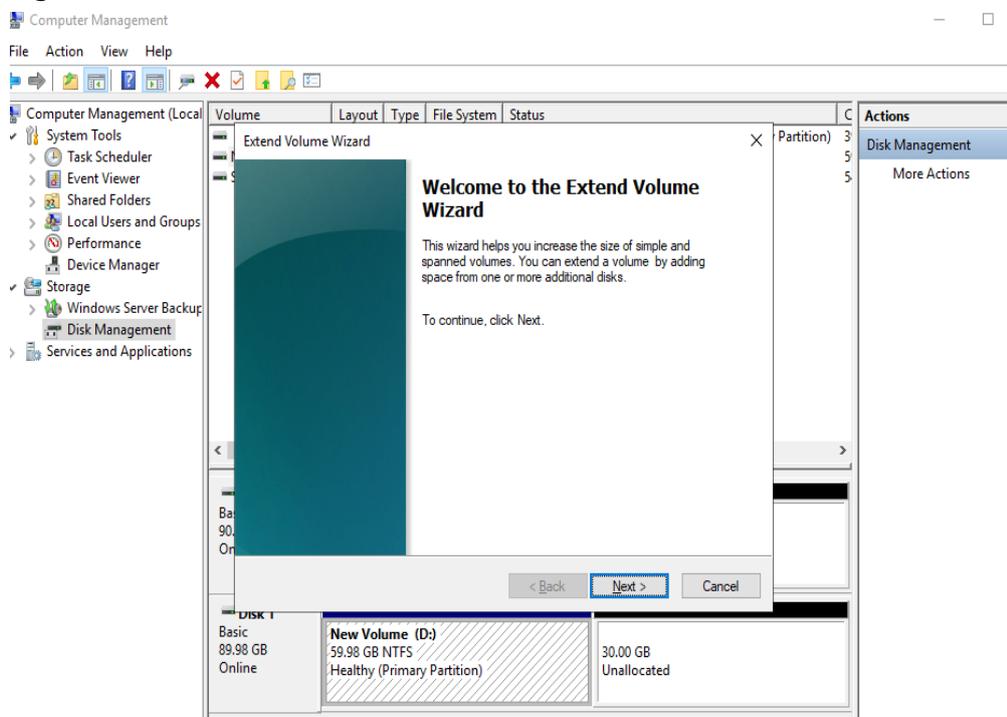
**Step 4** Right-click the volume and choose **Extend Volume** from the shortcut menu. In this example, right-click **New Volume (D:)**.

**Figure 3-3** Choosing Extend Volume



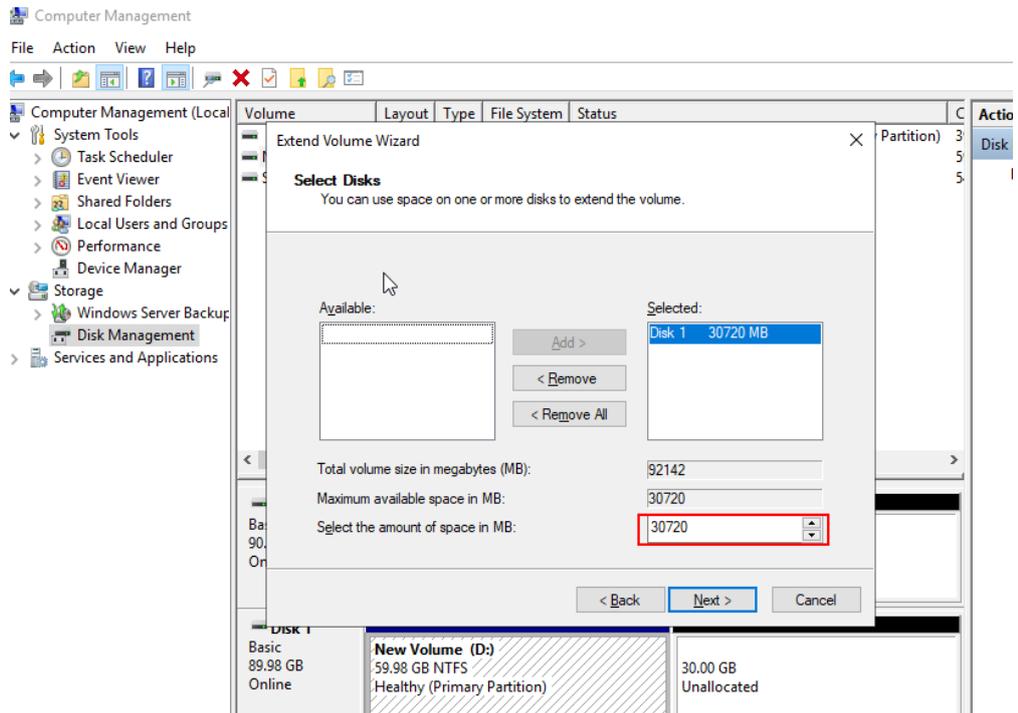
On the displayed **Extend Volume Wizard** windows, click **Next**.

**Figure 3-4** Extend Volume Wizard



**Step 5** 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.

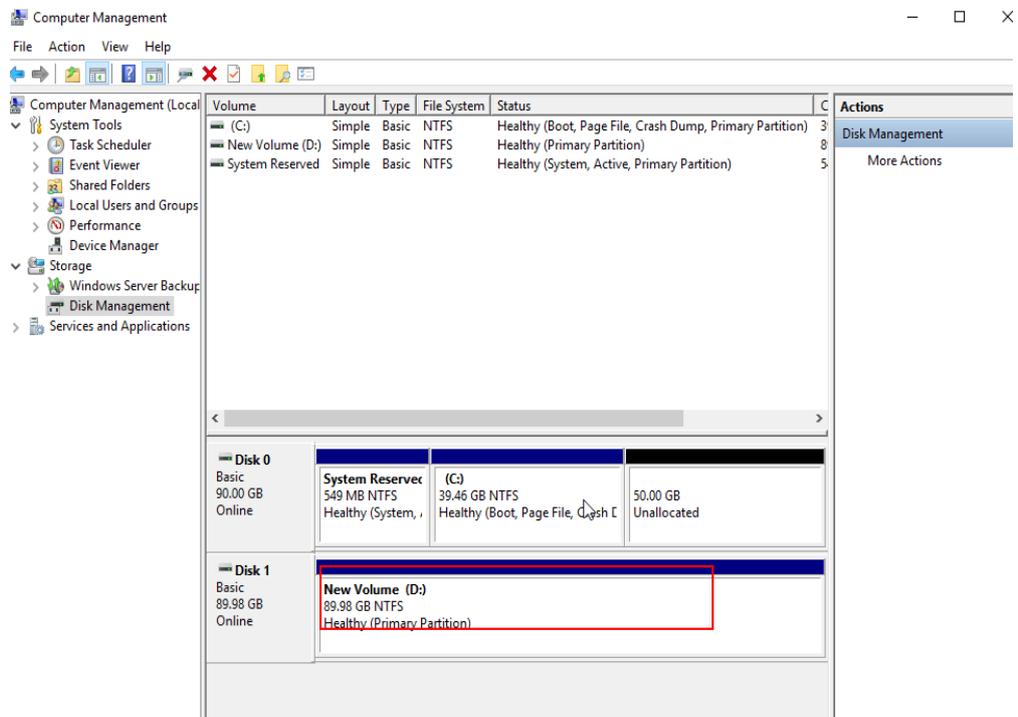
**Figure 3-5** Selecting the amount of space you want to add



**Step 6** Click **Finish**.

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

**Figure 3-6** Extension succeeded



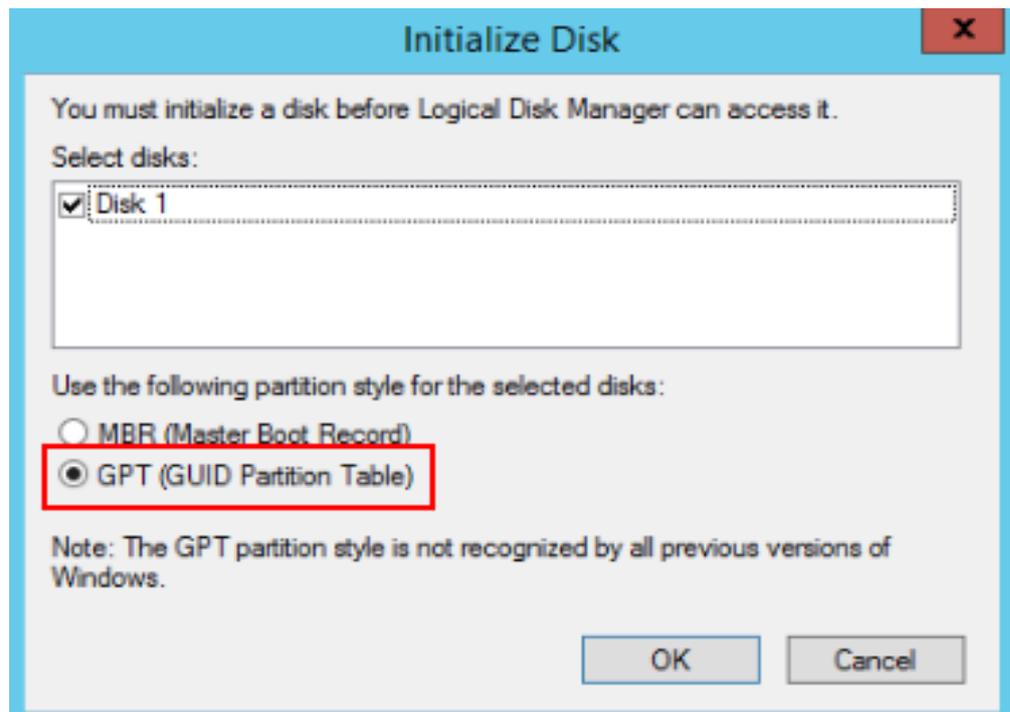
----End

## 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.

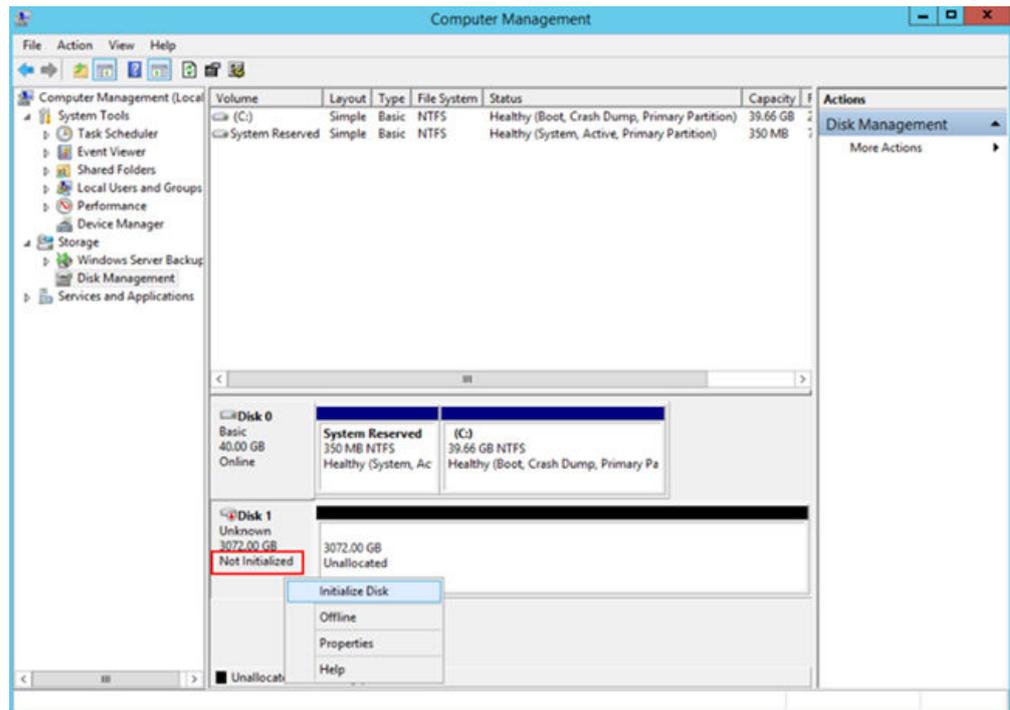
- Step 1** Log in to the server. On the server desktop, right-click  and choose **Disk Management**. The disk list is displayed.
- Disks are displayed in the right pane. If there is a disk that is not initialized, the system will prompt you with the **Initialize Disk** dialog box.

Figure 3-7 Disk list



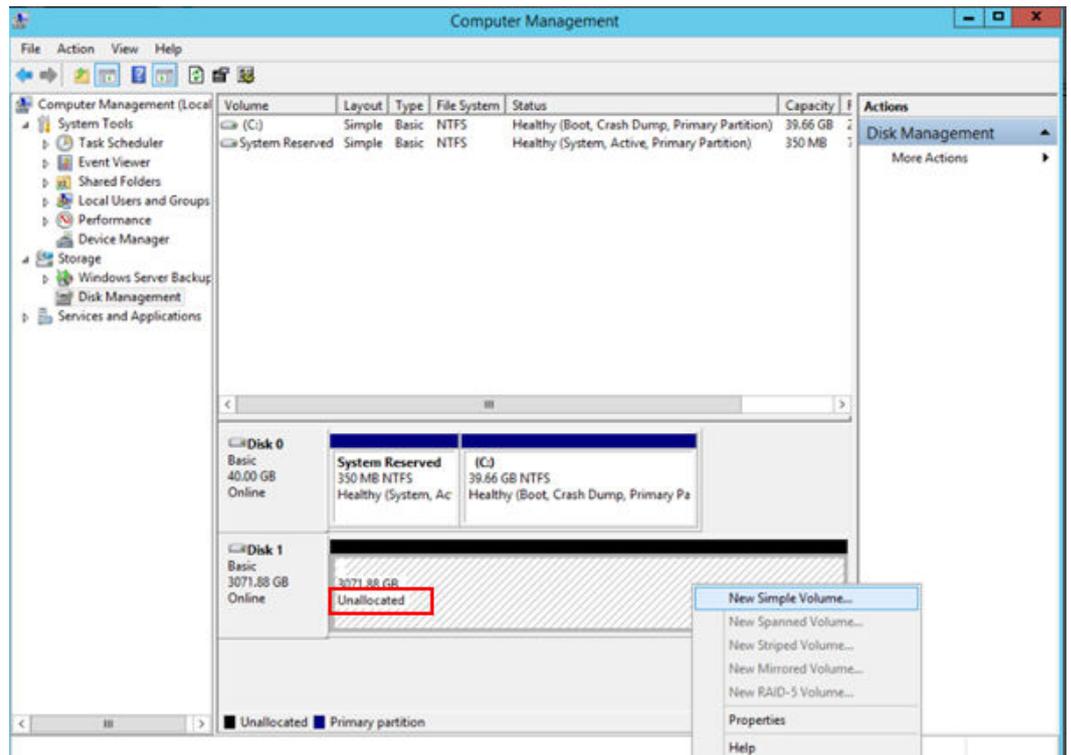
- If the **Initialize Disk** dialog box is not prompted up and the disk has no partitions (entire disk shown as **Unallocated**), right-click the area where the to-be-initialized disk is and choose **Initialize Disk** from the shortcut menu.

Figure 3-8 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 2](#) and subsequent steps.
  - To allocate the additional space to an existing partition and file system, go to [Extending an Existing Partition](#).

**Figure 3-9** Disk expanded but additional space not allocated



**NOTE**

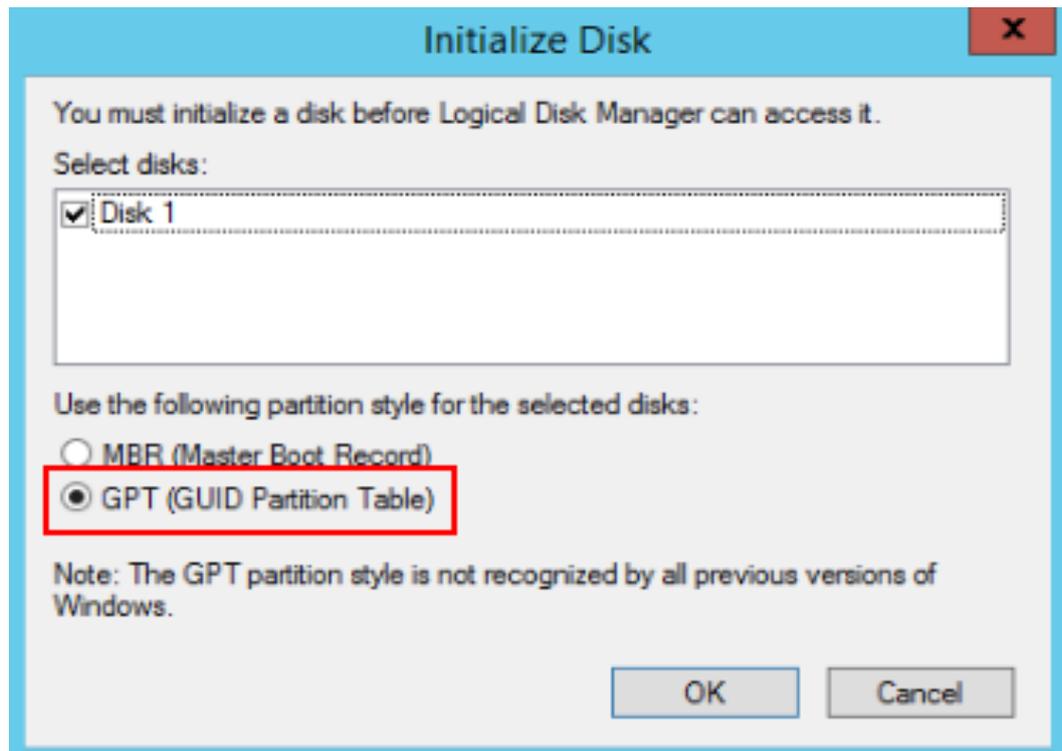
If the disk is offline, you need to **bring it online** before initializing it.

**Step 2** On the **Initialize Disk** dialog box, select **GPT (GUID Partition Table)** and click **OK** to go back to **Computer Management**.

**NOTE**

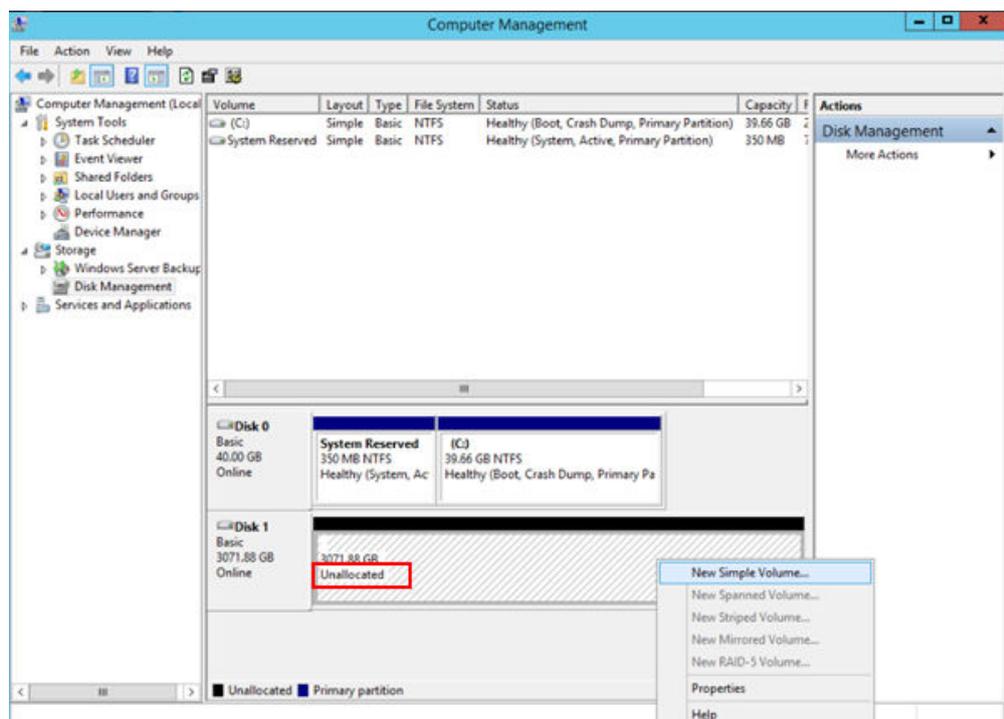
If your disk size is greater than 2 TiB or you plan to expand it to more than 2 TiB, select **GPT (GUID Partition Table)**.

Figure 3-10 Selecting GPT



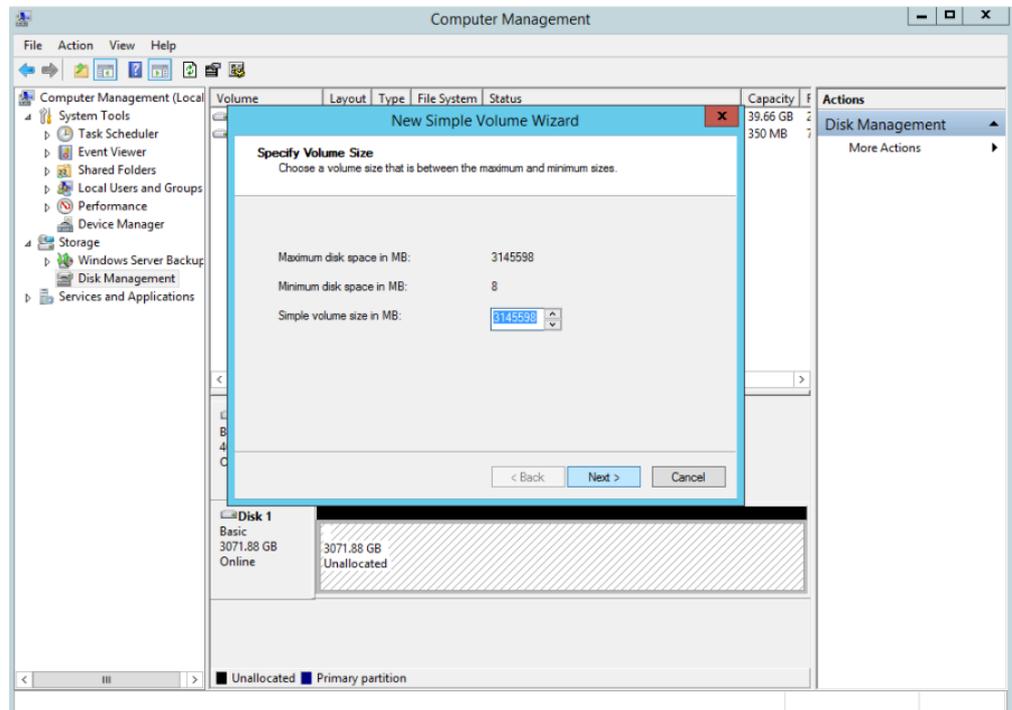
**Step 3** In the **Unallocated** area of **Disk 1**, right-click and choose **New Simple Volume** from the shortcut menu and initialize the disk as prompted.

Figure 3-11 New Simple Volume



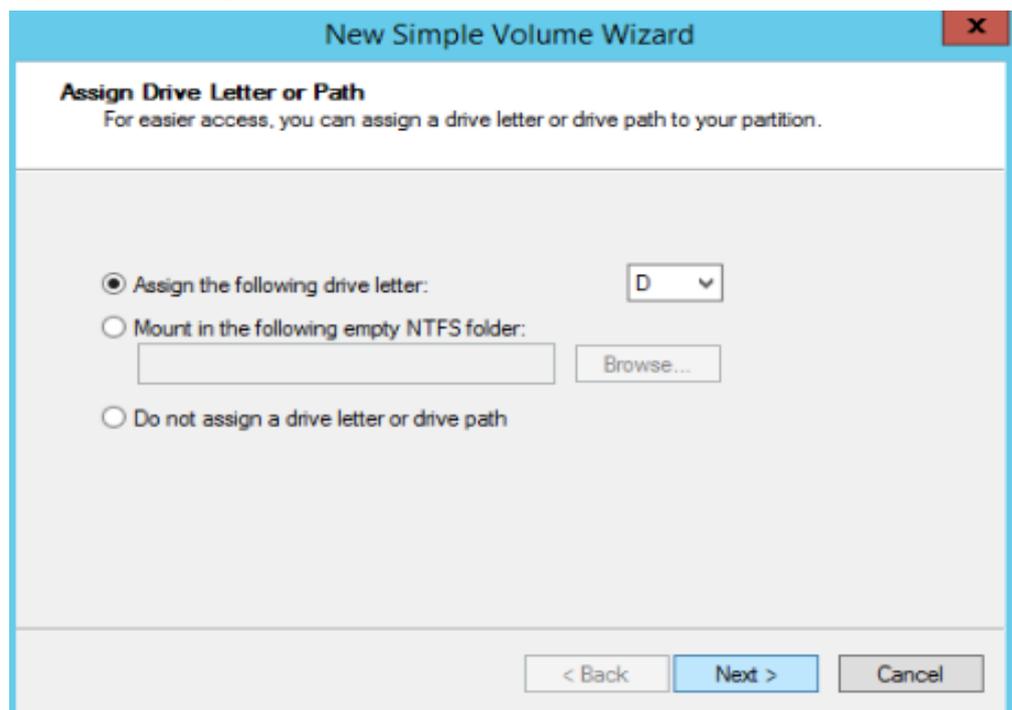
1. On the **Specify Volume Size** page, retain the default settings and click **Next**. The system uses the maximum disk space as the default volume size. You can specify a volume size as needed.

Figure 3-12 Specify Volume Size



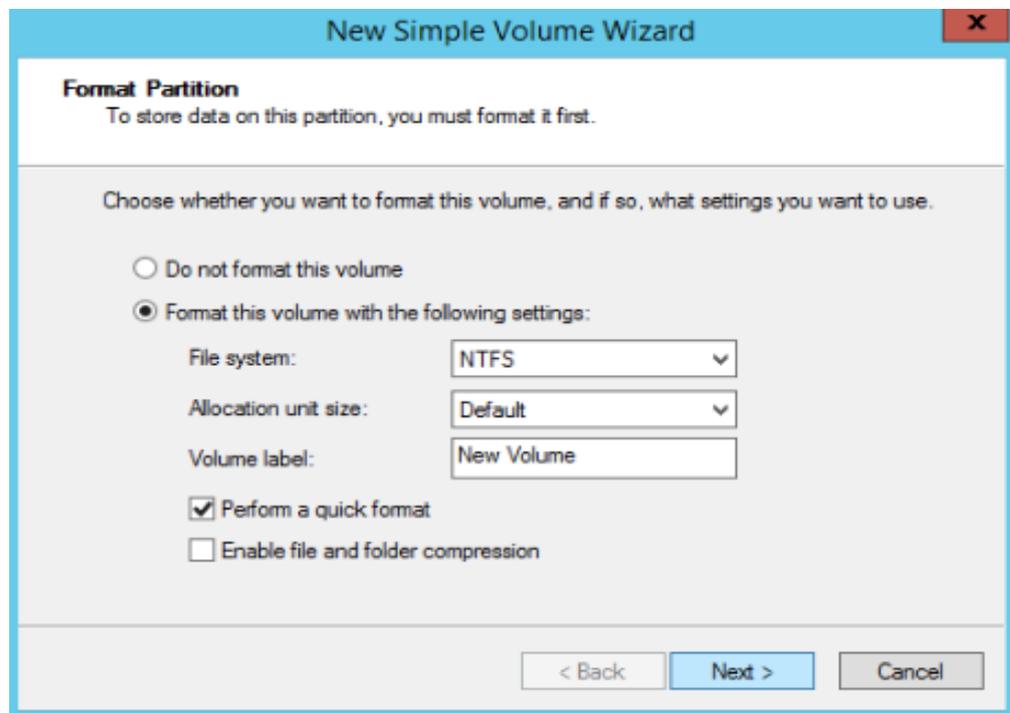
2. On the **Assign Drive Letter or Path** page, retain the default settings and click **Next**.

Figure 3-13 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 3-14** 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 3-15 Completing the New Simple Volume Wizard

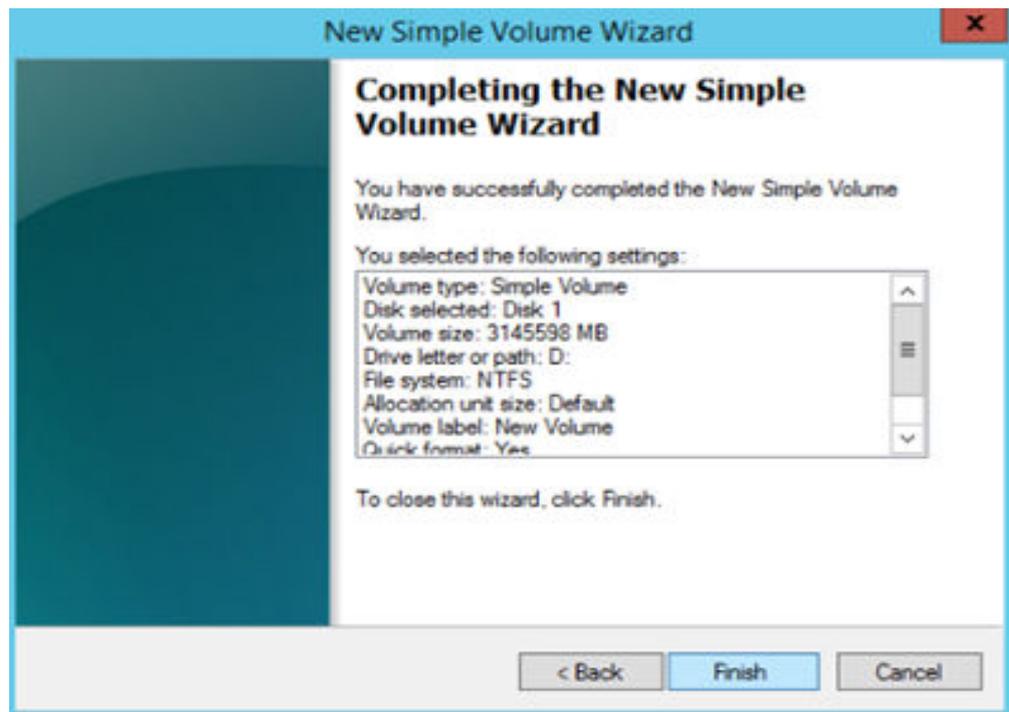
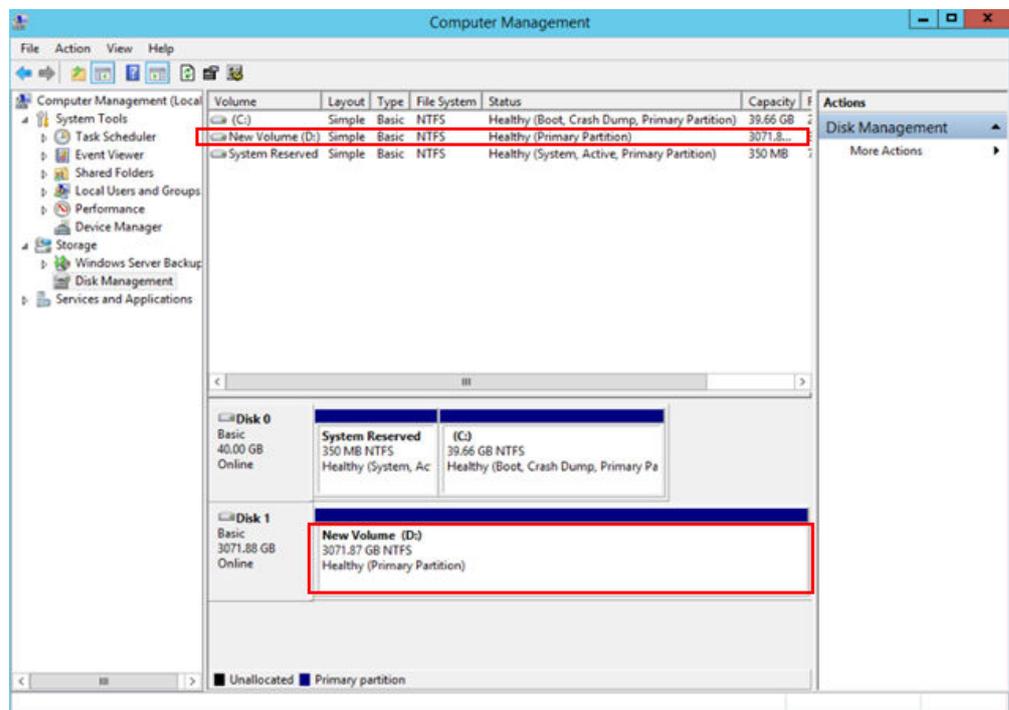


Figure 3-16 Viewing the initialization results

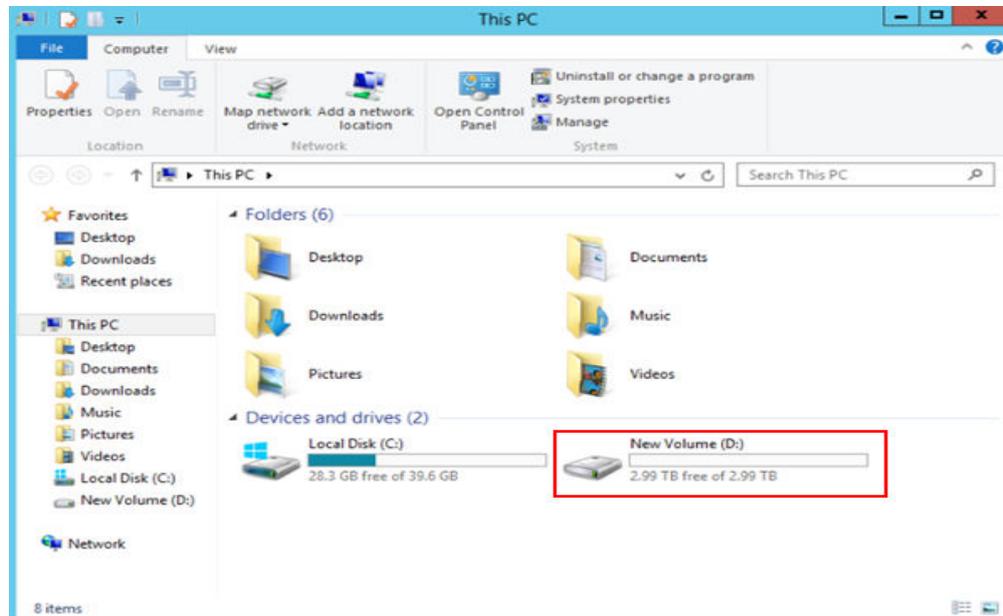


**Step 4** (Optional) Alternatively, choose **Server Manager > File and Storage Services > Volumes > Disks** to view the disk status, capacity, and partition style.

**Step 5** 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 3-17** File Explorer



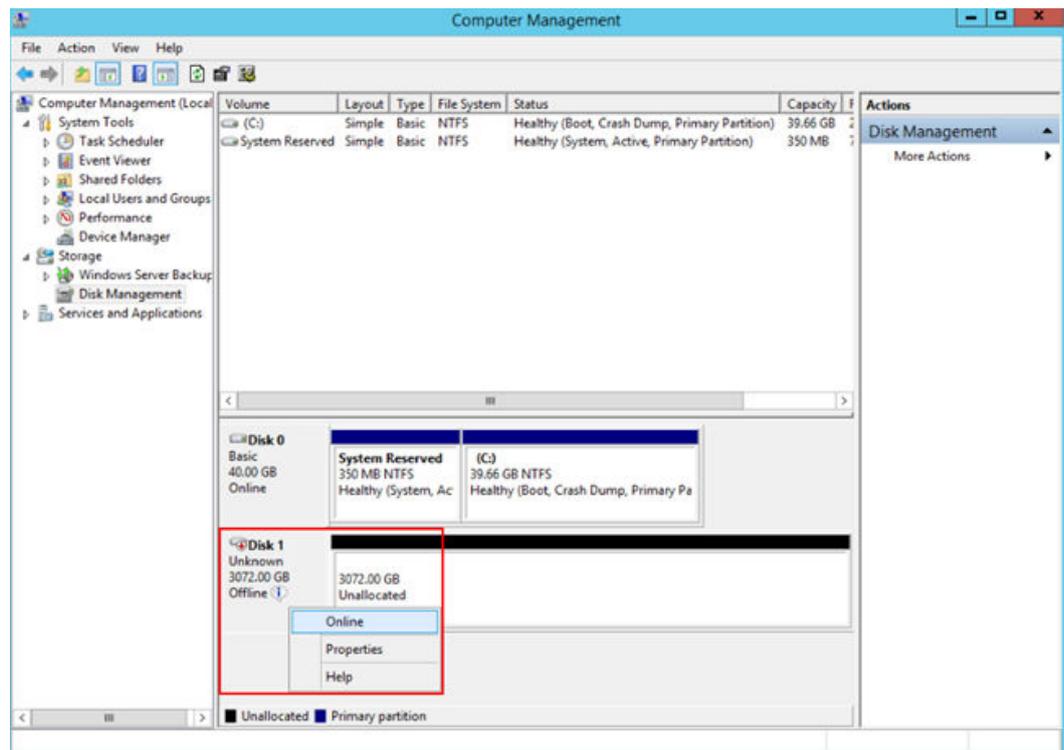
----End

## Related Operations

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

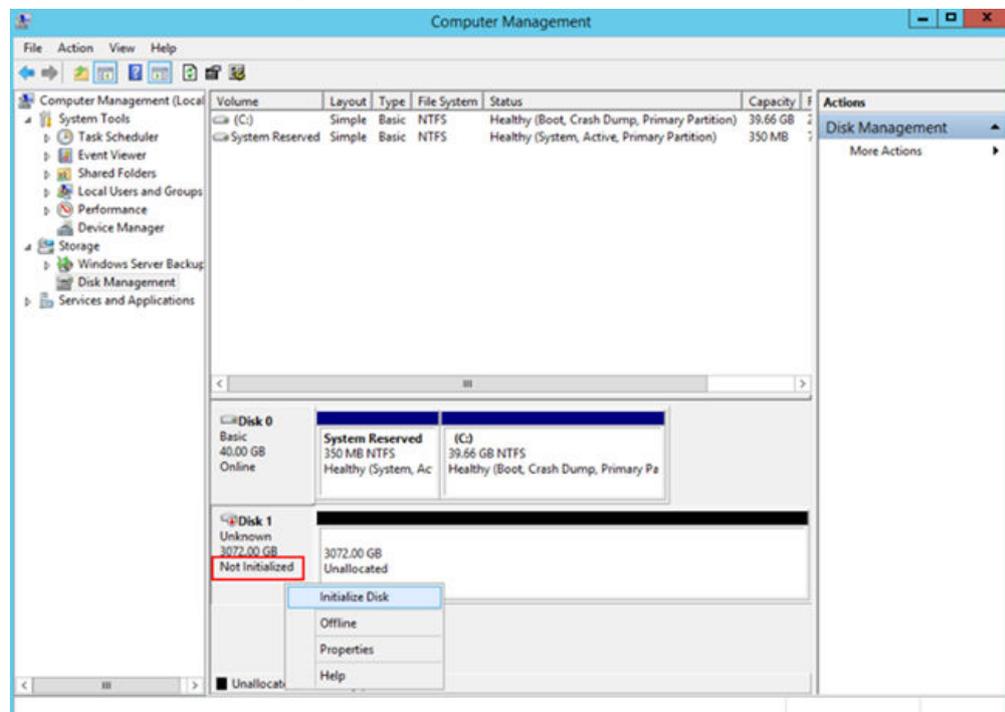
**Step 1** In the **Disk 1** area, right-click and choose **Online** from the shortcut menu.

Figure 3-18 Online



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

Figure 3-19 Disk brought online



----End

## Related Links

- For more expansion FAQs, see [Capacity Expansion](#).
- If the capacity expansion fails, you can roll back data using a snapshot. For details, see [Rolling Back Disk Data from a Snapshot](#).

# 4 Managing EVS Snapshots

## 4.1 EVS Snapshot Overview

### Snapshot 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.

EVS is now deploying the snapshot function commercially in regions one by one. You may see snapshots in OBT (legacy snapshots) or commercial use (standard snapshots) in some regions. [Table 4-1](#) describes their differences.

 **NOTE**

Standard snapshots are available only in some regions. You can check the supported regions on the console.

**Table 4-1** Snapshot types

Type	Description
Standard snapshot (snapshot function in commercial use)	Snapshot data is stored in OBS. When creating disks from a standard snapshot, you can change the device type (SCSI or VBD), encryption attribute, AZ, or type of the disks on the console.
Legacy snapshots (snapshot function in OBT)	Snapshot data is stored together with disk data on the disks. New legacy snapshots cannot be created in some regions. Existing legacy snapshots are retained, but you can only use them to restore data or create new disks in the same AZ, or delete them.

**Table 4-2** Comparison between standard snapshots and legacy snapshots

Dimension	Function	Description	Legacy Snapshot	Standard Snapshot
Key functions	Saving to object storage	By default, data is saved to object storage for redundancy.	No. Data is saved in EVS.	Yes
	Number of snapshots	The maximum number of snapshots that you can create for an EVS disk.	7	256
	Incremental snapshots	The first snapshot is a full snapshot. Subsequent snapshots are incremental snapshots.	Yes	Yes
	Data rollback to disk	If data is lost due to a system fault or misoperation, you can roll back disk data from a snapshot to quickly restore data.	Yes	Yes
	Creating new disks	You can create disks from a snapshot to quickly copy the snapshot data to disks.	Yes. You can only create disks in the AZ of the snapshot.	Yes. You can create disks in a different AZ of the snapshot.
	Billing	Pricing policies	Free	Billed based on the storage usage of snapshot chains and the usage period. For details, see <a href="#">Billing for EVS Snapshots</a> .
Data protection and disaster recovery	Snapshot encryption	Snapshots can be encrypted using algorithms such as AES 256.	Yes	Yes
	Instant Snapshot Restore	Snapshots can be used to perform operations, such as creating new disks, before the snapshot data upload is complete.	Yes	Yes

Dimension	Function	Description	Legacy Snapshot	Standard Snapshot
Lifecycle management	Snapshot consistent group	A snapshot consistency group allows you to create snapshots for multiple disks at the same time point to ensure application consistency.	No	Yes

**Table 4-3** Snapshot-related operations

Operation	Description	Reference
Creating snapshots	You can create a snapshot or snapshot consistency group (taking snapshots for multiple disks at the same time) to save the disk data at a specific time. <b>NOTE</b> Snapshots are read-only. After snapshots are created, data in the snapshots cannot be modified.	<a href="#">Creating an EVS Snapshot</a> <a href="#">Creating a Snapshot Consistency Group</a>
Rolling back data	If data on a disk is incorrect or damaged, you can roll back data from a snapshot to the source disk.	<a href="#">Rolling Back Disk Data from a Snapshot</a> <a href="#">Rolling Back Disk Data Using a Snapshot Consistency Group</a>
Creating disks from a snapshot	You can create disks from a snapshot to quickly copy the snapshot data to disks.	<a href="#">Creating a Disk from a Snapshot</a>
Using Instant Snapshot Restore	After you create standard snapshots, you can only use them until the system has uploaded the snapshot data to OBS. As more data is saved to an EVS disk, the upload takes a longer time. To address this issue, you can enable Instant Snapshot Restore. Then, you no longer need to wait for the upload to complete before using a standard snapshot to roll back data or create a new disk. The rollback and creation speed is fast and does not affect data integrity.	<a href="#">Enabling or Disabling Instant Snapshot Restore (for Standard Snapshots)</a>

Operation	Description	Reference
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.	<a href="#">Checking the EVS Snapshot Storage Usage (for Standard Snapshots)</a> <a href="#">Checking EVS Snapshot Details</a>
Deleting snapshots	If you no longer require certain snapshots or the snapshot quantity reaches the maximum allowed, you can delete some snapshots.	<a href="#">Deleting an EVS Snapshot</a> <a href="#">Deleting a Snapshot Consistency Group</a>

## Snapshot Usage Scenarios

The snapshot function helps address your following needs:

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

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

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

### NOTE

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

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

For example, you can create the **snapshot-01** snapshot for the **volume-01** disk in AZ1, then use this snapshot to create the **volume-02** disk in AZ2. In this way, the disk data is available in both AZ1 and AZ2.

 **NOTE**

Only standard snapshots support cross-AZ data migration.

## Constraints

**Table 4-4** Constraints

Item	Description
Snapshot quotas	<ul style="list-style-type: none"><li>You can manually create a maximum of seven legacy snapshots for a disk.</li><li>You can manually create a maximum of 256 standard snapshots for a disk, of which up to seven can have Instant Snapshot Restore enabled.</li></ul>
Disk types	<ul style="list-style-type: none"><li>When standard snapshots are created for Common I/O and High I/O disks, Instant Snapshot Restore cannot be enabled.</li><li>High I/O and Common I/O disks do not support Instant Snapshot Restore.</li><li>Instant Snapshot Restore is not supported when you create snapshot consistency groups for Common I/O and High I/O disks.</li></ul>
Snapshot retention policy	<p>The system does not automatically delete snapshots. Snapshots are deleted in the following scenarios:</p> <ul style="list-style-type: none"><li>You delete snapshots.</li><li>You delete EVS disks. Then, the legacy snapshots of the disks are automatically deleted. The standard snapshots of the disks are not affected.</li></ul>

## Snapshot Principle (Video)

When backing up disk data using snapshots, creating legacy snapshots establishes relationships between the snapshots and disk data. The snapshots are stored on the physical disks of the corresponding EVS disks. Standard snapshots are stored in OBS. They do not occupy the EVS disk space. The principles are introduced as follows:

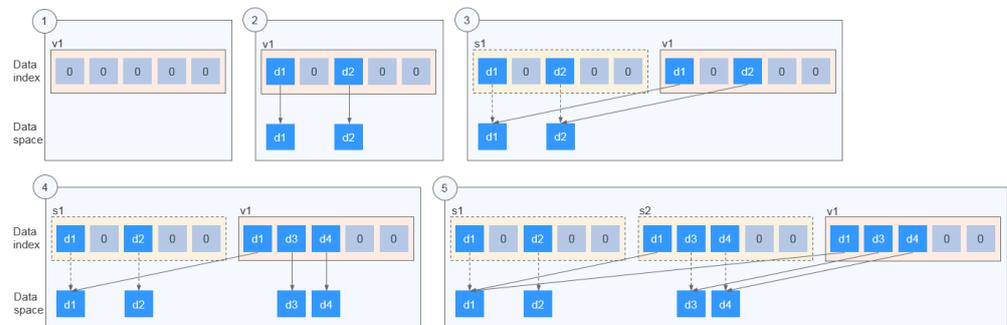
## Legacy Snapshot Principles

Creating legacy snapshots establishes relationships between snapshots and the disk data. Snapshots are stored on the physical disks that provide storage resources for EVS disks.

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

1. Disk v1 was created, which contains no data.
2. Data d1 and d2 were written to disk v1, and they were written to new spaces.
3. Snapshot s1 was created for disk v1 modified in step 2. Data d1 and d2 were not saved as another copy elsewhere. Instead, a relationship between snapshot s1 and data d1 and d2 was established.
4. Data d3 was written to disk v1, and data d2 was changed to d4. Data d3 and d4 were written to new spaces, without overwriting data d2. The relationship between snapshot s1 and data d1 and d2 remained valid. Snapshot s1 can be used to restore data if needed.
5. Snapshot s2 was created for disk v1 modified in step 4, and a relationship between snapshot s2 and data d1, d3, and d4 was established.

**Figure 4-1** 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 4-2** 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 was written to an EVS disk at 09:30 and 10:30. Snapshot 1 was created at 09:00, snapshot 2 at 10:00, and snapshot 3 at 11:00.

- At 09:00, snapshot 1 was created for the disk. This was the first time that a snapshot was created for this disk, so snapshot 1 was a full snapshot and it contained all the data on the disk, including data blocks A, B, and C. The metadata file of snapshot 1 recorded information about the disk's full data blocks: A, B, and C.
- After snapshot 1 was created, data block A was changed to A1, data block B was changed to B1, and data block D was added. Then, snapshot 2 was created at 10:00. It was an incremental snapshot. Compared with snapshot 1, data blocks A1, B1, and D were changed data blocks. The metadata file of snapshot 2 recorded information about the disk's full data blocks: A1, B1, C, and D, among which data block C was inherited from snapshot 1.
- After snapshot 2 was created, data block A1 was changed to A2, data block C was changed to C1, and data block E was added. Then, snapshot 3 was created at 11:00. It was an incremental snapshot. Compared with snapshot 2, data blocks A2, C1, and E were changed data blocks. The metadata file of snapshot 3 recorded information about the disk's full data blocks: A2, B1, C1, D, and E, among which data blocks B1 and D were 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. The storage usage of a single snapshot will not be greater than the disk capacity. As more snapshots are created for the disk, the storage usage of the snapshot chain may be greater than the disk capacity.

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

**Figure 4-3** Snapshot chain with snapshots added



Take the scenario in **Figure 4-3** 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

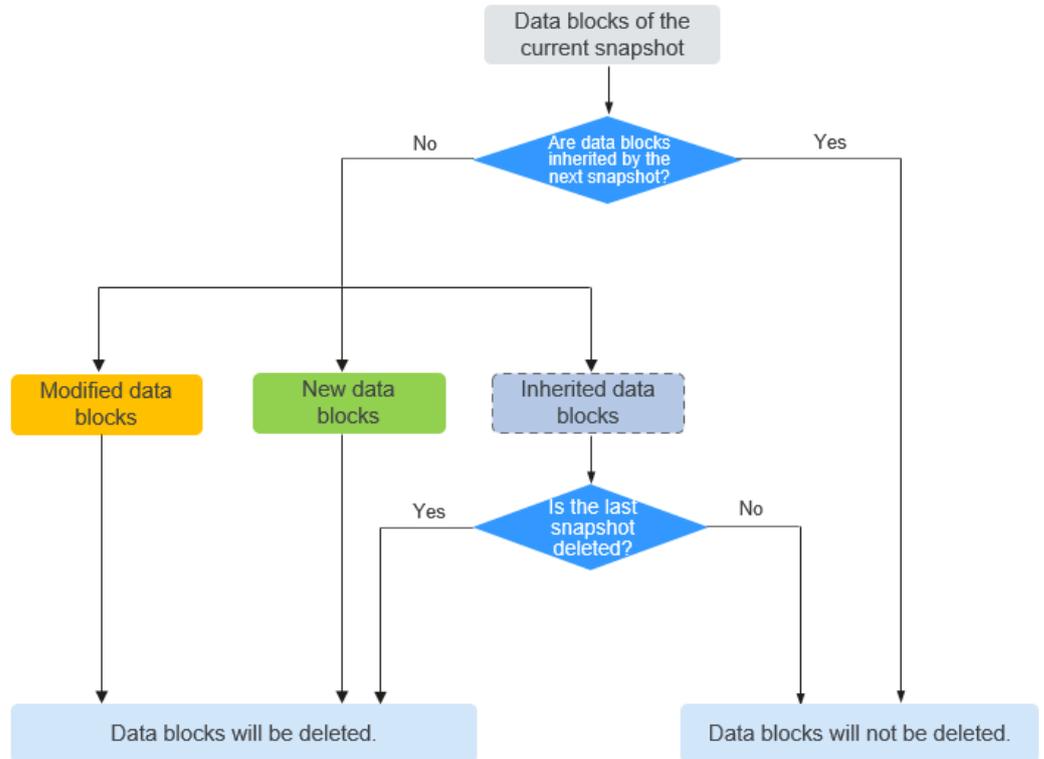
- **Snapshot chain's storage usage calculation after snapshots are deleted**

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 from which the data block is inherited 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 4-4** Snapshot data block deletion rules



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

**Figure 4-5** Snapshot chain with snapshots deleted



Take the scenario in **Figure 4-5** 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 \(for Standard Snapshots\)](#).

## Differences Between Disk Backups and Disk Snapshots

Both disk backups and disk snapshots provide redundancies for improved disk data reliability. [Table 4-5](#) lists the differences between them.

**Table 4-5** Differences between backups and snapshots

Item	Definition	Billing	Storage Solution	Data Synchronization	DR Range	Service Recovery
Disk Backups	Provides backup for Huawei Cloud servers, cloud disks, cloud databases, cloud desktops, SFS Turbo file systems, as well as cloud and local files and directories, protecting against threats such as viruses, accidental deletion, and software or hardware failures.	You need to buy disk backup vaults to store disk backups. You pay for the purchased vault capacity.	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 the same region, but can be in different AZs. Backups can be replicated across regions or clouds for data recovery.	You can use a backup to roll back data to its source disk or create a new disk. The data durability is high.

Item	Definition	Billing	Storage Solution	Data Synchronization	DR Range	Service Recovery
Legacy Snapshots	An EVS snapshot is a complete copy or image of the disk data taken at a specific time. Snapshot is a major disaster recovery approach, and you can use a snapshot to restore disk data to the time when the snapshot was created.	Free	<p>Snapshots are stored on the physical disks that provide storage resources for EVS disks. Therefore, snapshots do not use the EVS disk space.</p> <p><b>NOTE</b> Creating a backup requires a certain amount of time because data needs to be transferred to OBS. Creating a snapshot or rolling back data from a snapshot consumes less time than creating a backup.</p>	<p>A snapshot is the state of a disk at a specific point in time and is stored on the same disk. If the disk is deleted, all its snapshots will also be deleted. For example, if you reinstalled or changed the server OS, snapshots of the system disk were also automatically deleted. Snapshots of the data disks can be used as usual.</p>	<p>A snapshot and its source disk reside in the same AZ. Snapshot data cannot be replicated across regions or clouds for data recovery.</p>	<p>You can use a snapshot to roll back data to its source disk or create a new disk.</p>

Item	Definition	Billing	Storage Solution	Data Synchronization	DR Range	Service Recovery
Standard Snapshots	An EVS snapshot is a complete copy or image of the disk data taken at a specific time. Snapshot is a major disaster recovery approach, and you can use a snapshot to restore disk data to the time when the snapshot was created.	Snapshots of a disk are billed based on the storage usage of the snapshot chain and the storage usage period. A snapshot chain collects the storage space used by data blocks of all the snapshots of a disk.	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. When you delete a disk, all its snapshots will not be deleted.	A snapshot and its source disk reside in the same region, but can be in different AZs. Snapshot data cannot be replicated across regions or clouds for data recovery.	You can use a snapshot to roll back data to its source disk or create a new disk. The data durability is high.

## 4.2 Using EVS Snapshots

### 4.2.1 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**

Standard snapshots are available only in some regions. You can check the supported regions on the console.

## Constraints

**Table 4-6** Constraints

Item	Description
General constraints	<ul style="list-style-type: none"><li>• Snapshots can be created for both system disks and data disks.</li><li>• Snapshots of encrypted disks are stored encrypted, and those of non-encrypted disks are stored non-encrypted.</li></ul>
Legacy snapshots	<ul style="list-style-type: none"><li>• You can manually create a maximum of seven legacy snapshots for a disk.</li><li>• Huawei Cloud reserves the right to restrict user snapshots created during OBT.</li><li>• The enterprise project of a snapshot is the same as that of the snapshot's source disk.</li></ul>

Item	Description
Standard snapshots	<ul style="list-style-type: none"><li>• You can manually create a maximum of 256 standard snapshots for a disk, of which up to seven can have Instant Snapshot Restore enabled.</li><li>• 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.</li><li>• When a snapshot is being created, do not perform operations, such as stopping or restarting the compute instance (ECS or any other), that may change the instance status, or the snapshot creation will fail.</li><li>• Standard snapshots cannot be created for the disks in edge AZs. For details about the differences between edge AZs and general AZs, see <a href="#">CloudPond User Guide</a>.</li><li>• When standard snapshots are created for Common I/O and High I/O disks, Instant Snapshot Restore cannot be enabled.</li><li>• 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 amount of changed data compared with each last snapshot. The more the changed data, the longer the time required.</li><li>• If data on a disk is rolled back from a snapshot, the next standard snapshot created for this disk will be a full snapshot.</li><li>• During the creation of a standard snapshot, any incremental data written to the disk will not be backed up to the snapshot created.</li><li>• During the creation of a snapshot, deleting the snapshot's source disk will affect the snapshot creation. Delete the source disk after the snapshot is created.</li><li>• You can add a maximum of 20 tags to a snapshot.</li><li>• Tag keys of the same snapshot must be unique.</li></ul>

## Impacts on Performance

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.

## Billing

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

Standard snapshots are billed based on the storage usage of snapshot chains and the usage period. For details, see [Billing for EVS Snapshots](#).

## Prerequisites

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

## Legacy Snapshots

Legacy snapshots are stored on the physical disks that provide storage resources for EVS disks. They do not use the EVS disk space. If a disk is deleted, all legacy snapshots of this disk will also be deleted.

## Creating a Legacy Snapshot on the Disks Page

**Step 1** Sign in to the [EVS console](#).

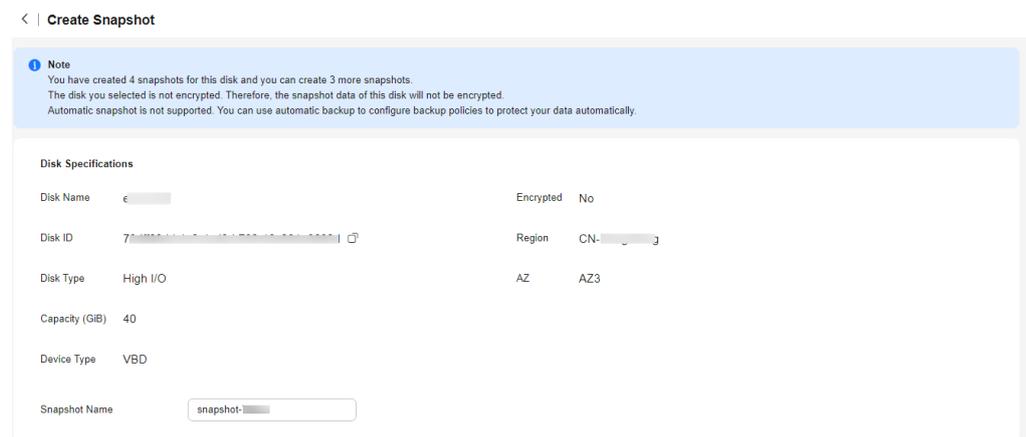
**Step 2** In the disk list, locate the target disk and click **Create Snapshot** in the **Operation** column.

Configure the snapshot parameter according to [Table 4-7](#).

**Table 4-7** Snapshot parameter

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

**Figure 4-6** Create Snapshot



**Step 3** Click **Create Now**.

**Step 4** Go back to the **Snapshots** page and view the created snapshot in the snapshot list.

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

----End

## Creating a Legacy Snapshot on the Snapshots Page

**Step 1** Sign in to the [EVS console](#).

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

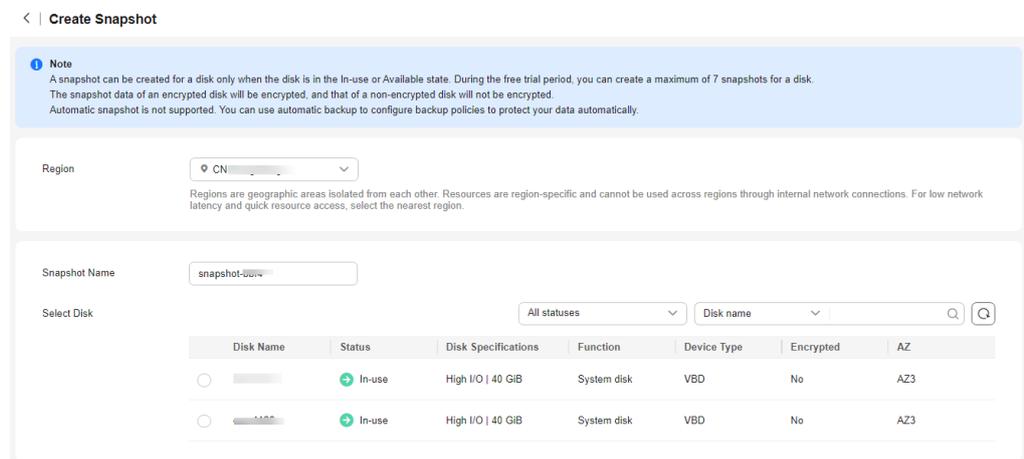
On the **Snapshots** page, click **Create Snapshot**.

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

**Table 4-8** Snapshot parameters

Parameter	Description	Example Value
Region	After you select a region, disks in the selected region will be displayed for you to choose from.	-
Snapshot Name	<ul style="list-style-type: none"> <li>It can contain only letters, digits, periods (.), hyphens (-), and underscores (_).</li> <li>It can contain a maximum of 64 characters.</li> </ul>	snapshot-01
Select Disk	Select a disk based on which the snapshot will be created.	volume-01

**Figure 4-7** Create Snapshot



**Step 3** Click **Create Now**.

**Step 4** Go back to the **Snapshots** page and view the created snapshot in the snapshot list.

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

----End

## Standard Snapshots

Standard snapshots are stored in OBS, instead of on disks. They can be used to restore data when the disk is damaged.

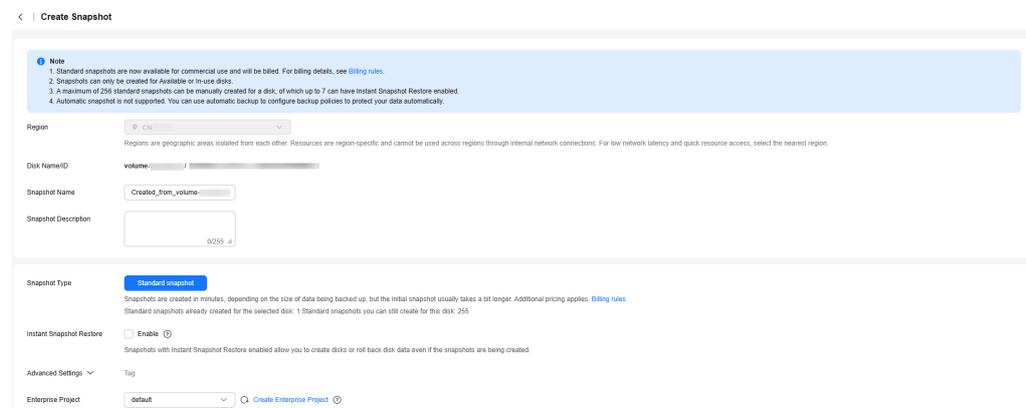
### Creating a Standard Snapshot on the Disks Page

**Step 1** Sign in to the [EVS console](#).

**Step 2** In the disk list, locate the target disk and click **Create Snapshot** in the **Operation** column.

Configure the snapshot parameters according to [Table 4-9](#).

**Figure 4-8** Create Snapshot



**Table 4-9** Snapshot parameters

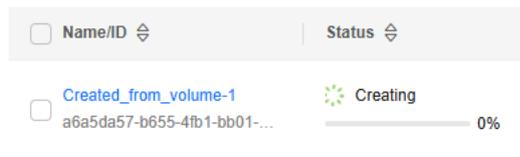
Parameter	Description	Example Value
Region	The region to which the snapshot belongs. The snapshot must be in the same region as its source disk.	CN-Hong Kong
Disk Name/ID	The name and ID of the source disk. You can obtain them on the disk list page.	-
Snapshot Name	<ul style="list-style-type: none"> <li>It can contain only letters, digits, periods (.), hyphens (-), and underscores (_).</li> <li>It can contain a maximum of 64 characters.</li> </ul>	snapshot-01Created_from_evstest
Snapshot Description	Optional The description can contain up to 255 characters.	-

Parameter	Description	Example Value
Snapshot Type	<p>The type of the snapshot. Only standard snapshot is supported currently.</p> <p>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.</p>	Standard snapshot
Instant Snapshot Restore	<p>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 <a href="#">Enabling or Disabling Instant Snapshot Restore (for Standard Snapshots)</a>.</p>	-
Advanced Settings > Tag	<p>Optional</p> <p>You can add tags when creating standard snapshots. Tags can help you identify, classify, and search for your snapshots.</p> <p>A tag consists of a tag key and a tag value.</p> <ul style="list-style-type: none"> <li>• A tag key cannot start or end with a space, or start with <b>_sys_</b>. It can contain a maximum of 128 characters and contain letters, digits, spaces, and the following special characters: <code>_.:=+@</code></li> <li>• A tag value can contain a maximum of 255 characters and contain letters, digits, spaces, and the following special characters: <code>./:=+@</code></li> </ul>	-

**Step 3** Click **Create Now**.

**Step 4** Go back to the **Snapshots** page and view the snapshot creation progress in the snapshot list.

When the upload progress reaches 100%, the snapshot is created.



----End

## Creating a Standard Snapshot on the Snapshots Page

**Step 1** Sign in to the [EVS console](#).

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

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

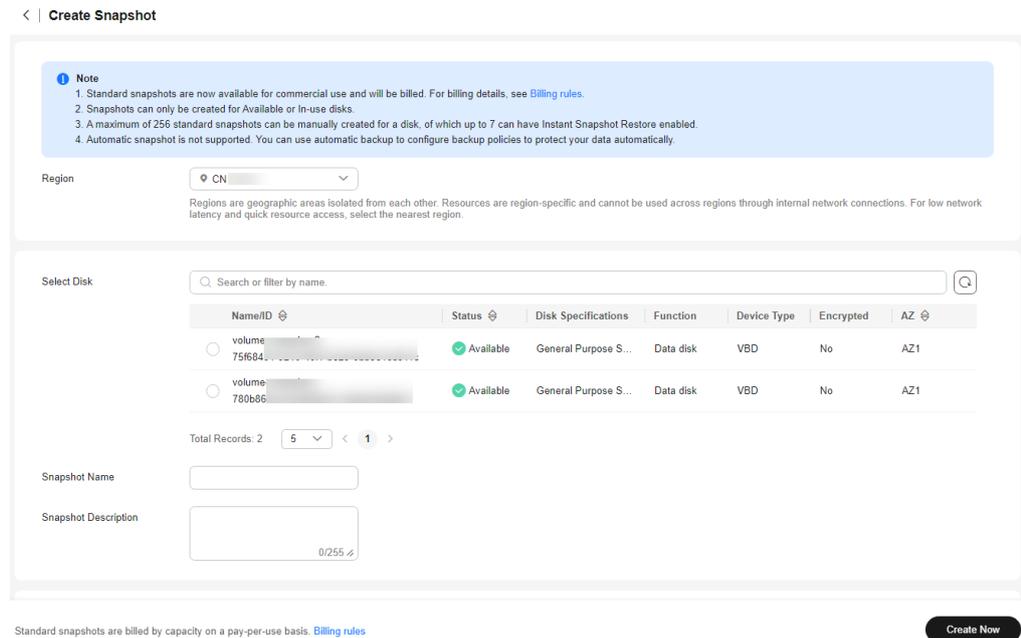
Configure the snapshot parameters according to [Table 4-10](#).

**Table 4-10** Snapshot parameters

Parameter	Description	Example Value
Region	After you select a region, disks in the selected region will be displayed for you to choose from.	CN-Hong Kong
Select Disk	Select a disk for which you want to create a snapshot.	-
Snapshot Name	The name can contain a maximum of 64 characters.	snapshot-01Created_from_evstest
Snapshot Description	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 <a href="#">Enabling or Disabling Instant Snapshot Restore (for Standard Snapshots)</a> .	Enable

Parameter	Description	Example Value
Advanced Settings > Tag	<p>Optional</p> <p>You can add tags when creating standard snapshots. Tags can help you identify, classify, and search for your snapshots.</p> <p>A tag consists of a tag key and a tag value.</p> <ul style="list-style-type: none"> <li>• A tag key cannot start or end with a space, or start with <b>_sys_</b>. It can contain a maximum of 128 characters and contain letters, digits, spaces, and the following special characters: <code>_:+=+@</code></li> <li>• A tag value can contain a maximum of 255 characters and contain letters, digits, spaces, and the following special characters: <code>_:./=+-@</code></li> </ul>	-
Enterprise Project	<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 <b>default</b>.</p>	default

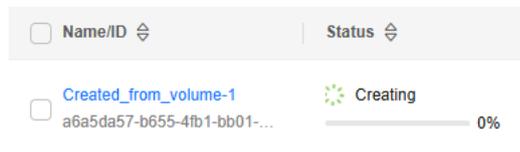
**Figure 4-9 Create Snapshot**



**Step 3** Click **Create Now**.

**Step 4** Go back to the **Snapshots** page and view the snapshot creation progress in the snapshot list.

When the upload progress reaches 100%, the snapshot is created.



----End

## Related Links

- You can refer to [Rolling Back Disk Data from a Snapshot](#) to restore data on EVS disks.
- After a snapshot is created, you can [check the snapshot bill](#) in Billing Center.
- You are advised to delete unnecessary snapshots on a regular basis to reduce snapshot costs. For details, see [Deleting an EVS Snapshot](#).
- To check more snapshot FAQs, see [Snapshot](#).
- To increase the snapshot quota, see [Managing EVS Quotas](#).

## 4.2.2 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.

---

 **CAUTION**

After the data rollback, all the changes made after the snapshot was created will be lost.

---

## Constraints

- Snapshot data can only be rolled back to source EVS disks. Rollback to a different disk is not possible.
- If the snapshot status is **Creating**, it cannot be used to roll back disk data.
- If the standard snapshot status is **Available** and the snapshot has Instant Snapshot Restore enabled, you can use the snapshot to roll back data even if data upload is in progress.
- 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.

## Prerequisites

- At least **a snapshot has been created** for the disk.
- 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 it first.

## Rolling Back Disk Data from a Snapshot

**Step 1** Sign in to the [EVS console](#).

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

The **Snapshots** page is displayed.

**Step 3** In the snapshot list, locate the target snapshot and click **Roll Back Disk** in the **Operation** column.

**Step 4** 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 5** 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

## Follow-up Operations

- A rollback operation restores disk data to the time when the snapshot was taken. Any changes made concerning system configuration, software, files and others after that time will be lost. Reconfigure them after the rollback.

- If you expand the disk capacity after creating a snapshot for a data disk, the partition and file system created on the expanded capacity will be lost after you roll back data from this snapshot. You need to expand the capacity again. For details, see [Extending Disk Partitions and File Systems \(Linux\)](#) and [Extending Disk Partitions and File Systems \(Windows\)](#).

## Related Links

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

## 4.2.3 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 creating a disk or specifying parameter **Create from snapshot** when creating a new disk.

### Constraints

Table 4-11 Constraints

Item	Description
Legacy snapshots	<ul style="list-style-type: none"><li>• Batch disk creation from a snapshot is not supported.</li><li>• A disk created from a snapshot has the same device type (SCSI or VBD), encryption attribute, AZ, region, and disk type as the snapshot's source disk.</li><li>• A snapshot whose name starts with <b>autobk_snapshot_vbs_</b>, <b>manualbk_snapshot_vbs_</b>, <b>autobk_snapshot_csbs_</b>, or <b>manualbk_snapshot_csbs_</b> is automatically generated during backup. Such a snapshot can only be viewed. It cannot be used to create new disks.</li></ul>
Standard snapshots	<ul style="list-style-type: none"><li>• A standard snapshot with Instant Snapshot Restore disabled can only be used to create disks when the snapshot status is <b>Available</b>.</li><li>• 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.</li><li>• You can use a standard snapshot to batch create disks after its data upload is complete.</li><li>• After a standard snapshot has been uploaded, you can change the device type (SCSI or VBD), encryption attribute, AZ, or type of the disks when using this snapshot to create disks on the console or through the V5 version API. Snapshots created using the V2 version API cannot be used across AZs.</li></ul>

Item	Description
General constraints	<ul style="list-style-type: none"><li>When you create a disk from a snapshot, the disk capacity must be greater than or equal to the snapshot size. In the condition that you do not specify a disk capacity, if the snapshot size is smaller than 10 GiB, the default capacity 10 GiB will be used as the disk capacity; if the snapshot size is greater than 10 GiB, the snapshot size will be used as the disk capacity.</li></ul>

 **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.

## Creating an EVS Disk from a Snapshot

**Step 1** Sign in to the [EVS console](#).

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

The **Snapshots** page is displayed.

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

**Step 4** Configure the disk parameters. For details, see parameter descriptions and operations provided in [Purchasing an EVS Disk](#).

 **NOTE**

To specify a disk capacity larger than the snapshot size, simply enter your desired capacity in the **Disk Specifications** area.

**Step 5** Click **Next**.

**Step 6** Confirm the configuration and click **Submit**.

**Step 7** If you are buying a yearly/monthly disk, make the payment and click **OK**.

The disk list page is displayed.

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

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

**Step 9** Follow-up operation: A disk created from a snapshot already has partitions and file systems. You need to attach the disk to an ECS and then mount the partitions before you can use the disk.

- If you choose not to attach to server during the disk creation, attach the disk to an ECS by referring to [Attaching an EVS Disk](#).
- If you choose to use existing partitions instead of re-initializing the disk, mount the partitions in the OS.
  - In Linux, mount the partitions on desired mount points and configure auto mount at system startup. For details, see the steps for mounting

partitions and configuring auto mount in [Initializing a Linux Data Disk \(Less Than or Equal to 2 TiB\)](#).

- In Windows, no further action is required. You can simply use the existing partitions.
- If you choose to re-initialize the disk, do as follows:
  - Note that 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 details about how to initialize a disk, see [Initializing EVS Data Disks](#).

----End

## Related Links

- To check more snapshot FAQs, see [Snapshot](#).
- To buy new data disks, see [Purchasing an EVS Disk](#).

## 4.2.4 Enabling or Disabling Instant Snapshot Restore (for Standard Snapshots)

### Scenarios

After you create standard snapshots, you can only use them until the system has uploaded the snapshot data to OBS. As more data is saved to an EVS disk, the upload takes a longer time. To address this issue, you can enable Instant Snapshot Restore. Then, you no longer need to wait for the upload to complete before using a standard snapshot to roll back data or create a new disk. The rollback and creation speed is fast and does not affect data integrity.

### Billing

You can use Instant Snapshot Restore for free. The snapshots you create are billed based on the storage usage of snapshot chains and the usage period. For details, see [Billing for EVS Snapshots](#).

### Constraints

- High I/O and Common I/O disks do not support Instant Snapshot Restore.
- 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 not be 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** Sign in to the [EVS console](#).
- Step 2** In the navigation pane on the left, choose **Elastic Volume Service > Snapshots**.  
The **Snapshots** page is displayed.
- Step 3** In the snapshot list, locate the target snapshot and click **Disable Instant Snapshot Restore** in the **Operation** column.
- Step 4** In the disabled dialog box, click **OK**.  
----End

## Related Links

- To roll back data after Instant Snapshot Restore is enabled, see [Rolling Back Disk Data from a Snapshot](#).
- To check more snapshot FAQs, see [Snapshot](#).

## 4.2.5 Checking the EVS Snapshot Storage Usage (for Standard Snapshots)

### 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.

### Constraints

- The size of a single snapshot is equal to or 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

The total snapshot storage usage of an EVS disk is calculated by snapshot chain. A snapshot chain measures the storage space used by data blocks of all the snapshots of a disk. For details about the measurement principles, see [Calculating the Standard Snapshot Storage Usage](#).

- Step 1** Sign in to the [EVS console](#).
- Step 2** Locate the disk that you want to check its total snapshot usage and click  to copy the disk ID.

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

The **Snapshots** page is displayed.

**Step 4** Click the **Snapshot Chains** tab.

**Step 5** In the search box above the list, select Disk ID, paste the copied disk ID, and click .

**Step 6** View the capacity displayed in the **Snapshot Storage Usage** column.

**Step 7** (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

Perform the following operations on the console to view the total snapshot storage usage in a specified period in the current region.

**Step 1** Sign in to the [EVS console](#).

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

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

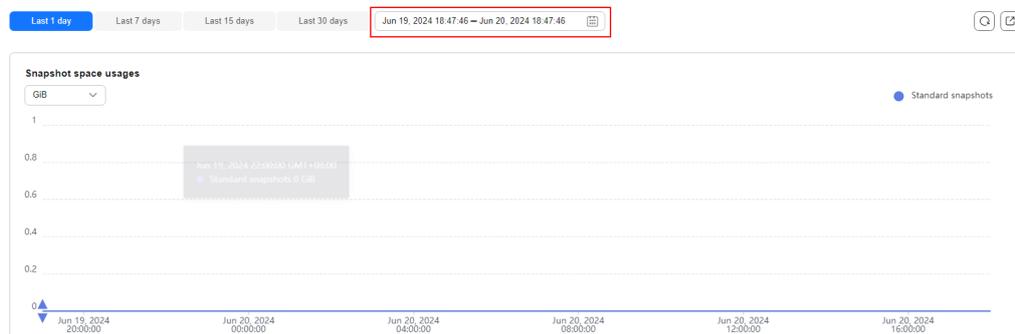
The **Snapshots** page is displayed.

**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 4-10** Querying the total snapshot storage usage in a specified period



----End

## Related Links

- Standard snapshots are billed by snapshot chain space usage based on the usage duration. For details, see [Billing for EVS Snapshots](#).

- You are advised to periodically delete snapshots that are no longer used. This helps you avoid unnecessary billing on the snapshots. For details, see [Deleting an EVS Snapshot](#).

## 4.2.6 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** Sign in to the [EVS console](#).

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

The **Snapshots** page is displayed.

**Step 3** In the snapshot list, find the target snapshot and click the snapshot name. On the displayed **Basic Information** page, view the snapshot details.

----End

### Snapshot Statuses

An EVS snapshot has several statuses. [Table 4-12](#) lists the EVS snapshot statuses, the meaning of each status, and the operations a snapshot in each status allows.

**Table 4-12** Snapshot status details

Snapshot Status	Description	Allowed Operation
Creating	The snapshot is being created.	None
Available	The snapshot is successfully created.	<ul style="list-style-type: none"><li>Creating EVS disks using snapshots</li><li>Deleting snapshots</li><li>Rolling back data to EVS disks using snapshots</li></ul>
Deleting	The snapshot is being deleted.	None
Error	An error occurs when you try to create a snapshot.	Deleting snapshots
Deletion failed	An error occurs when you try to delete a snapshot.	Deleting snapshots

Snapshot Status	Description	Allowed Operation
Rolling back	The snapshot is being used to roll back data. <b>NOTE</b> <ul style="list-style-type: none"><li>When you roll back data from a snapshot, you can only roll back data to the source EVS disk. Rollback to a specific disk is not supported.</li><li>A snapshot can only be used for rollback when its source disk is in the <b>Available</b> or <b>Rollback failed</b> state.</li></ul>	None
Backing up	This status only shows up for 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. <b>NOTE</b> Temporary snapshots are created through the CBR service. Do not perform any operation on these snapshots.	None

## Related Links

- To check more snapshot FAQs, see [Snapshot](#).
- To check the snapshot space usage, see [Checking the EVS Snapshot Storage Usage \(for Standard Snapshots\)](#).

## 4.2.7 Deleting an EVS Snapshot

### Scenarios

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

### Prerequisites

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

## Constraints

**Table 4-13** Constraints

Item	Description
Legacy snapshots	<ul style="list-style-type: none"><li>• If a snapshot's source disk is deleted, all legacy snapshots of this disk are also deleted.</li><li>• If you reinstall or change the server OS, snapshots of the system disk are automatically deleted. Those of the data disks can be used as usual.</li><li>• A snapshot whose name starts with <b>autobk_snapshot_vbs_</b>, <b>manualbk_snapshot_vbs_</b>, <b>autobk_snapshot_csbs_</b>, or <b>manualbk_snapshot_csbs_</b> is automatically generated during backup. You can check details of such snapshots, but cannot delete them.</li></ul>
Standard snapshots	<ul style="list-style-type: none"><li>• Standard snapshots are not deleted even if their source disks are deleted.</li><li>• When you delete a disk whose standard snapshots have Instant Snapshot Restore enabled, the standard snapshots will not be deleted.</li><li>• If you reinstall or change the server OS, standard snapshots will not be deleted, but Instant Snapshot Restore will be disabled automatically if it has been enabled for the standard snapshots of the system disk.</li></ul>
General constraints	<ul style="list-style-type: none"><li>• If a snapshot is deleted, disks rolled back or created from this snapshot are not affected.</li></ul>

## Procedure

**Step 1** Sign in to the [EVS console](#).

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

The **Snapshots** page is displayed.

**Step 3** In the snapshot list, locate the target snapshot and click **Delete** in the **Operation** column.

**Step 4** 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

## Related Links

- To check more snapshot FAQs, see [Snapshot](#).

## 4.3 Using EVS Snapshot Consistency Groups

### 4.3.1 Creating a Snapshot Consistency Group

#### Scenarios

If you create snapshots for multiple disks one by one, the snapshots are created at different times. This may result in inconsistent data. Snapshot consistency group allows you to create snapshots for disks on one or more servers at the same time and ensures crash consistency.

#### Impacts on Performance

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.

#### Billing

Creating a snapshot consistency group is free, but the snapshots created will be billed based on their storage used and usage duration. For details, see [Billing for EVS Snapshots](#).

#### Constraints

- You can use a snapshot consistency group to create snapshots for up to 64 disks at a time.
- Instant Snapshot Restore is not supported when you create snapshot consistency groups for Common I/O and High I/O disks.
- A snapshot consistency group can only contain disks that are in the same AZ.
- If a snapshot is being created for a disk, you can only create a snapshot consistency group after the snapshot is created.

#### Prerequisites

You can only create snapshot consistency groups for disks in the **Available** or **In-use** state. Check the disk status in the status column on the disk list page.

#### Procedure

- Step 1** Sign in to the [EVS console](#).
- Step 2** In the navigation pane on the left, choose **Elastic Volume Service > Snapshots**.  
The **Snapshots** page is displayed.
- Step 3** On the **Snapshot Consistency Groups** tab, click **Create Snapshot Consistency Group** in the upper right corner.
- Step 4** Configure the required parameters.

**Table 4-14** Parameter description

Parameter	Description
Region	The region to which the snapshot consistency group belongs. The snapshot consistency group must be in the same region as its source disks.
AZ	The AZ to which the snapshot consistency group belongs.
Select Servers	You can select multiple servers. The servers must be in the same region and AZ.
Select Disks	You can select multiple disks. The disks must be in the same region and AZ.
Group Name	Mandatory It can contain a maximum of 64 characters.
Group Description	The description can contain up to 255 characters.
Tag	Optional You can add tags to snapshots when creating a snapshot consistency group. A tag consists of a tag key and a tag value. <ul style="list-style-type: none"><li>• A tag key cannot start or end with a space, or start with <code>_sys_</code>. It can contain a maximum of 128 characters and contain letters, digits, spaces, and the following special characters: <code>._:=-@</code></li><li>• A tag value can contain a maximum of 255 characters and contain letters, digits, spaces, and the following special characters: <code>._:/=-@</code></li></ul>
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 <a href="#">Enabling or Disabling Instant Snapshot Restore (for Standard Snapshots)</a> .

Parameter	Description
Enterprise Project	When creating a snapshot consistency group, you can add the group to an existing enterprise project or a new one. An enterprise project facilitates project-level management and grouping of cloud resources and users. The default project is <b>default</b> .

**Step 5** Click **Submit**.

After the snapshot consistency group is created, you can view it on the **Snapshot Consistency Groups** tab.

----End

## Related Links

- To check more snapshot FAQs, see [Snapshot](#).
- To create an individual snapshot, see [Creating an EVS Snapshot](#).
- To roll back disk data from a snapshot consistency group, see [Rolling Back Disk Data Using a Snapshot Consistency Group](#).
- You are advised to delete unnecessary snapshots on a regular basis to reduce snapshot costs. For details, see [Deleting a Snapshot Consistency Group](#).

## 4.3.2 Rolling Back Disk Data Using a Snapshot Consistency Group

### Scenarios

If data on EVS disks is incorrect or damaged, you can roll back data using a snapshot consistency group to one or multiple source disks.

### Constraints

- Snapshot data can only be rolled back to source EVS disks. Rollback to a different disk is not possible.
- If a snapshot is being created, it cannot be used to roll back disk data.

### Prerequisites

- At least a snapshot consistency group has been created for disks.
- You can only roll back disk data from a snapshot consistency group when the source disk statuses are **Available** (not attached to any server) or **Rollback failed**. If any source disk is attached, detach it first.
- If the snapshot of a disk in the snapshot consistency group is deleted, data on that disk cannot be rolled back. Snapshots that are not deleted can be used to roll back data.

## Procedure

- Step 1** Sign in to the [EVS console](#).
- Step 2** In the navigation pane on the left, choose **Elastic Volume Service > Snapshots**.  
The **Snapshots** page is displayed.
- Step 3** Click the **Snapshot Consistency Groups** tab.
- Step 4** Find the target snapshot consistency group and click **Roll Back Disk** in the **Operation** column.
- Step 5** In the displayed dialog box, select the disks whose data you want to roll back.
- Step 6** Click **OK**.

After the rollback is successful, the system will display a successful rollback message.

You can log in to the ECS and check whether the EVS disk data has been rolled back to the state when the snapshots were created.

----End

## Related Links

- To check more snapshot FAQs, see [Snapshot](#).
- To roll back data from a single snapshot, see [Rolling Back Disk Data from a Snapshot](#).

### 4.3.3 Deleting a Snapshot Consistency Group

#### Scenarios

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

---

**CAUTION**

All the snapshots in the snapshot consistency group will be deleted together with the group.

---

## Procedure

- Step 1** Sign in to the [EVS console](#).
- Step 2** In the navigation pane on the left, choose **Elastic Volume Service > Snapshots**.  
The **Snapshots** page is displayed.
- Step 3** Click the **Snapshot Consistency Groups** tab.
- Step 4** Find the target snapshot consistency group and click **Delete** in the **Operation** column.

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

If the snapshot consistency group disappears from the list, it is deleted successfully.

----**End**

## Related Links

- To check more snapshot FAQs, see [Snapshot](#).
- To delete an individual snapshot, see [Deleting an EVS Snapshot](#).

# 5 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.

## Constraints

**Table 5-1** Constraints on the disk type change

Phase	Description
Before the change	<ul style="list-style-type: none"><li>You can only change the disk type when the disk status is <b>Available</b> or <b>In-use</b>.</li><li>The disk type cannot be changed when any snapshot of the disk is being deleted.</li><li>A disk having more than 128 snapshots cannot have its disk type changed. You can delete some snapshots and then perform the change.</li><li>In rare cases, the disk type may fail to be changed due to a background resource issue. If this happens, submit a service ticket.</li><li>A disk protected by Business Recovery Service (BRS) cannot have its disk type changed.</li></ul>

Phase	Description
During the change	<ul style="list-style-type: none"> <li>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.</li> <li>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.</li> <li>In rare cases, the change may fail due to resource problems. In this case, you are advised to perform the change again.</li> <li>You can have a maximum of 10 disks with their types being changed at the same time.</li> <li>The OS cannot be changed if you are changing the disk type of a system disk.</li> </ul>
After the change	<ul style="list-style-type: none"> <li>In rare cases, the disk type may fail to be changed due to a background resource issue. If this happens, try again later.</li> <li>In rare cases, the disk type may fail to be changed after <a href="#">a data rollback from a snapshot</a>. If this happens, <a href="#">submit a service ticket</a>.</li> </ul>

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 5-2** Supported changes between disk types

Source Disk Type	New Disk Type
Extreme SSD V2	Extreme SSD V2 (IOPS changed)
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

Source Disk Type	New Disk Type
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
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, but the disk performance may be affected. Perform the change during off-peak hours.

## Billing

After a disk specifications change, the billing of the disk will also be changed:

- For a pay-per-use disk, the disk will be billed based on the billing standards of the new disk type.
- For a yearly/monthly disk, the disk expiration time remains unchanged, and the system will calculate the price difference based on the remaining days and the price difference between the original and new disk specifications.

## Prerequisites

You are advised to create a snapshot for the disk to back up the disk data before changing the disk specifications. For details, see [Creating an EVS Snapshot](#).

## Procedure

**Step 1** Sign in to the [EVS console](#).

**Step 2** 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 3** Select a disk type from the drop-down list. The system shows you the new price based on your selection.

To change to the General Purpose SSD V2 type, you also need to specify the disk IOPS and throughput.

To change to the Extreme SSD V2 type, you also need to specify the disk IOPS.

**Step 4** 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.

 **NOTE**

If the disk is a yearly/monthly disk, the costs may change. Handle resource change orders in a timely manner.

----End

## Related Links

- To learn more about the disk performance, see [Disk Types and Performance](#).
- After a disk type change, the performance of the disk is not only restricted by the new disk type. The instance that the disk is attached to also restricts the disk performance. For details, see [Instance QoS](#).

# 6 Viewing EVS Disk Details

---

## Scenarios

When using EVS disks, you may need to check the disk information, such as the disk status, type, and capacity. This section describes how to view disk details. Methods are provided as follows:

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

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

### NOTE

Disks in the [grace period or retention period](#) can still be viewed in the disk list even if your account is in arrears.

## Viewing Disk Details from the EVS Console

**Step 1** Sign in to the [EVS console](#).

**Step 2** In the disk list, view disk information including the disk status, type, size, function, and device type.

In the search box above the list, you can search for disks by project, status, disk name, tag, or other attributes.

**Step 3** 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 4** (Optional) Export disk information.

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

----End

## Viewing Disk Details from the ECS Console

**Step 1** Sign in to the [ECS console](#).

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

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

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

The ECS 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

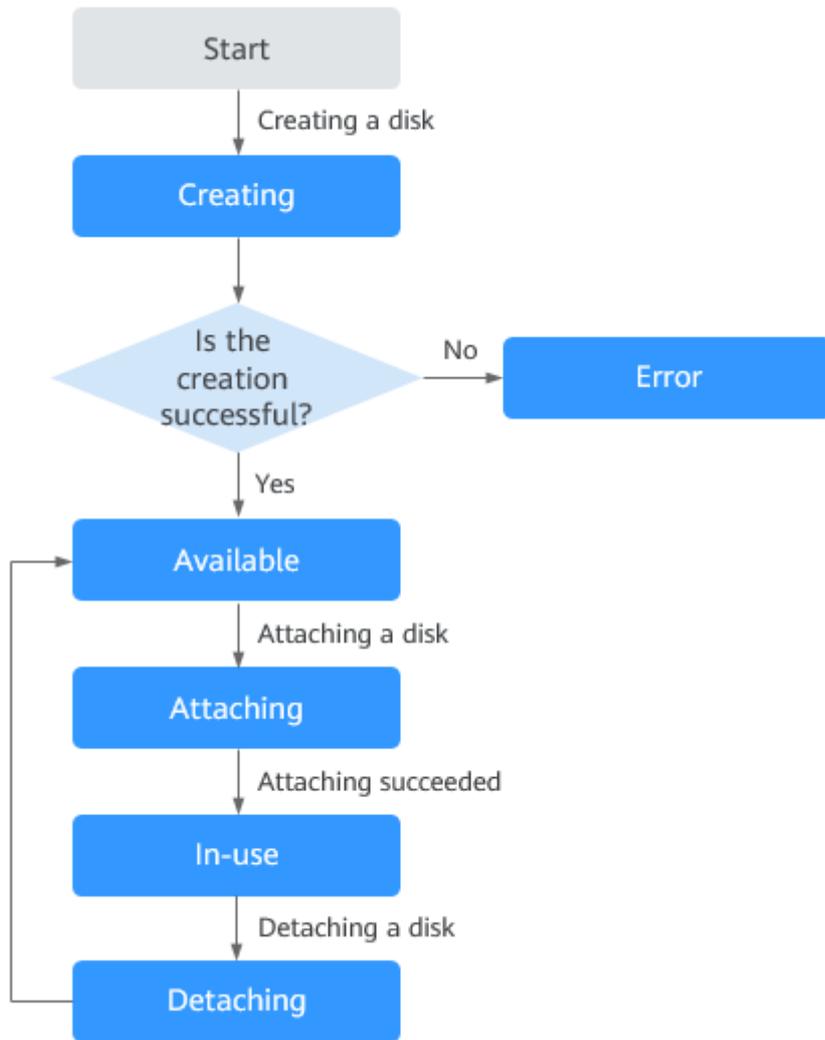
## EVS Disk Status

**Table 6-1** EVS disk status details

Status	Description	Allowed Operation
In-use	The EVS disk has been attached to a server and is in use.	<ul style="list-style-type: none"> <li>• Detaching</li> <li>• Creating backups</li> <li>• Expanding capacity</li> <li>• Creating snapshots</li> <li>• Modifying disk specifications</li> </ul>
Available	The EVS disk has not been attached to any server, so you can attach it.	<ul style="list-style-type: none"> <li>• Attaching</li> <li>• Expanding capacity</li> <li>• Deleting pay-per-use disks or unsubscribing from yearly/monthly disks</li> <li>• Creating backups</li> <li>• Rolling back data to EVS disks using snapshots</li> <li>• Creating snapshots</li> <li>• Modifying disk specifications</li> </ul>
Creating	The EVS disk is being created.	-
Attaching	The EVS disk is being attached to a server.	-
Detaching	The EVS disk is being detached from a server.	-
Deleting	The EVS disk is being deleted.	-

Status	Description	Allowed Operation
Restoring	A backup is being used to restore the EVS disk.	-
Expanding	The capacity of the EVS disk is being expanded.	-
Uploading	Data on the EVS disk is being uploaded to an image. This status occurs when you create an image from a server.	-
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. <b>NOTE</b> <ul style="list-style-type: none"><li>When you roll back data from a snapshot, you can only roll back data to the source EVS disk. Rollback to a specific disk is not supported.</li><li>A snapshot can only be used for rollback when its source disk is in the <b>Available</b> or <b>Rollback failed</b> state.</li></ul>	-
Rollback failed	An error occurs when the EVS disk data is rolled back from a snapshot.	<ul style="list-style-type: none"><li>Deleting</li><li>Rolling back data to EVS disks using snapshots</li></ul>
Awaiting transfer	The EVS disk is waiting for transfer.	-

**Figure 6-1** 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.

### Related Links

To learn more about the disk capacity usage, check on the ECS or install the agent on the ECS to obtain the disk usage from the monitoring data. For details, see [How Can I View My Disk Usage?](#)

# 7 Detaching and Deleting an EVS Disk

## 7.1 Detaching an EVS Disk

### Scenarios

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

 NOTE

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

## Billing

A detached EVS disk will not be automatically deleted, so it will still be billed, as described in [Billing for EVS Disks](#). If you no longer need a disk, [delete or unsubscribe from](#) it in a timely manner.

## Constraints

- After a system disk is detached, some operations cannot be performed on the original server and the system disk. The unsupported operations are as follows:
  - Server: starting the server, remote login, resetting the password, changing the 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
- A shared data disk can be detached from ECSs in a batch.

## Prerequisites

- Before detaching an EVS disk from a running Windows server, ensure that no programs are reading data from or writing data to the disk. Otherwise, data will be lost.
- Before detaching an EVS disk from a running Linux server, you must log in to the server and run the **umount** command to cancel the association between the disk and the file system, and ensure that no programs are reading data from or writing data to the disk. Otherwise, you will not be able to detach the disk.
- You are advised to back up data. You can [create snapshots](#) or [use CBR to create disk backups](#).

## Detaching a System Disk

**Step 1** Sign in to the [EVS console](#).

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

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

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

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

**Step 4** Click the name of this ECS.

The ECS details page is displayed.

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

**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 succeeds, the detached system disk is no longer displayed under the **Disks** tab.

**Step 8** (Optional) **Re-attach** the bootable disk to another 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** Sign in to the [EVS console](#).

**Step 2** Choose a way to detach the disk by determining whether you want to check the server information first.

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

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

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

----End

## Detaching a Shared Data Disk

**Step 1** Sign in to the [EVS console](#).

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

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

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

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

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

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

----End

## Related Links

- After disks are detached, you can [attach](#) them to other servers in the same AZ.
- To detach disks using the API, see [Detaching an EVS Disk from an ECS](#).
- To check out more detachment FAQs, see [Detachment](#).

## 7.2 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. After you permanently delete an EVS disk, EVS immediately destroys

the metadata to ensure that data can no longer be accessed. In addition, the physical storage space of the EVS 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.

If you have enabled the EVS recycle bin, when you delete a disk, whether the disk will be deleted permanently or moved to the recycle bin depends on your recycle bin policy. The recycle bin is disabled by default. If you need to use it, 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 7-1](#).

**Table 7-1** 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	<a href="#">How Do I Unsubscribe from ECSs?</a>
	Unsubscribing from data disks separately	<a href="#">Unsubscribing from a Yearly/Monthly Disk on the EVS Console</a> <a href="#">Unsubscribing from a Yearly/Monthly Disk on the ECS Console</a>
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	<a href="#">Unsubscribing from a Yearly/Monthly Disk on the EVS Console</a>
Unsubscribing from yearly/monthly data disks that were purchased on the EVS console	Unsubscribing from data disks separately	<a href="#">Unsubscribing from a Yearly/Monthly Disk on the EVS Console</a>

## 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 Business Recovery Service (BRS). 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.

---

## Billing

- For a yearly/monthly disk, the billing stops after the disk is successfully unsubscribed from, and the refund is calculated as follows: Refund = Your actual payment - Amount due - Handling fees. For more information, see [How Do I View the Refund for My Resource Unsubscription?](#)
- For a pay-per-use disk, the billing stops after the disk is successfully deleted.

## Prerequisites

- Ensure that separately purchased disks are detached. For details, see [Detaching an EVS Disk](#).
- You are advised to back up data. You can [create standard snapshots](#) or [use CBR to create disk backups](#).

## Deleting Pay-per-Use EVS Disks

**Step 1** Sign in to the [EVS console](#).

**Step 2** In the disk list, locate the target disk and choose **More > Delete** in the **Operation** column.

**Step 3** (Optional) If multiple disks are to be deleted, select  in front of each target disk and click **Delete** in the upper left area of the list.

**Step 4** On the displayed page, confirm the information and click **OK**.

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

Supported verification methods include SMS, email, and 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 5** Click **OK**.

- If the EVS recycle bin is disabled, the disk will be deleted immediately.  
When the disk disappears from the disk list, the disk has been deleted.
- If the EVS recycle bin is enabled, the disk will be moved to the recycle bin if the days passed since the disk creation is greater than that configured in the recycle bin policy. To permanently delete the disk, go to the recycle bin, find the disk, and [delete it from the bin](#).

When the disk disappears from the recycle bin, the disk has been deleted.

----End

## Unsubscribing from a Yearly/Monthly Disk on the EVS Console

**Step 1** Sign in to the [EVS console](#).

**Step 2** 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 3** On the resource unsubscription page, confirm the information, select required check boxes, and click **Confirm**.

- If the EVS recycle bin is disabled, the disk will be unsubscribed from immediately.  
When the disk disappears from the disk list, the disk has been unsubscribed from.
- If the EVS recycle bin is enabled, the disk will be moved to the recycle bin if the days passed since the disk creation is greater than that configured in the recycle bin policy. To permanently delete the disk, go to the recycle bin, find the disk, and [delete it from the bin](#).

When the disk disappears from the recycle bin, the disk has been deleted.

----End

## Unsubscribing from a Yearly/Monthly Disk on the ECS Console

### NOTE

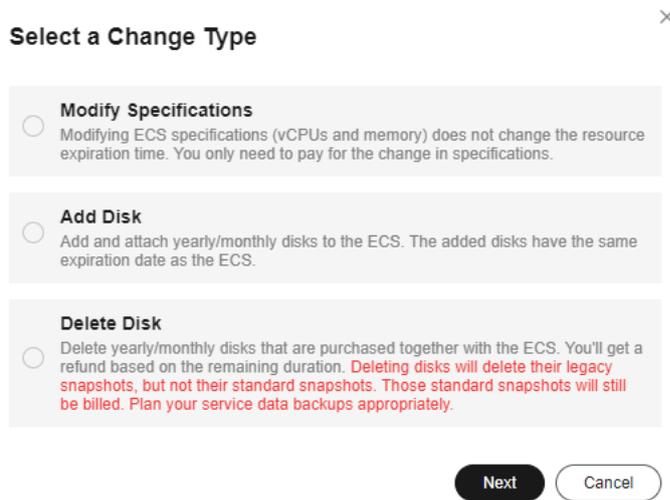
- When you unsubscribe from a server, the yearly/monthly data disks purchased together with and later added to the server will also be unsubscribed from. For details, see [How Do I Unsubscribe from ECSs?](#)
- To unsubscribe from non-shared data disks separately, perform the following steps:

**Step 1** Sign in to the [ECS console](#).

**Step 2** In the ECS list, locate the target ECS, choose **More > Change** in the **Operation** column.

**Step 3** In the displayed dialog box, select **Delete Disk**.

**Figure 7-1** Selecting a change type



**Step 4** Select the disks you want to delete and click **Next**.

**Step 5** On the deletion page, confirm the information, select "I understand a handling fee will be charged for this unsubscription", and click **Submit**.

**Step 6** After the unsubscription is submitted:

- If the EVS recycle bin is disabled, the disk will be unsubscribed from immediately.  
When the disk disappears from the disk list, the disk has been unsubscribed from.
- If the EVS recycle bin is enabled, the disk will be moved to the recycle bin if the days passed since the disk creation is greater than that configured in the recycle bin policy. To permanently delete the disk, go to the recycle bin, find the disk, and **delete it from the bin**.  
When the disk disappears from the recycle bin, the disk has been deleted.

----End

## Related Links

To delete disks using the API, see [Deleting an EVS Disk](#).

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

# 8 Managing EVS Recycle Bin

---

## 8.1 Recycle Bin Overview

EVS recycle bin is disabled by default. You need to manually [enable EVS recycle bin](#) before you can use it.

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

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

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

 **NOTE**

EVS recycle bin is only available in some regions. You can check the supported regions on the console.

### 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.
- You can recover or permanently delete the disks in the recycle bin. After the disks expire, they are permanently deleted and cannot be recovered.
- You can permanently delete up to 50 disks at a time.
- If the ECS recycle bin is enabled, EVS disks deleted or unsubscribed from together with ECSs will be retained in the recycle bin for the same period of time as the ECSs and can only be recovered or permanently deleted on the ECS side. After the disks expire, they are permanently deleted and cannot be recovered.

## Billing

- 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 [Will I Be Billed for the Disks in the Recycle Bin?](#)
- If you already have disks 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 in the bin before the system permanently deletes them. For details, see the billing example provided in [Billing for EVS Recycle Bin](#).

## 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 expire.
- 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 passes 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](#).

If ECS recycle bin is also supported, and you configure both the ECS and EVS 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 8-2](#). Assume that an ECS (including system and data disks) was created 8 days ago. If you delete or unsubscribe from the ECS, the resources would be processed as follows.

**Table 8-1** Scenarios and rules for processing resources

Scenario	Processing Rule	Processing Results for ECSs	Processing Results for EVS Disks
<p>If the minimum number of days set in the EVS recycle bin policy is greater than that set in the ECS recycle bin policy, for example:</p> <ul style="list-style-type: none"> <li>ECS recycle bin policy: 7 days</li> <li>EVS recycle bin policy: 15 days</li> </ul>	Both ECSs and EVS disks comply with the ECS recycle bin policy.	Moved to the recycle bin	Moved to the recycle bin
<p>If the minimum number of days set in the EVS recycle bin policy is smaller than that set in the ECS recycle bin policy, for example:</p> <ul style="list-style-type: none"> <li>ECS recycle bin policy: 15 days</li> <li>EVS recycle bin policy: 7 days</li> </ul>	<p>ECSs comply with the ECS recycle bin policy.</p> <p>EVS disks comply with the EVS recycle bin policy.</p>	Deleted or unsubscribed from	Moved to the recycle bin

## Recycle Bin Operations

**Table 8-2** 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.	<a href="#">Enabling the Recycle Bin</a>
Disable the recycle bin.	<p>You can disable the recycle bin if you no longer need it.</p> <p>You must empty the recycle bin before disabling it.</p>	<a href="#">Disabling the Recycle Bin</a>
Configure a recycle bin policy.	You can configure a recycle bin policy to define when to move deleted disks to the recycle bin.	<a href="#">Configuring a Recycle Bin Policy</a>

Operation	Description	Reference
Recover disks from the recycle bin.	You can recover disks from the recycle bin.	<a href="#">Recovering Disks from the Recycle Bin</a>
Permanently delete disks from the recycle bin.	You can permanently delete the EVS disks from the recycle bin at any time.	<a href="#">Permanently Deleting Disks from the Recycle Bin</a>

## 8.2 Enabling the Recycle Bin

### Scenarios

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

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

### Constraints

- To learn when EVS disks will 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, 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.
- By default, disks created at least 7 days ago will be moved to the recycle bin upon deletion or unsubscription. You can customize the value, up to a maximum of 1,000 days.

### Procedure

- Step 1** Sign in to the [EVS console](#).
- Step 2** Click the **Recycle Bin** tab.
- Step 3** On the **Recycle Bin** tab, click **Enable Recycle Bin** to open the **Configure Recycle Bin Policy** page.
- Step 4** Configure the recycle bin policy.

### Configure Recycle Bin Policy ×

- 1 • Disks in the EVS recycle bin are billed based on a pay-per-use basis. [Billing details](#)
- Disks can be moved to the recycle bin if they are not in the retention period and their statuses are normal and their accounts are not in arrears, restricted, or frozen. [Learn more](#)
- If your account is in arrears, disks in the recycle bin will be kept shorter than the duration you configured.
- There are no limits on the capacity and quantity of disks in the recycle bin.
- If the ECS recycle bin is enabled, the ECS recycle bin policy will apply to any EVS disks deleted or unsubscribed from together with the ECSs.

Move disks to the recycle bin upon deletion or unsubscription if they were created at least this number of days ago:

If disks were created more than this number of days ago, they will be moved to the recycle bin upon deletion or unsubscription.

Permanently delete disks after they are kept in the recycle bin for more than this number of days:

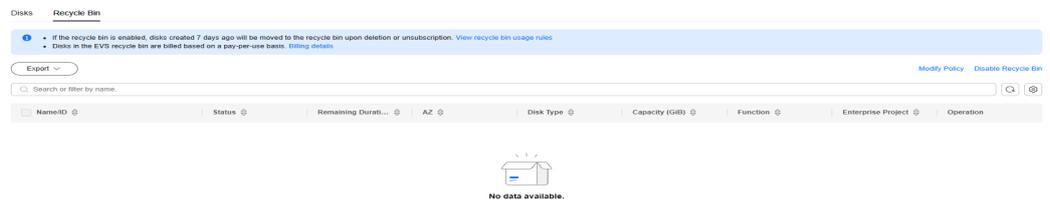
If disks are kept in the recycle bin for more than the duration you specify, they will be deleted permanently. This rule only applies to the disks that are moved to the recycle bin after the rule is configured.

**Table 8-3** Recycle bin policy parameters

Parameter	Example Value	Description
Number of days since disk creation	1	<p>If disks were created more than this number of days ago, they will be moved to the recycle bin upon deletion or unsubscription.</p> <p>For example, if you set this parameter to <b>1</b>, disks created within 1 day will not be moved to the recycle bin upon deletion, but deleted permanently. Disks created more than 1 day ago will be moved to the recycle bin upon deletion and billed on a pay-per-use basis.</p> <p><b>NOTE</b> If you use both ECS and EVS 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. For details, see <a href="#">Recycle Bin Policy Configuration Suggestions</a>.</p>
Number of days disks can be retained	10	<p>The maximum number of days that disks can be retained in the recycle bin. When a disk stays longer than the days you specify, it will be automatically deleted.</p> <p>For example, if you set this parameter to <b>10</b>, disks can stay in the recycle bin for up to 10 days. After that, the disks will be automatically deleted and cannot be recovered.</p>

- Step 5** Click **OK**. Deleted or unsubscribed disks that meet the recycle bin policy will be displayed on the **Recycle Bin** tab.

**Figure 8-1** Enable Recycle Bin



----End

## Related Links

- You can **disable** the recycle bin if you no longer need it.
- You can customize how many days after disk creation the disks will be moved to the recycle bin when unsubscribed from or deleted. For details, see [Configuring a Recycle Bin Policy](#).

## 8.3 Configuring a Recycle Bin Policy

### Scenarios

After enabling the recycle bin, you can customize the recycle bin policy to determine when to move unsubscribed or deleted disks to the bin.

#### NOTE

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 bin, you can configure an appropriate 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.
- 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.

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

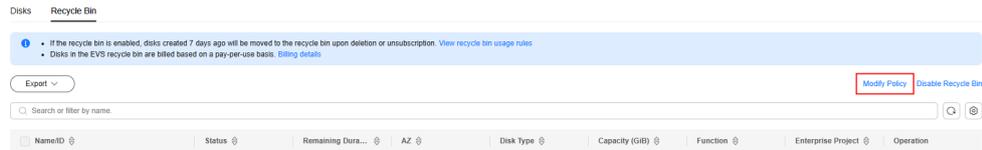
### Constraints

- By default, disks created at least 7 days ago will be moved to the recycle bin upon deletion or unsubscription. You can customize the value, up to a maximum of 1,000 days.
- You can configure a retention duration up to 365 days for the disks in the recycle bin. The default retention duration is 7 days.
- 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 passed since disk creation with the days configured in the recycle bin policy.

## Procedure

- Step 1** Sign in to the [EVS console](#).
- Step 2** Click the **Recycle Bin** tab.
- Step 3** In the upper right corner of the **Recycle Bin** tab page, click **Modify Policy**.  
The **Configure Recycle Bin Policy** page is displayed.

Figure 8-2 Modify Policy



- Step 4** Configure the recycle bin policy.

### Configure Recycle Bin Policy

- Disks in the EVS recycle bin are billed based on a pay-per-use basis. [Billing details](#)
- Disks can be moved to the recycle bin if they are not in the retention period and their statuses are normal and their accounts are not in arrears, restricted, or frozen. [Learn more](#)
- If your account is in arrears, disks in the recycle bin will be kept shorter than the duration you configured.
- There are no limits on the capacity and quantity of disks in the recycle bin.
- If the ECS recycle bin is enabled, the ECS recycle bin policy will apply to any EVS disks deleted or unsubscribed from together with the ECSs.

Move disks to the recycle bin upon deletion or unsubscription if they were created at least this number of days ago:

If disks were created more than this number of days ago, they will be moved to the recycle bin upon deletion or unsubscription.

Permanently delete disks after they are kept in the recycle bin for more than this number of days:

If disks are kept in the recycle bin for more than the duration you specify, they will be deleted permanently. This rule only applies to the disks that are moved to the recycle bin after the rule is configured.

**Table 8-4** Recycle bin policy parameters

Parameter	Example Value	Description
Number of days since disk creation	1	<p>If disks were created more than this number of days ago, they will be moved to the recycle bin upon deletion or unsubscription.</p> <p>For example, if you set this parameter to <b>1</b>, disks created within 1 day will not be moved to the recycle bin upon deletion, but deleted permanently. Disks created more than 1 day ago will be moved to the recycle bin upon deletion and billed on a pay-per-use basis.</p> <p><b>NOTE</b> If you use both ECS and EVS 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. For details, see <a href="#">Recycle Bin Policy Configuration Suggestions</a>.</p>
Number of days disks can be retained	10	<p>The maximum number of days that disks can be retained in the recycle bin. When a disk stays longer than the days you specify, it will be automatically deleted.</p> <p>For example, if you set this parameter to <b>10</b>, disks can stay in the recycle bin for up to 10 days. After that, the disks will be automatically deleted and cannot be recovered.</p>

**Step 5** 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 Links

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

## 8.4 Recovering Disks from the Recycle Bin

### Scenarios

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

### Constraints

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

- If the ECS recycle bin is enabled, EVS disks deleted together with ECSs will be retained in the recycle bin for the same period of time as the ECSs and can only be recovered on the ECS side.

## Billing

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](#).

## Procedure

**Step 1** Sign in to the [EVS console](#).

**Step 2** Click the **Recycle Bin** tab.

**Step 3** 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 4** 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

## 8.5 Permanently Deleting Disks from the Recycle Bin

### Scenarios

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

### Constraints

- You can recover or permanently delete the disks in the recycle bin. After the disks expire, they are permanently deleted and cannot be recovered.



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

---

## Permanently Deleting an EVS Disk

**Step 1** Sign in to the [EVS console](#).

**Step 2** Click the **Recycle Bin** tab.

**Step 3** Locate the disk you want to permanently delete and click **Delete** in the **Operation** column.

The confirmation dialog box is displayed.

**Step 4** Click **Yes**.

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

----End

## Permanently Deleting Multiple Disks in a Batch

**Step 1** Sign in to the [EVS console](#).

**Step 2** Click the **Recycle Bin** tab.

**Step 3** In the disk list, click  in front of the desired disks and click **Delete** in the upper left corner of the list.

The confirmation dialog box is displayed.

**Step 4** Click **Yes**.

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

----End

# 8.6 Disabling the Recycle Bin

## Scenarios

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

## 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** Sign in to the [EVS console](#).

**Step 2** Click the **Recycle Bin** tab.

**Step 3** In the upper right corner of the **Recycle Bin** tab page, click **Disable Recycle Bin**.

A dialog box is displayed.

**Step 4** Click **OK**.

When message "Recycle bin is disabled" is displayed, the recycle bin is disabled successfully.

----**End**

# 9 Managing Encrypted EVS Disks

---

## 9.1 EVS Encryption Overview

### What Is EVS Encryption?

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

EVS encryption uses the industry-standard XTS-AES-256 algorithm and Key Management Service (KMS) keys provided by Data Encryption Workshop (DEW) for encryption. With EVS encryption, you do not need to establish and maintain your own 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.

### How EVS Encryption Works

The encryption system uses a two-layer key structure. The first-layer key is the customer master key (CMK), and the second-layer key is the data key (DK). The CMK encrypts and decrypts the DK to ensure their security in transit and at rest. The DK encrypts and decrypts service data. The details are as follows:

#### 1. Encrypt the DK

Before being used to encrypt service data, a DK is first encrypted by a CMK. Only encrypted DKs can be stored or transferred. If an attacker gains access to an encrypted DK and service data, it cannot decrypt data due to the lack of the CMK.

#### 2. Encrypt data in transit and at rest

To read encrypted data, a decryption request is first sent to KMS to obtain the plaintext DK. KMS verifies the request validity and then uses the CMK to decrypt the DK and returns the plaintext DK. The decryption is done in the memory, so the plaintext DK will not be persistently stored on any storage medium. The system then uses the plaintext DK in the memory to decrypt disk I/O data to ensure the security of data in transit and at rest.

## 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 keys. 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 DK to the host memory for use. EVS uses the plaintext DK to encrypt and decrypt disk I/Os. The plaintext DK 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 scheduled for deletion in KMS, the disk encrypted using this custom key can still use the plaintext DK stored in the host memory. If this disk is later detached, the plaintext DK 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 available.

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"> <li>• For an encrypted disk already attached: Reads and writes to the disk are normal. If the disk is detached, it cannot be attached again.</li> <li>• For an unattached encrypted disk: The disk cannot be attached anymore.</li> </ul>	Enable the custom key. For details, see <a href="#">Creating a Custom Key</a> .
Scheduled deletion		Cancel the scheduled deletion for the custom key. For details, see <a href="#">Creating a Custom Key</a> .
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.

## Relationships Between EVS Encryption, Snapshots, Backups, and Images

You can use EVS encryption to encrypt system disks, data disks, snapshots, backups, and images. Find the details below:

- System disk encryption relies on the image that is used to create the server.
  - If an encrypted image is used to create the server, the system disk will be encrypted by default, and the system disk and image share the same encryption method. For details, see [Encrypting Images](#).
  - If a non-encrypted image is used to create the server, you can determine whether to encrypt the system disk or not during the server creation. For details, see "Getting Started" > "Creating an ECS" > "Step 1: Configure Basic Settings" in the *Elastic Cloud Server User Guide*.
  - If a non-encrypted image is used and you want an encrypted system disk, first replicate the non-encrypted image to be an encrypted one, create the server, and then create the encrypted system disk. For details, see [Replicating Images Within a Region](#).
- If an empty disk is created, you can determine whether to encrypt the disk or not. The encryption attribute of the disk cannot be changed after the disk has been created.
- If a disk is created from a backup, the encryption attribute of the disk does not need to be the same as that of the backup.
- If a disk is created from an image, the encryption attribute of the disk will be the same as that of the image's source disk.
- If a backup is created for a disk, the encryption attribute of the backup will be the same as that of the disk.
- If a snapshot is created for a disk, the encryption attribute of the snapshot is the same as that of the disk.

## Relationships Between EVS Encryption and Backups

You can use EVS encryption to encrypt system disks, data disks, and backups. Find the details below:

- System disk encryption relies on the image that is used to create the server.
  - If an encrypted image is used to create the server, the system disk will be encrypted by default, and the system disk and image share the same encryption method. For details, see "Managing Private Images" > "Encrypting Images" in the *Image Management Service User Guide*.
  - If a non-encrypted image is used to create the server, you can determine whether to encrypt the system disk or not during the server creation. For details, see "Getting Started" > "Creating an ECS" > "Step 1: Configure Basic Settings" in the *Elastic Cloud Server User Guide*.
- If an empty disk is created, you can determine whether to encrypt the disk or not. The encryption attribute of the disk cannot be changed after the disk has been created.
- If a disk is created from a backup, the encryption attribute of the disk does not need to be the same as that of the backup.
- If a backup is created for a disk, the encryption attribute of the backup will be the same as that of the disk.

## 9.2 Creating an Encrypted EVS Disk

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

### Disk 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** Relationship between images and system disk encryption

Whether to Encrypt System Disk When Purchasing Server	Whether to Create Server from an Encrypted Image	Whether System Disk Will Be Encrypted	Description
Yes (key A)	Yes (key B)	Yes (key A)	<ul style="list-style-type: none"><li>• To encrypt system disks during the server purchase, see <a href="#">Purchasing an ECS in Custom Config Mode</a>.</li><li>• For details about how to encrypt images, see <a href="#">Encrypting Images</a>.</li></ul>
Yes (key A)	No	Yes (key A)	To encrypt system disks during the server purchase, see <a href="#">Purchasing an ECS in Custom Config Mode</a> .
No	Yes (key B)	Yes (key B)	For details about how to encrypt images, see <a href="#">Encrypting Images</a> .
No	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 <a href="#">Replicating Images Within a Region</a> .

- **Data disk encryption** (default encryption disabled)

Data disks can be purchased along with servers or separately. Whether data disks are encrypted or not depends on their data sources. See the following table for details.

**Table 9-3** Relationship between backups, snapshots, images, and data disk encryption

Buy Disk On	Method of Purchase	Whether Data Disk Will Be Encrypted	Description
ECS console	Buying together with a 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> .
EVS console	No data source selected	Yes/No	When an empty disk is created, you can choose whether to encrypt the disk or not. The encryption attribute of the disk cannot be changed after the disk has been created.
	Creating from a backup	Yes/No	<ul style="list-style-type: none"> <li>When a disk is created from a backup, you can choose whether to encrypt the disk or not. The encryption attributes of the disk and backup do not need to be the same.</li> <li>When you create a backup for a system or data disk, the encryption attribute of the backup will be the same as that of the disk.</li> </ul>
	Creating from a snapshot (The snapshot's source disk is encrypted.)	Yes	A snapshot created from an encrypted disk is also encrypted.

Buy Disk On	Method of Purchase	Whether Data Disk Will Be Encrypted	Description
	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	-

- **Data disk encryption** (default encryption enabled)

Data disks can be purchased with servers or separately. Whether data disks are encrypted or not depends on their creation scenarios and data sources. For how to enable default encryption, see [Configuring Default Encryption](#).

**Table 9-4** Relationship between backups, snapshots, images, and data disk encryption

Creation Scenario	Whether Data Source Is Encrypted	Whether Data Disk Will Be Encrypted	Description
Creating together with a server	Empty disk	Yes	You can use the key preset for default encryption or change the key as required.
No data source selected	Empty disk	Yes	You can use the key preset for default encryption or change the key as required.
Creating from a legacy snapshot	Non-encrypted	No	Encryption is not supported.
	Encrypted	Yes	The key of the snapshot is inherited.

Creation Scenario	Whether Data Source Is Encrypted	Whether Data Disk Will Be Encrypted	Description
Creating from a standard snapshot (Instant Snapshot Restore is enabled, but data upload is not complete)	Non-encrypted	No	Encryption is not supported.
	Encrypted	Yes	The key of the snapshot is inherited.
Creating from a standard snapshot (data upload is complete)	Non-encrypted	Yes	You can use the key preset for default encryption or change the key as required.
	Encrypted	Yes	You can use the key preset for default encryption or change the key as required.
Creating from a private image	Non-encrypted	Yes	You can use the key preset for default encryption or change the key as required.
	Encrypted	Yes	You can use the key preset for default encryption or change the key as required.
Creating from a public image	Non-encrypted	Yes	You can use the key preset for default encryption or change the key as required.
	Encrypted	Yes	You can use the key preset for default encryption or change the key as required.
Creating from a disk backup (shared or non-shared)	Non-encrypted	Yes	You can use the key preset for default encryption or change the key as required.
	Encrypted	Yes	You can use the key preset for default encryption or change the key as required.

## Constraints

**Table 9-5** Constraints on disk encryption

Item	Description
Disk types supporting encryption	All disk types support encryption, but the encryption attribute of an existing disk cannot be changed.
Disk encryption	<ul style="list-style-type: none"><li>• The encryption attribute of a disk cannot be changed after the disk is created, meaning that:<ul style="list-style-type: none"><li>- An encrypted disk cannot be changed to a non-encrypted disk.</li><li>- A non-encrypted disk cannot be changed to an encrypted disk.</li></ul></li></ul>
User permissions	<p>When a user uses encryption, the condition varies depending on whether the user is the first one ever in the current region or project to use this function.</p> <ul style="list-style-type: none"><li>• If the user is the first user, the user needs to follow the prompt to create an agency, which grants <b>EVS KMSAccess</b> permissions to EVS. Then, the user can create and obtain keys to encrypt and decrypt disks.</li><li>• If the user is not the first user, the user can use encryption directly.</li></ul>
Image encryption	<ul style="list-style-type: none"><li>• Encrypted images cannot be replicated across regions.</li><li>• Encrypted images cannot be changed to non-encrypted images.</li><li>• Encrypted images cannot be exported.</li></ul>

## Billing

If KMS encryption is used, what you use beyond the free quota given by KMS will be billed. For details, see [DEW Billing](#).

## Creating an Encrypted EVS Disk

Before you use the encryption function, KMS access permissions 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](#).

You can use encrypted system disks immediately after they are created. You need to attach and initialize encrypted data disks after they are created.

**Table 9-6** Follow-up operations

Step	Description
Step 1: Attach the disk.	If you choose not to attach the disk when purchasing the disk, you need to manually <b>attach</b> it later.
Step 2: Initialize the disk.	The procedure for initializing a newly created empty data disk differs from that for a data disk with data on it. For details, see <b>Initialization Overview</b> .

## 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**.

## Related Links

- To learn more about KMS keys, see **KMS Overview**.
- To learn more about encryption principles, see **EVS Encryption Overview**.

# 9.3 Configuring Default Encryption for EVS Disks (OBT)

## Scenarios

Default encryption is a region-specific setting. You can configure your account to enforce the encryption of the newly created EVS disks and replicated images that you create. Properly configuring default encryption helps you ensure consistent data encryption across regions, simplify data security management, and improve data security.

## Precautions

- Default encryption is in OBT. **Submit a service ticket** to apply for using it.
- After default encryption is enabled, whether newly created disks are encrypted or not depends on their data sources. For details, see **Disk Encryption Scenarios**.

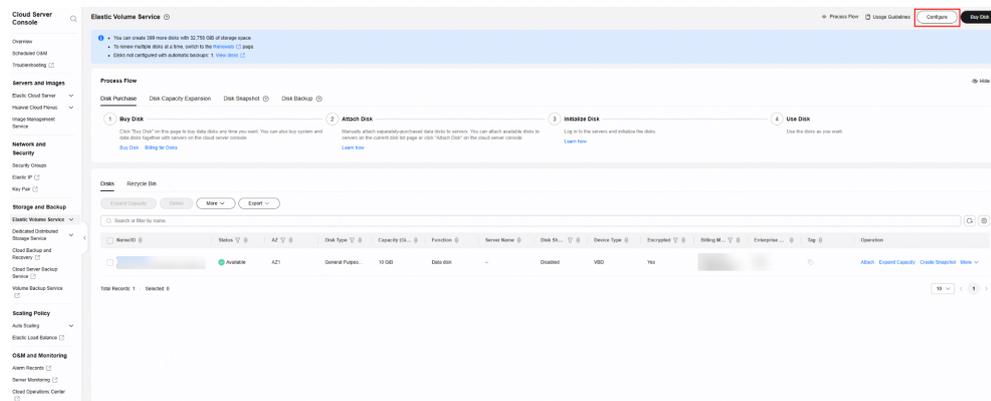
- Before using encryption, you need to create an agency to grant KMS access permissions to EVS. If you have the right to grant the permission, grant KMS access permissions to EVS directly. After KMS access permissions are granted, follow-up operations do not require the permissions to be granted again. If you do not have this permission, contact a user with the security administrator permissions to grant KMS access permissions to EVS, then repeat the preceding operations.

## Configuring Default Encryption

You can enable or disable default encryption or change the key on the **Data Encryption Protection** tab. The following steps describe how to enable default encryption:

- Step 1** On the disk list page, click **Configure** in the upper right corner.

**Figure 9-1** Configuring default encryption



- Step 2** On the **Data Encryption Protection** tab, toggle on the switch.

**Figure 9-2** Enabling default encryption

### Configure

Data Encryption Protection 

#### Default Encryption

Enabled

Once enabled, encryption will be enabled for all new disks. You will be unable to disable encryption.

 If KMS encryption is used, what you use beyond the free quota given by KMS will be billed.

KMS Key Name

evs/default   [View KMS Key](#)

KMS Key ID

57b76c10-b29c-4f15-b867-fd8f371edf39 

#### Default Encryption Enabled

KMS Key Name

--

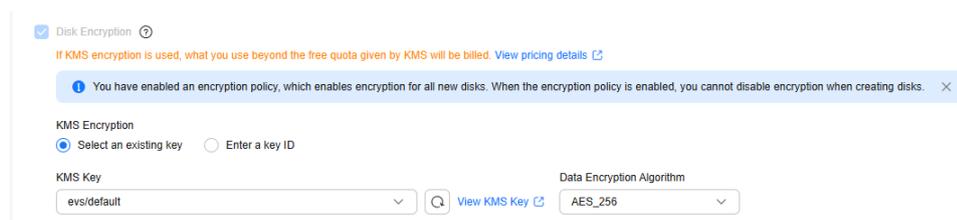
KMS Key ID

-- 

**Step 3** Select a key and click **OK**.

**Step 4** Go back to the disk list page and click **Buy Disk** in the upper right corner.

If you buy a new empty disk, the encryption option is preselected by default and you cannot deselect it. For other scenarios with default encryption enabled, see [Table 9-4](#).



### NOTE

You can change the key if needed.

----End

## Related Links

To learn more about encryption principles, see [EVS Encryption Overview](#).

## FAQs

With default encryption enabled, on the **Buy Disk** page, the console prompts me with the message "Check that the configured preset key is available and then try again" and the disk creation fails.

Possible cause: The key is disabled.

Solution:

Solution 1: On the disk list page, click **Configure** in the upper right corner. On the displayed page, click the **Data Encryption Protection** tab and change the key.

Solution 2: Enable the key. For details, see [Enabling a Custom Key](#).

---

# 10 Managing Shared EVS Disks

---

## What Is Disk Sharing?

Disk sharing allows you to create shared EVS disks. A shared EVS disk is a block storage device that can be attached to multiple cloud servers for concurrent reads/writes. Shared disks feature high concurrency, high performance, and high reliability. They are typically 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. To share files, you need to deploy a shared file system or a cluster management system first, such as Windows MSCS, Veritas VCS, or CFS.

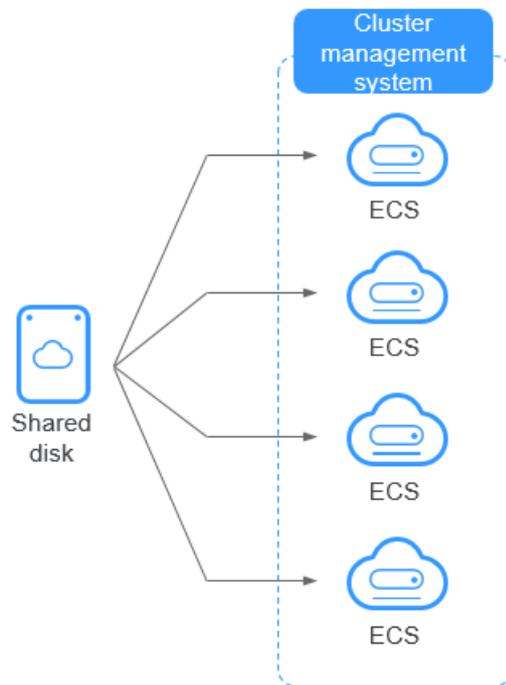
---

### 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.

## Billing

The sharing function is free. The EVS disks and ECSs required when you use the sharing function will be billed. For details about EVS disk billing, see [Billing for EVS Disks](#). For details about the ECS billing, see [ECS Billing](#).

## Creating a Shared EVS Disk

You need to select **Share** when [buy disks](#) on the console.

**Figure 10-2** Selecting Share

## 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](#).

## 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.

---

## Related 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.

- Backup policies: To perform automatic backups, configure a backup policy by setting the execution times of backup tasks, the backup frequency, and the retention rule, and then apply the policy to a vault.
- Replication policies: To automatically replicate backups or vaults, configure a replication policy by setting the execution times of replication tasks, the replication frequency, and the retention rule, and then apply the policy to a vault. Backup replicas are stored in replication vaults.

## Backup Mechanism

The first backup is a full backup. It 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. If a disk already has a backup, after another backup, the old snapshot will be deleted and the latest one will be retained.

CBR stores backups in OBS to ensure data security.

## Backup Options

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

**Table 11-1** One-off backup and periodic backup

Item	One-Off Backup	Periodic Backup
Backup policy	Not required	Required
Number of backup tasks	One manual backup task	Periodic tasks driven by a backup policy
Backup name	User-defined backup name, which is <b>manualbk_XXXX</b> by default	System-assigned backup name, which is <b>autobk_XXXX</b> by default

Item	One-Off Backup	Periodic Backup
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
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.

This section describes how to quickly purchase a disk backup vault and perform backup.

### Constraints

- Backups can be created only when the disks are in the **Available** or **In-use** state.
- Backup data can only be restored to original disks.
- Only users with the CBR FullAccess permissions can use the cloud disk backup function. If the user does not have the permissions, contact the account administrator to grant the permissions first.

### Creating a Disk Backup

Before backing up an EVS disk using CBR, you need to buy a disk backup vault, associate the disk with the vault, and then create a backup. [Table 11-2](#) describes the procedure.

**Table 11-2** Procedure of creating a disk backup

Procedure	Description	Operation Instruction
Step 1: Purchase a Disk Backup Vault	Buy a disk backup vault to store disk backups.	<a href="#">Creating a Disk Backup</a>
Step 2: Associate Disks with the Vault	Associate your disks with the vault to back up and replicate the disk data.	

Procedure	Description	Operation Instruction
Step 3: Create a Disk Backup	Create disk backups to protect data.	

## Related Information

- After the backup is complete, you can use it to restore disk data on the **Backups** tab. For details, see [Restoring from a Cloud Disk Backup](#).
- You can also use the disk backup to create a new disk. The newly created disk will contain the same data at the time of the backup. For details, see [Creating a Disk from a Cloud Disk Backup](#).

# 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 through APIs only. For more information, see [EVS Transfer](#).

## Constraints

- Monthly/yearly EVS disks cannot be transferred.
- Frozen EVS disks cannot be transferred.
- 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.
- DSS disks cannot be transferred.
- EVS disks cannot be transferred across regions.

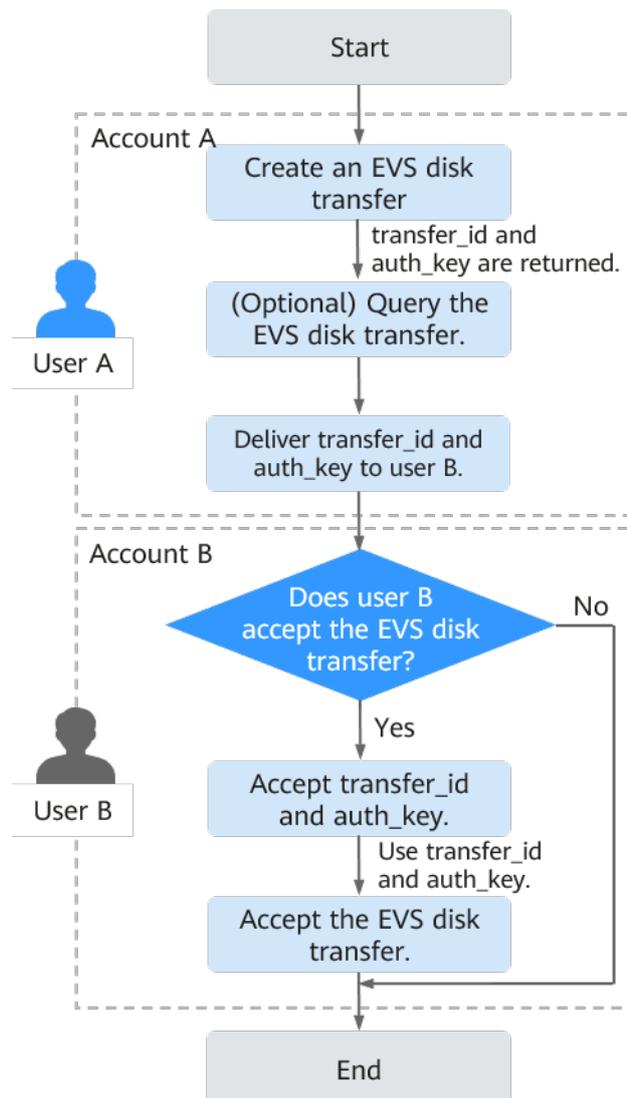
## Procedure

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 an EVS 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 an EVS Disk Transfer](#). If multiple disk transfers have been created, user A can query all disk transfers. For details, see [Querying All EVS Disk Transfers](#) or [Querying Details of All EVS 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 an EVS 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 an EVS 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
<a href="#">Adding a Tag</a>	Add tags for existing disks or during disk creations.
<a href="#">Modifying a Tag</a>	Change tag values for existing disks. Tag keys of existing disks cannot be changed.
<a href="#">Deleting a Tag</a>	Delete tags that are no longer needed for existing disks.
<a href="#">Searching for Disks by Tag</a>	After tags are added, search for disks by tags.

## 13.2 Adding a Tag

### Scenarios

You can add tags for an existing EVS disk. You can also add tags when creating a disk.

### 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.

## Constraints

- You can add a maximum of 20 tags for a single EVS disk.
- Tag keys of the same EVS disk must be unique.

## Procedure

**Step 1** Sign in to the [EVS console](#).

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

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

The disk details page is displayed.

**Step 4** Click the **Tags** tab.

**Step 5** Click **Add Tag**.

The **Add Tag** page is displayed.

**Step 6** Enter a key and a value for a tag and click **OK**.

- Tag key: This parameter is mandatory.
- Tag value: This parameter is optional.

The **Tags** tab is displayed, and you can view the newly added tag.

----End

## Follow-up Operations

After adding tags, you can search for disks by tag. For details, see [Searching for Disks by Tag](#).

## 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.

### Constraints

- You can add a maximum of 20 tags for a single EVS disk.
- Tag keys of the same EVS disk must be unique.

### Procedure

**Step 1** Sign in to the [EVS console](#).

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

**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

## Follow-up Operations

After adding tags, you can search for disks by tag. For details, see [Searching for Disks by Tag](#).

# 13.4 Deleting a Tag

## Scenarios

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

## Procedure

**Step 1** Sign in to the [EVS console](#).

**Step 2** In the disk list, locate the desired disk and click the disk name.

The disk details page is displayed.

**Step 3** Click the **Tags** tab.

**Step 4** Locate the target tag and click **Delete** in the **Operation** column.

The **Delete Tag** page is displayed.

**Step 5** 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** Sign in to the [EVS console](#).

**Step 2** 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

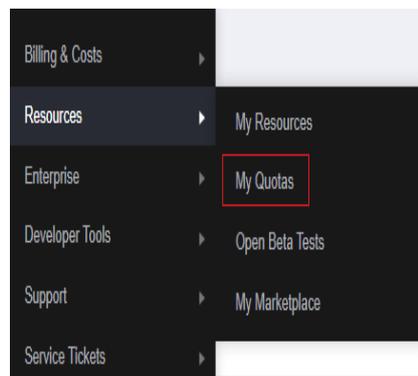
Quotas are enforced for cloud service resources to prevent unforeseen spikes in resource usage. There are preset quotas on the EVS disk quantity, EVS disk capacity, and EVS snapshot quantity.

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

### Procedure

- Step 1** Sign in to the [EVS console](#).
- Step 2** Click  in the upper left corner and select the desired region.
- Step 3** Click  in the upper left corner and choose **Storage > Elastic Volume Service**.
- Step 4** Choose **Resources > My Quotas** in the upper right corner of the page.  
The **Quotas** page is displayed.

**Figure 14-1** My Quotas



- Step 5** 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

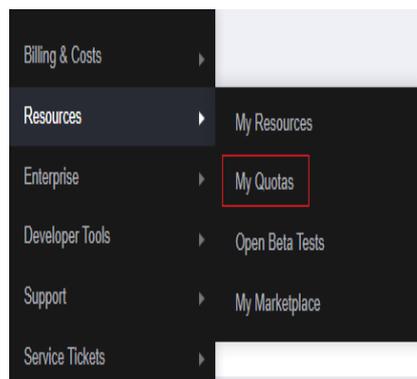
### Scenarios

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 **Quotas** 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

The image shows a screenshot of the 'Service Quota' page. At the top right, there is a red button labeled 'Increase Quota'. Below it is a table with the following columns: Service, Resource Type, Used Quota, and Total Quota.

Service	Resource Type	Used Quota	Total Quota
Auto Scaling	AS group	0	
	AS configuration	0	
Image Management Service	Image	0	
Cloud Container Engine	Cluster	0	
FunctionGraph	Function	0	
	Code storage(MB)	0	
Elastic Volume Service	Disk	3	
	Disk capacity(GB)	120	
Storage Disaster Recovery Service	Snapshots	4	
	Protection group	0	
Cloud Server Backup Service	Replication pair	0	
	Backup Capacity(GB)	0	
Scalable File Service	Backup	0	
	File system	0	
CCN	File system capacity(GB)	0	
	Domain name	0	
	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 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 how to configure alarm rules, see [Setting Alarm Rules](#).

### Namespace

SYS.EVS

### Metrics

If you find that your disk throughput or IOPS is high or exceeds the upper limit, you are advised to change the disk to a type with high specifications by referring to [Changing the EVS Disk Type \(OBT\)](#). To learn more about disk types and performance, see [Disk Types and Performance](#).

**Table 15-1** EVS metrics

Metric ID	Metric Name	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Period (Raw Data)
disk_device_read_bytes_rate	Disk Read Bandwidth	Number of bytes read from the monitored disk per second	$\geq 0$	bytes/s	1024 (IEC)	EVS disk	5 minutes (instantaneous value of the collection point)
disk_device_write_bytes_rate	Disk Write Bandwidth	Number of bytes written to the monitored disk per second	$\geq 0$	bytes/s	1024 (IEC)	EVS disk	5 minutes (instantaneous value of the collection point)
disk_device_read_requests_rate	Disk Read IOPS	Number of read requests sent to the monitored disk per second	$\geq 0$	requests/s	N/A	EVS disk	5 minutes (instantaneous value of the collection point)
disk_device_write_requests_rate	Disk Write IOPS	Number of write requests sent to the monitored disk per second	$\geq 0$	requests/s	N/A	EVS disk	5 minutes (instantaneous value of the collection point)

Metric ID	Metric Name	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Period (Raw Data)
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	≥ 0	count	N/A	EVS disk	5 minutes (instantaneous value of the collection point)
disk_device_io_util	Disk I/O Utilization	Percentage of time the monitored disk was performing read or write operations in the monitoring period	0-100	%	N/A	EVS disk	5 minutes (instantaneous value of the collection point)
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	≥ 0	KB/op	N/A	EVS disk	5 minutes (instantaneous value of the collection point)

Metric ID	Metric Name	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Period (Raw Data)
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	$\geq 0$	KB/opp	N/A	EVS disk	5 minutes (instantaneous value of the collection point)
disk_device_write_await	Disk Write Await	Average await time per I/O write for the monitored disk in the monitoring period	$\geq 0$	ms/opp	N/A	EVS disk	5 minutes (instantaneous value of the collection point)
disk_device_read_await	Disk Read Await	Average await time per I/O read for the monitored disk in the monitoring period	$\geq 0$	ms/opp	N/A	EVS disk	5 minutes (instantaneous value of the collection point)
disk_device_io_service_time	Disk I/O Service Time	Average service time per I/O read or write for the monitored disk in the monitoring period	$\geq 0$	ms/opp	N/A	EVS disk	5 minutes (instantaneous value of the collection point)

Metric ID	Metric Name	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Period (Raw Data)
disk_device_iops_qos_number	IOPS Upper Limit Reached (Count)	Number of times that the IOPS of the monitored disk has reached the upper limit	$\geq 0$	count	N/A	EVS disk	5 minutes (accumulated value)
disk_device_iobw_qos_number	Bandwidth Upper Limit Reached (Count)	Number of times that the bandwidth of the monitored disk has reached the upper limit	$\geq 0$	count	N/A	EVS disk	5 minutes (accumulated value)

## Dimension

Key	Value
disk_name	<ul style="list-style-type: none"> <li><i>ECS ID-Drive letter</i>, for example, <b>6f3c6f91-4b24-4e1b-b7d1-a94ac1cb011d-vda</b> (vda is the drive letter). It is suitable for KVM ECSs.</li> <li><i>ECS ID-volume-Disk ID</i>, for example, <b>6f3c6f91-4b24-4e1b-b7d1-a94ac1cb011d-volume-31f45764-38b3-44ad-aaca-4015c83371e6</b> Volume ID. It is suitable for Qingtian ECSs.</li> </ul> <p>You can obtain the ECS ID, volume ID, and drive letter by referring to <a href="#">Querying Details of a Single Disk Attached to an ECS</a>.</p>

## Viewing Monitoring Data

**Step 1** Sign in to the [EVS 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 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.

### Prerequisites

You have installed the Agent.

For instructions about how to install and configure the Agent, see [Installing and Configuring the Agent](#).

### Namespace

AGT.ECS

## Monitoring Metrics

Table 15-2 Disk metrics

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Period (Raw Data)
disk_free	(Agent) Available Disk Space	<p>Free space on the disks</p> <ul style="list-style-type: none"> <li>Linux: Run the <b>df -h</b> command to check the value in the <b>Avail</b> column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</li> <li>Windows: Use the WMI interface to call <code>GetDiskFreeSpaceExW</code> API to obtain disk space data. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</li> </ul>	≥0	GB	N/A	instance_id,mount_point	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Period (Raw Data)
disk_total	(Agent) Disk Storage Capacity	<p>Total space on the disks, including used and free</p> <ul style="list-style-type: none"> <li>Linux: Run the <b>df -h</b> command to check the value in the <b>Size</b> column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</li> <li>Windows: Use the WMI interface to call <code>GetDiskFreeSpaceExW</code> API to obtain disk space data. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</li> </ul>	≥0	GB	N/A	instance_id,mount_point	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Period (Raw Data)
disk_used	(Agent) Used Disk Space	<p>Used space on the disks</p> <ul style="list-style-type: none"> <li>Linux: Run the <b>df -h</b> command to check the value in the <b>Used</b> column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</li> <li>Windows: Use the WMI interface to call <code>GetDiskFreeSpaceExW</code> API to obtain disk space data. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</li> </ul>	≥0	GB	N/A	instance_id,mount_point	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Period (Raw Data)
disk_usedPercent	(Agent) Disk Usage	<p>Percentage of total disk space that is used, which is calculated as follows:  <b>Disk Usage = Used Disk Space/Disk Storage Capacity</b></p> <ul style="list-style-type: none"> <li>Linux: It is calculated as follows: Used/Size. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</li> <li>Windows: Use the WMI interface to call GetDiskFreeSpaceExW API to obtain disk space data. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</li> </ul>	0-100	%	N/A	instance_id,mount_point	1 minute

**Table 15-3** Disk I/O metrics

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Period (Raw Data)
disk_agt_read_bytes_rate	(Agent) Disks Read Rate	<p>Number of bytes read from the monitored disk per second</p> <ul style="list-style-type: none"> <li>Linux:                     <p>The disk read rate is calculated based on the data changes in the sixth column of the corresponding device in file <code>/proc/diskstats</code> in a collection period. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</p> </li> <li>Windows:                     <ul style="list-style-type: none"> <li>The disk I/O data is obtained through the Win32_PerformanceData_PerfDisk_LogicalDisk object in WMI. The object is obtained once in each collection period. The instantaneous value returned by the object indicates the</li> </ul> </li> </ul>	≥ 0	byte/s	1024(IEC)	<ul style="list-style-type: none"> <li>instance_id, disk</li> <li>instance_id, mount_point</li> </ul>	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Period (Raw Data)
		<p>metric value in a collection period.</p> <ul style="list-style-type: none"> <li>- The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</li> <li>- When the CPU usage is high, monitoring data obtaining timeout may occur and result in the failure of obtaining monitoring data.</li> </ul>					

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Period (Raw Data)
disk_agent_read_requests_rate	(Agent) Disks Read Requests	<p>Number of read requests sent to the monitored disk per second</p> <ul style="list-style-type: none"> <li>Linux: The disk read requests are calculated based on the data changes in the fourth column of the corresponding device in file <b>/proc/diskstats</b> in a collection period. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</li> <li>Windows: <ul style="list-style-type: none"> <li>The disk I/O data is obtained through the Win32_PerformanceData_PerfDisk_LogicalDisk object in WMI. The object is obtained once in each collection period. The instantaneous value returned by the object</li> </ul> </li> </ul>	≥ 0	Request/s	N/A	<ul style="list-style-type: none"> <li>instance_id, disk</li> <li>instance_id, mount_point</li> </ul>	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Period (Raw Data)
		<p>indicates the metric value in a collection period.</p> <ul style="list-style-type: none"> <li>- The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</li> <li>- When the CPU usage is high, monitoring data obtaining timeout may occur and result in the failure of obtaining monitoring data.</li> </ul>					

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Period (Raw Data)
disk_agent_write_bytes_rate	(Agent) Disks Write Rate	<p>Number of bytes written to the monitored disk per second</p> <ul style="list-style-type: none"> <li>Linux:                     <p>The disk write rate is calculated based on the data changes in the tenth column of the corresponding device in file <b>/proc/diskstats</b> in a collection period. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</p> </li> <li>Windows:                     <ul style="list-style-type: none"> <li>The disk I/O data is obtained through the Win32_PerformanceData_PerfDisk_LogicalDisk object in WMI. The object is obtained once in each collection period. The instantaneous value returned by the object indicates the</li> </ul> </li> </ul>	≥ 0	byte/s	1024(IEC)	<ul style="list-style-type: none"> <li>instance_id, disk</li> <li>instance_id, mount_point</li> </ul>	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Period (Raw Data)
		<p>metric value in a collection period.</p> <ul style="list-style-type: none"> <li>- The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</li> <li>- When the CPU usage is high, monitoring data obtaining timeout may occur and result in the failure of obtaining monitoring data.</li> </ul>					

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Period (Raw Data)
disk_agent_write_requests_rate	(Agent) Disks Write Requests	<p>Number of write requests sent to the monitored disk per second</p> <ul style="list-style-type: none"> <li>Linux: The disk write requests are calculated based on the data changes in the eighth column of the corresponding device in file <b>/proc/diskstats</b> in a collection period. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</li> <li>Windows: <ul style="list-style-type: none"> <li>The disk I/O data is obtained through the Win32_PerformanceData_PerfDisk_LogicalDisk object in WMI. The object is obtained once in each collection period. The instantaneous value returned by the object</li> </ul> </li> </ul>	≥ 0	Request/s	N/A	<ul style="list-style-type: none"> <li>instance_id, disk</li> <li>instance_id, mount_point</li> </ul>	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Period (Raw Data)
		<p>indicates the metric value in a collection period.</p> <ul style="list-style-type: none"> <li>- The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</li> <li>- When the CPU usage is high, monitoring data obtaining timeout may occur and result in the failure of obtaining monitoring data.</li> </ul>					

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Period (Raw Data)
disk_readTime	(Agent) Average Read Request Time	<p>Average amount of time that read requests have waited on the disks</p> <ul style="list-style-type: none"> <li>Linux: The average read request time is calculated based on the data changes in the seventh column of the corresponding device in file / <b>proc/diskstats</b> in a collection period. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</li> <li>Windows is not supported currently.</li> </ul>	≥ 0	ms/Count	N/A	<ul style="list-style-type: none"> <li>instance_id, disk</li> <li>instance_id, mount_point</li> </ul>	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Period (Raw Data)
disk_writeTime	(Agent) Average Write Request Time	<p>Average amount of time that write requests have waited on the disks</p> <ul style="list-style-type: none"> <li>Linux: The average write request time is calculated based on the data changes in the eleventh column of the corresponding device in file <code>/proc/diskstats</code> in a collection period. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</li> <li>Windows is not supported currently.</li> </ul>	$\geq 0$	ms/Count	N/A	<ul style="list-style-type: none"> <li>instance_id, disk</li> <li>instance_id, mount_point</li> </ul>	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Period (Raw Data)
disk_io Utils	(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> <ul style="list-style-type: none"> <li>Linux: The disk I/O usage is calculated based on the data changes in the thirteenth column of the corresponding device in file <code>/proc/diskstats</code> in a collection period. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</li> <li>Windows is not supported currently.</li> </ul>	0-100	%	N/A	<ul style="list-style-type: none"> <li>instance_id, disk</li> <li>instance_id, mount_point</li> </ul>	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Period (Raw Data)
disk_queue_length	(Agent) Disk Queue Length	<p>This metric reflects the disk usage in a specified period and can be used to evaluate the disk I/O performance. A larger value indicates a busier disk and poorer I/O performance.</p> <ul style="list-style-type: none"> <li>Linux: The metric value is calculated by dividing the data changes in the fourteenth column of the corresponding device in <b>/proc/diskstats</b> in a collection period by the metric collection period. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</li> <li>Windows is not supported currently.</li> </ul>	≥ 0	Count	N/A	<ul style="list-style-type: none"> <li>instance_id, disk</li> <li>instance_id, mount_point</li> </ul>	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Period (Raw Data)
disk_write_bytes_per_operation	(Agent) Average Disk Write Size	<p>Average number of bytes in an I/O write for the monitored disk in the monitoring period</p> <ul style="list-style-type: none"> <li>Linux: The average disk write size is calculated based on the data changes in the tenth column of the corresponding device to divide that of the eighth column in file <b>/proc/diskstats</b> in a collection period. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</li> <li>Windows is not supported currently.</li> </ul>	≥ 0	Byte/operation	N/A	<ul style="list-style-type: none"> <li>instance_id, disk</li> <li>instance_id, mount_point</li> </ul>	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Period (Raw Data)
disk_read_bytes_per_operation	(Agent) Average Disk Read Size	<p>Average number of bytes in an I/O read for the monitored disk in the monitoring period</p> <ul style="list-style-type: none"> <li>Linux: The average disk read size is calculated based on the data changes in the sixth column of the corresponding device to divide that of the fourth column in file / <b>proc/diskstats</b> in a collection period. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</li> <li>Windows is not supported currently.</li> </ul>	≥ 0	Byte/operation	N/A	<ul style="list-style-type: none"> <li>instance_id, disk</li> <li>instance_id, mount_point</li> </ul>	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Period (Raw Data)
disk_io_svctm	(Agent) Disk I/O Service Time	<p>Average time in an I/O read or write for the monitored disk in the monitoring period</p> <ul style="list-style-type: none"> <li>Linux: The average disk I/O service time is calculated based on the data changes in the thirteenth column of the corresponding device to divide the sum of data changes in the fourth and eighth columns in file <code>/proc/diskstats</code> in a collection period. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).</li> <li>Windows is not supported currently.</li> </ul>	≥ 0	ms/op	N/A	<ul style="list-style-type: none"> <li>instance_id, disk</li> <li>instance_id, mount_point</li> </ul>	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Period (Raw Data)
disk_device_used_percent	Block Device Usage	<p>Percentage of the physical disk usage of the monitored object. Calculation formula: Used storage space of all mounted disk partitions/Total disk storage space</p> <ul style="list-style-type: none"> <li>Collection method for Linux ECSs: Obtain the disk usage of each mount point, calculate the total disk storage space based on the disk sector size and the number of sectors, and then you can calculate the used storage space in total.</li> <li>Windows is not supported currently.</li> </ul>	0-100	%	N/A	instance_id,disk	1 minute

## Dimension

Key	Value
Server	<p>instance_id Server ID</p> <p>You can obtain the value by referring to <a href="#">Querying the Original Dimension Values in Server Monitoring</a>.</p>
Server - Disk	<p>disk Server disk</p> <p>You can obtain the value by referring to <a href="#">Querying the Original Dimension Values in Server Monitoring</a>.</p>

## 15.3 EVS Events Supported by Event Monitoring

Table 15-4 Supported EVS events

Event Source	Name space	Event Name	Event ID	Event Severity	Description	Handling Suggestion	Impact
EVS	SYS.EVS	Update disk	updateVolume	Minor	Update the name and description of an EVS disk.	No action is required.	None
		Expand disk	extendVolume	Minor	Expand the capacity of an EVS disk.	No action is required.	None
		Delete disk	deleteVolume	Major	Delete an EVS disk.	No action is required.	Deleted disks cannot be recovered.
		QoS upper limit reached	reachQoS	Major	The I/O latency increases as the QoS upper limits of the disk are frequently reached and flow control triggered.	Change the disk type to one with a higher specification.	The current disk may fail to meet service requirements.

## 15.4 Setting an Alarm Rule

### Scenarios

You can set alarm rules for EVS to customize the monitoring scope and notification policies and stay aware of the EVS disk statuses.

An alarm rule includes the alarm rule name, monitoring scope, monitoring metrics, alarm thresholds, monitoring interval, and whether to send notifications. This section describes how to set alarm rules.

## Procedure

1. Sign in to the [console](#).
2. Click  in the upper left corner and select your region and project.
3. Choose **Management & Governance > Cloud Eye**.
4. In the navigation pane on the left, choose **Alarm Management > Alarm Rules**.
5. On the **Alarm Rules** page, click **Create Alarm Rule** to create one, or modify an existing alarm rule.

The following describes how you can modify an existing alarm rule:

- a. Click the name of the target alarm rule.
- b. Click **Modify** in the upper right corner of the page.
- c. On the **Modify Alarm Rule** page, set parameters as required.
- d. Click **OK**.

After the alarm rule is modified, you will be notified when an alarm is triggered.

### NOTE

For more information about alarm rules, see [Cloud Eye User Guide](#).

# 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](#).