Elastic Volume Service

FAQs

Issue	01
Date	2023-08-04

Contents

1 Summary

General

- **How Do I Start Using a Newly Purchased Disk?**
- **Can EVS Disks Be Used Directly for Storage?**
- **Can EVS Disks Be Used Alone?**
- **How Can I View My Disk Details?**
- **Can I Change the AZ of My Disk?**
- **Can I Change the Disk Type, Device Type, or Sharing Attribute of My Disk?**
- **What Should I Do If an Error Occurs on My EVS Disk?**
- **Why Do Some of My EVS Disks Not Have WWN Information?**
- **How Can I Migrate Data from an EVS Disk?**
- **What Are the Differences Between System Disks and Data Disks?**
- **How Can I Upload Files to My EVS Disk?**

Billing

- **Will I Be Billed If I Have Purchased an EVS Disk But Not Used It?**
- **How Can I Stop Being Billed for My Disk?**
- **Can I Recover My Disk Data If the Disk Is Unsubscribed from, Deleted by Mistake, or Released After the Retention Period Ends?**
- **Will My EVS Disk Be Deleted When I Delete Its Server?**

Capacity Expansion

- **Can I Reduce or Temporarily Expand the Disk Capacity?**
- **What Are the Differences Between Expanding Capacity by Expanding an EVS Disk and Creating a New EVS Disk?**
- **Will My Disk Data Be Lost After I Expand the Disk Capacity?**
- **Can I Use Backups or Snapshots Created Before Capacity Expansion to Restore Data on Expanded Disks?**
- **Do I Need to Restart the Server After Expanding the Disk Capacity?**
- **Do I Need to Detach an EVS Disk Before Expanding Its Capacity?**

- **How Can I Improve My Disk Performance?**

Sharing

- **Do I Have to Deploy a Cluster to Use Shared Disks?**
- **How Many Servers Can I Attach a Shared Disk to?**
- **How Can I Attach a Shared Disk to Multiple Servers?**
- **Can I Attach a Shared Disk to Servers Running Different OSs?**

Snapshot

- **What Are the Typical Causes of a Snapshot Creation Failure?**
- **Does EVS Support Automatic Snapshot Creation?**
- **Can I Create Snapshots for Multiple Disks at a Time?**
- **How Is a Snapshot Created for My Disk?**
- **Why Can't I Roll Back My Disk Data from a Snapshot?**
- **Can I Roll Back the Data from a Snapshot After Reinstalling the OS or Formatting the Disk?**
- **How Is the Snapshot Size Calculated?**
- **Do Snapshots Take Space on the Disk?**

Backup

- **Do I Need to Stop the Server Before Performing a Backup?**
- **Can I Back Up and Restore My EVS Disk to a Different Region?**

2 General

2.1 How Do I Start Using a Newly Purchased Disk?

A newly purchased disk must be attached to a server and then initialized in the server OS before you can use it.

For more information, see **Quickly Creating and Using an EVS Disk**.

2.2 Can EVS Disks Be Used Directly for Storage?

No.

EVS disks must be attached to cloud servers before use. You cannot use EVS disks alone to store data.

2.3 Can EVS Disks Be Used Alone?

No.

EVS disks must be attached to servers before you can use them.

2.4 How Can I View My Disk Details?

To do so, perform the following operations:

Step 1 Log in to the management console.

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

The disk list page is displayed.

Step 3 Locate the row that contains the target disk and view the disk specifications, attributes, and billing mode.

Step 4 (Optional) Click the disk name to view more information, such as the disk backup and snapshot information.

View more information on the **Summary** tab.

----End

2.5 Can I Change the AZ of My Disk?

No.

The AZ of a disk cannot be changed after you have purchased the disk. If you want to change the AZ, unsubscribe or delete the disk and then purchase a new one.

2.6 Can I Change the Disk Type, Device Type, or Sharing Attribute of My Disk?

The following table describes whether the disk type, device type, sharing, and encryption attributes of a disk can be changed.

Table 2-1 EVS disk change description

Attribute	Allow Change	Change Direction
Disk type	No	- From common I/O to ultra-high I/O - From ultra-high I/O to common I/O - From high I/O to common I/O The preceding descriptions are examples only. EVS disk type cannot be changed.
Sharing	No	- From shared to non-shared - From non-shared to shared
Device type	No	- From SCSI to VBD - From VBD to SCSI
Encryption	No	- From non-encrypted to encrypted - From encrypted to non-encrypted

However, you can:

1. Create a backup for the disk.

Figure 2-1 Create Backup
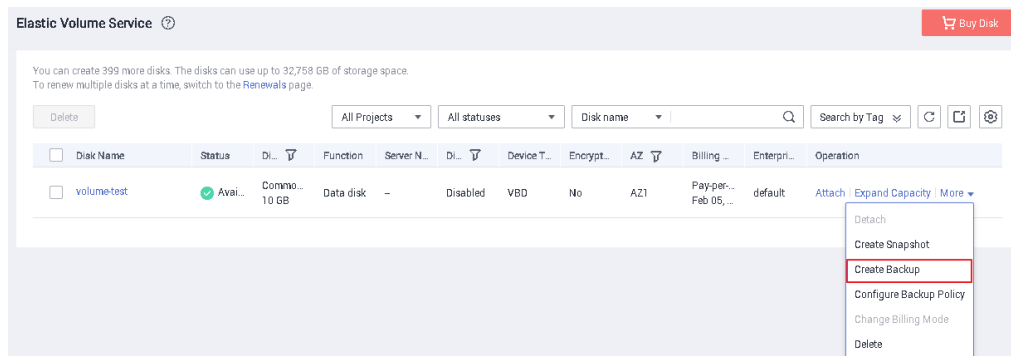
2. Create a new disk from this backup. During the creation, select a new disk type and configure advanced settings (sharing, SCSI, and encryption) based on your service requirements.

Figure 2-2 Create from backup
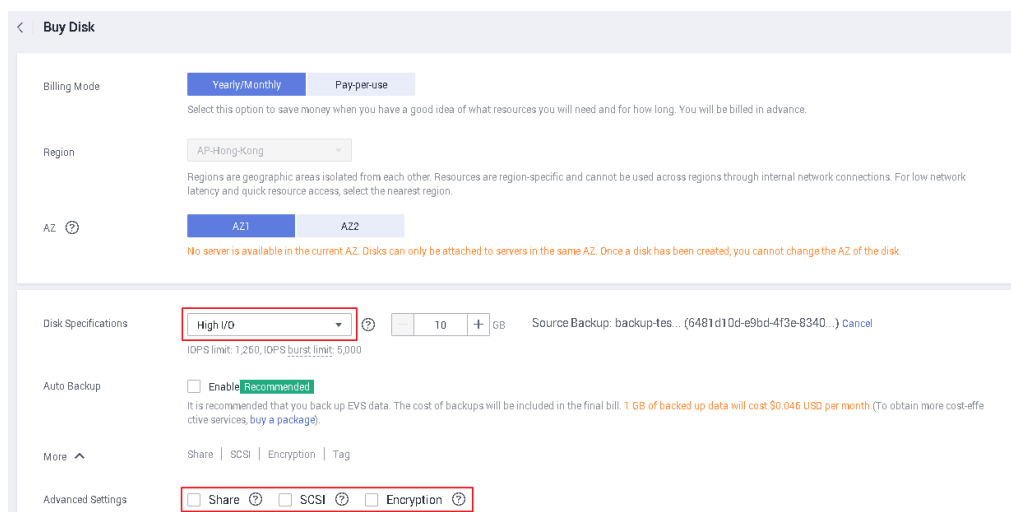
2.7 What Should I Do If an Error Occurs on My EVS Disk?

If an error occurs, the disk may show one of the states listed in the following table. Take the measures described in the table to handle the exceptions.

Table 2-2 Solutions for disk errors

Error Status	Handling Suggestion
Error	Delete the disk in the **Error** state and create another one.
Deletion failed	Contact customer service.

Error Status	Handling Suggestion
Expansion failed	Our customer service personnel will contact you and help you handle this error. Do not perform any operations on the disk before the customer service personnel contact you. If the issue is urgent, you can contact our customer service personnel directly.
Restoration failed	Our customer service personnel will contact you and help you handle this error. Do not perform any operations on the disk before the customer service personnel contact you. If the issue is urgent, you can contact our customer service personnel directly.
Rollback failed	You can roll back the snapshot data to the disk again.

2.8 How Can I Obtain ECS NIC Information?

Scenarios

This topic shows how to obtain the ECS NIC information, such as the ECS IP address.

Procedure

Step 1 Log in to the management console.

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

The ECS console is displayed.

Step 3 In the ECS list, click the target ECS name.

The ECS details page is displayed.

Step 4 Click the **NICs** tab.

The NIC details page is displayed.

Step 5 On the **NICs** tab, click ∨ to view ECS NIC information.

Step 6 To view the virtual IP address, click **Manage Virtual IP Address**.

The virtual IP address details page is displayed. You can find the bound virtual IP address based on the ECS NIC information.

----End

2.9 Why Do Some of My EVS Disks Not Have WWN Information?

EVS disks have two device types: VBD and SCSI. WWNs are used as the unique identifiers for SCSI EVS disks, and VBD EVS disks do not have WWNs.

You can view the WWN of a SCSI EVS disk on the management console. The details are as follows:

- If the SCSI EVS disk is brand new, you can view the disk WWN on the disk details page.

  Figure 2-3 shows the query result.

  Figure 2-3 Queried WWN information

  WWN: ⑦   6888603000038430fa17a17502223655

- If the SCSI EVS disk was created before the WWN feature rollout, the disk WWN will fail to be obtained.

  Figure 2-4 shows the query result.

  Figure 2-4 No WWN information

  WWN: ⑦   --

2.10 How Can I Migrate Data from an EVS Disk?

Data migration involves the following scenarios:

- Cross-AZ data migration: Disk data can be migrated from one AZ to another through disk backups. You can create backups for your disks using the CBR service, and then use these backups to create new disks in the target AZ. For details, see **Creating a Cloud Disk Backup** and **Using a Backup to Create a Disk**.

- Cross-region data migration: You can create a data disk image from the data disk in the current region and replicate the image to the other region. Then you can use the data disk image to create data disks in that other region. For details, see **Creating a Data Disk Image from an ECS**.

- Cross-account data migration: You can create a data disk image from the data disk under one account and then share the image with another account. Then you can use the data disk image under that second account to create data disks. For details, see **Creating a Data Disk Image from an ECS**.

2.11 What Are the Differences Between System Disks and Data Disks?

- A system disk runs the server OS. It is like drive C in a PC.

  When a server is created, a system disk is automatically purchased and attached. You cannot purchase a system disk separately. The maximum size of a system disk is 1,024 GiB.

- Data disks store the server data. They are like drive D, drive E, and drive F in a PC.

  Data disks can be purchased during or after the server purchase. If you purchase data disks during the server purchase, the system will automatically attach the data disks to the server. If you purchase data disks after the server

purchase, you need to manually attach the data disks. The maximum size of a data disk is 32,768 GiB.

If one system disk already meets your business needs, you do not need to purchase data disks. As your business grows, you can purchase data disks when needed.

If the disk paths in your service systems cannot be changed or are difficult to change, you are advised to purchase data disks according to your system planning.

2.12 Will I Lose My Disk Data If I Reinstall ECS OS, Change the OS, or Change the ECS Specifications?

Table 2-3 Impact

Item	OS Reinstallation	OS Change	Specifications Modification
Application scenario	Initialize an ECS. The ECS OS remains unchanged after OS change.	Change the OS of an ECS by changing its image. For details about OS change constraints, see **Changing the OS**.	Change ECS specifications, such as increasing the number of vCPUs or adding memory, to meet your service requirements.
Billing	OS reinstallation is free of charge. The ECS price remains unchanged.	OS change is free of charge. However, you will be billed based on your new image type after OS change. .	Modifying ECS specifications is free of charge. However, you will be billed based on the new specifications after modification.
IP address	The private IP address, EIP, and MAC address remain unchanged.	The private IP address, EIP, and MAC address remain unchanged.	The private IP address, EIP, and MAC address remain unchanged.
System disk	Reinstalling OS will clear the data in all partitions of the ECS system disk. Back up data before reinstalling the OS.	Changing OS will clear the data in all partitions of the ECS system disk. Back up data before changing the OS.	No impact on system disk.
Data disk	No impact on data disk.	No impact on data disk.	No impact on data disk.

Item	OS Reinstallation	OS Change	Specifications Modification
Backup	Back up data before reinstalling the OS to prevent data loss.	Back up data before changing the OS to prevent data loss.	Create a system disk snapshot before modifying ECS specifications to prevent data loss.

2.13 How Can I Download My EVS Disk Data to a Local PC?

EVS disk data cannot be directly saved to a local PC. It is recommended that you use a third-party tool, such as FTP, to download the data.

2.14 How Can I Export the Original Data After I Changed My Server OS from Windows to CentOS?

Solution:

1. Install the ntfsprogs software to enable Linux to access the NTFS file system.

   yum install ntfsprogs

2. View the data disks previously attached to Windows.

   parted –l

3. Mount the data disks.

   mount -t ntfs-3g *Data disk path Mount point*

2.15 What Are the Differences Between MBR and GPT Partition Styles?

Table 2-4 lists the common disk partition styles. In Linux, different partition styles require different partitioning tools.

Table 2-4 Disk partition styles

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

NOTICE

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

If you change the partition style after the disk has been used, the data on the disk will be cleared. Therefore, select an appropriate partition style when initializing the disk.

2.16 What Does the "reserveVolume" Trace Mean in CTS?

Before an EVS disk is attached, the system will call the reserveVolume EVS API to check whether the disk can be attached. If it can be attached, the system then changes the disk status to **attaching** to avoid conflicts with other operations.

2.17 How Can I Upload Files to My EVS Disk?

EVS disks must be attached to servers before you can use them. For how to upload files, see **How Do I Upload Files to My ECS?**

3 Billing

3.1 How Can I Stop Being Billed for My Disk?

Pay-per-use disk: Just delete the disk. You will not be billed for a disk after it is deleted. For details, see **Deleting EVS Disks**.

3.2 Will I Be Billed If I Have Purchased an EVS Disk But Not Used It?

Yes, you will be billed by disk capacity after you have purchased the disk. The disk billing is irrelevant to whether the disk is attached, how much disk space is used, or whether the attached server is stopped or not.

If your disks are no longer needed, to stop being billed for their use, take the following actions:

Pay-per-use disk: Just delete the disk. You will not be billed for a disk after it is deleted. For details, see **Deleting EVS Disks**.

3.3 Can I Recover My Disk Data If the Disk Is Unsubscribed from, Deleted by Mistake, or Released After the Retention Period Ends?

Check whether a disk backup is available.

- If a backup is available, you can use it to restore data. If you use a backup to create a disk, the disk will contain the backup data in its initial state. For details, see **Using a Backup to Create a Disk**.

  📖 **NOTE**

  Disks created from backups do not need to be initialized. Initialization destroys any data that was on the disk.

- If no backups are available, the disk data cannot be recovered.

3.4 Will My EVS Disk Be Deleted When I Delete Its Server?

- For pay-per-use disks:
  - If such a disk is separately purchased and has been attached, the system will prompt you whether to delete the disk when you delete the server, and you can make the decision based on your requirements.
  - If such disks are purchased together with a server, the system disk will be deleted when you delete the server. For the data disks, the system will prompt you whether to delete the disks, and you can make the decision based on your service requirements.
- For yearly/monthly disks:

  If the disks are purchased together with a server, they will be unsubscribed from when you unsubscribe from the ECS.

4 Capacity Expansion

4.1 Can I Reduce or Temporarily Expand the Disk Capacity?

No. The disk capacity can only be expanded, and temporary capacity expansion is not supported.

4.2 What Are the Differences Between Expanding Capacity by Expanding an EVS Disk and Creating a New EVS Disk?

The differences are as follows:

- Expanding an EVS disk is when you expand the capacity of an existing EVS disk. Some systems let you expand the capacity of EVS disks in use. In this case, services are not interrupted.

- If you create a new EVS disk and attach it to a server that already has an existing EVS disk, the new EVS disk and the original EVS disk are attached to the same server but independent from each other.

4.3 Will My Disk Data Be Lost After I Expand the Disk Capacity?

Data will not be deleted during a system disk or data disk capacity expansion. However, incorrect operations during an expansion may result in data loss or exceptions. Exercise caution when performing capacity expansions. You are advised to back up the disk data before expanding capacity.

- For details about disk backups, see **Cloud Backup and Recovery User Guide**.

- For details about disk snapshots, see **Creating a Snapshot (OBT)**.

4.4 Can I Use Backups or Snapshots Created Before Capacity Expansion to Restore Data on Expanded Disks?

Yes. If backups or snapshots have been created for disks before capacity was expanded, you can restore your disk data using these backups or snapshots after the capacity is expanded. Expansion operations do not affect backups and snapshots.

After the disk data is restored, the disk capacity is increased, but the additional space still needs to be partitioned and formatted before it can be used. You must log in to the server to extend the disk partition or file system.

To extend disk partitions and file systems, see the following sections:

- **Extending Disk Partitions and File Systems (Windows)**
- **Partition and File System Extension Preparations (Linux)**

4.5 Do I Need to Restart the Server After Expanding the Disk Capacity?

An EVS disk can be expanded either in the Available or In-use state. Expanding the disk capacity on the management console enlarges the disk capacity, but you still need to log in to the server and extend the disk partitions and file systems to make that additional space usable. You may need to restart the server during the partition and file system extension. The details are as follows:

- After expanding an In-use disk on the management console, log in to the server and view the disk capacity.
  - If the additional space can be viewed, you can extend the partition and file system and a restart is not required.
  - If the additional space cannot be viewed, the server OS may not be included in the **compatibility list**. In this case, you should stop and then start the server (do not restart the server). When the additional space can be viewed, extend the partition and file system.
- After expanding an Available disk on the management console, attach the disk to the server and extend the partition and file system on the server. In this case, a server restart is not required.

4.6 Do I Need to Detach an EVS Disk Before Expanding Its Capacity?

An expansion consists of two phases:

1. Expand the disk capacity on the management console.
   - A shared, in-use disk cannot be expanded. You must detach the shared disk from all its servers and then expand its capacity.

–   A non-shared, in-use disk can be expanded, and you can leave the disk attached during expansion as long as the following conditions are met:

- The disk's server is in the **Running** or **Stopped** state.

- The disk's server OS supports the expansion of In-use disks.

  📖 **NOTE**

  Only some server OSs support the expansion of In-use disks. For details, see **Expanding Capacity for an In-use EVS Disk**.

2. Log in to the server and create a new partition or allocate the additional space to one that is already there.

   –   In Windows, there are no partition extensions that require the partition to be unmounted first.

   –   In Linux:

   - When allocating the additional space to an existing partition, that is, extending an existing partition, you must use the **umount** command to unmount the partition first.

   - When allocating the additional space to a new partition, that is, creating a new partition, you do not need to unmount the existing partition.

4.7 What Should I Do If My Disk Capacity Exceeds 2 TiB After Expansion?

An EVS system disk can be as large as 1 TiB (1,024 GiB). You can expand the capacity of a system disk to up to 1 TiB.

An EVS data disk can be as large as 32 TiB (32,768 GiB).

- With MBR, any disk space in excess of 2 TiB cannot be allocated and used, because the maximum disk capacity supported by MBR is 2 TiB (2,048 GiB).

  In this case, if you want to expand the disk capacity to over 2 TiB, change the partition style from MBR to GPT. Ensure that the disk data has been backed up before changing the partition style because services will be interrupted and data on the disk will be deleted during this change.

- With GPT, you can expand the capacity of a data disk to up to 32 TiB because the maximum disk capacity supported by GPT is 18 EiB (19,327,352,832 GiB).

  If the in-use partition style is GPT, see the following methods:

  –   Windows:

  Extending Disk Partitions and File Systems (Windows Server)

  –   Linux:

  Extending Partitions and File Systems for Data Disks (Linux)

4.8 How Can I Allocate Newly Added Space to a New Partition?

Windows

In Windows, if your disk already has a D: drive and you want to create an E: drive, see **System Disk: Add Additional Space to New Volume (F:)** and **Data Disk: Add Additional Space to New Volume (E:)** in the following section:

Linux

In Linux, if your disk already has, for example, partition **/dev/vdb1** and you want to create a new partition, **/dev/vdb2**, and allocate the additional space to the new partition,

See **Creating a New MBR Partition** and **Creating a New GPT Partition** in the following sections:

- **Extending Partitions and File Systems for System Disks (Linux)**
- **Extending Partitions and File Systems for Data Disks (Linux)**

4.9 How Can I Allocate Newly Added Space to an Existing Partition?

Windows

In Windows, if your disk already has, for example, a D: drive and you want to add space to this volume, see the following section:

See **System Disk: Add Additional Capacity to New Volume (C:)** and **Data Disk: Add Additional Capacity to New Volume (D:)** in the following section:

Linux

In Linux, if your disk already has, for example, partition **/dev/vdb1** and you want to add space to this partition,

See **Extending an Existing MBR Partition** and **Extending an Existing GPT Partition** in the following sections:

- **Extending Partitions and File Systems for System Disks (Linux)**
- **Extending Partitions and File Systems for Data Disks (Linux)**

4.10 Why Did My Disk Capacity Remain Unchanged on the Server After Capacity Expansion?

After expanding disk capacity on the management console, you must log in to the server and extend the disk partition and file system for the extra capacity to become available. Otherwise, you cannot view the additional space on the server.

To extend disk partitions and file systems, see the following sections:

- **Extending Disk Partitions and File Systems (Windows)**
- **Partition and File System Extension Preparations (Linux)**

4.11 What Can't I Expand Capacity for My Disk?

Symptom

Capacity expansion is not allowed for the disk.

Troubleshooting

Possible causes are listed here in order of their probability.

If the fault persists after you have ruled out one cause, move on to the next one in the list.

Figure 4-1 Troubleshooting
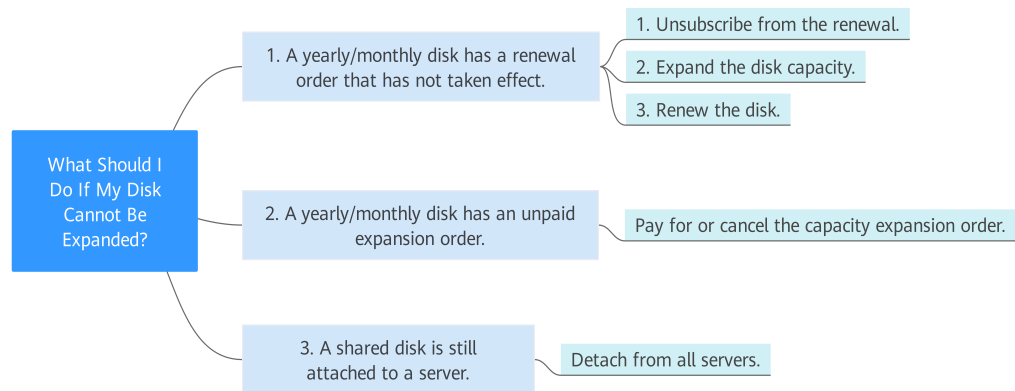
Table 4-1 Troubleshooting

Possible Cause	Solution
A yearly/monthly disk has a renewal order that has not taken effect.	See **Yearly/Monthly Disk Has a Renewal Order That Has Not Taken Effect**.
A yearly/monthly disk has an unpaid expansion order.	See **Yearly/Monthly Disk Has An Unpaid Expansion Order**.
A shared disk is still attached to a server.	See **Shared Disk Is Still Attached to a Server**.

Yearly/Monthly Disk Has a Renewal Order That Has Not Taken Effect

What Is a Renewal Order That Has Not Taken Effect?

For example, a user has purchased a yearly/monthly disk, and the disk expired on December 30, 2019, if the user renewed the disk on December 15, 2019, the order is a renewal order that will not take effect before December 30, 2019.

Symptom: A message is displayed in the upper right corner of the console, indicating that the disk cannot be expanded because it has a renewal order that has not taken effect.

Solution: Cancel the renewal order, expand the disk capacity, and renew the disk again.

Step 1 Log in to the management console and click **Billing Center** in the upper right corner.

Step 2 In the navigation pane on the left, choose **Unsubscriptions and Changes** > **Unsubscriptions**.

Step 3 Locate the target order and choose **More** > **Unsubscribe from Renewal** in the **Operation** column.

Step 4 After the unsubscription is complete, go to the EVS console and expand the disk capacity.

Step 5 Renew the disk after the capacity expansion succeeded.

----End

Yearly/Monthly Disk Has An Unpaid Expansion Order

Symptom: The **Expand Capacity** button is grayed out and the following hover tip is displayed: This operation cannot be performed because the disk has unpaid orders.

Solution: Pay for or cancel the capacity expansion order.

Step 1 On the EVS console, click **Unpaid** in the **Status** column.

The order details page is displayed.

Step 2 Click **Pay** or **Cancel Order**.

Step 3 Return to the EVS disk list. If the **Expand Capacity** button is available, the operation is successful.

----End

Shared Disk Is Still Attached to a Server

Symptom: The **Expand Capacity** button is grayed out. When you attempt to click the capacity expansion button, the following hover tip is displayed: This operation can be performed only when the shared disk is in the Available state.

Solution: Detach the disk from all servers. If the **Expand Capacity** button becomes available, you can expand the disk capacity.

4.12 How Do I Extend the File System of an Unpartitioned Data Disk in Linux?

Scenarios

If no partition but only a file system is created on a data disk, extend the file system according to the following operations:

Run the **lsblk** command. Information similar to the following is displayed:

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

In the command output, no partition but only a file system is created on data disk **/dev/vdb**.

In the following example, CentOS 7.4 64bit is used as the sample OS, data disk **/dev/vdb** has 10 GiB, no partition but only a file system is created on the disk, and additional 50 GiB has been added to this data disk on the management console. The following steps show how to extend this 50 GiB to the file system.

- **Extending the EXT* File System**
- **Extending the XFS File System**

The way you allocate additional space depends on the OS. This example is used for reference only. For the detailed operations and differences, see the corresponding OS documentations.

Extending the EXT* File System

Step 1 Run the following command to extend the file system:

resize2fs *Disk name*

In this example, run the following command:

resize2fs /dev/vdb

Information similar to the following is displayed:

```
[root@ecs-test ~]# resize2fs /dev/vdb
resize2fs 1.42.9 (28-Dec-2013)
Filesystem at /dev/vdb is mounted on /root/test; on-line resizing required
old_desc_blocs = 2, old_desc_blocs = 8
[17744.521535] EXT4-fs (vdb): resizing filesystem from 26214400 to 15728640 blocks
[17744.904470] EXT4-fs (vdb): resized filesystem to 15728640
The filesystem on /dev/vdb is now 15728640 blocks long.
```

Step 2 Run the following command to view the result:

df -TH

Information similar to the following is displayed:

```
[root@ecs-test ~]# df -TH
Filesystem    Type       Size  Used Avail Use% Mounted on
```

```
/dev/vda1     ext4     43G 1.9G  39G  5% /
devtmpfs      devtmpfs 2.0G   0 2.0G  0% /dev
tmpfs         tmpfs    2.0G   0 2.0G  0% /dev/shm
tmpfs         tmpfs    2.0G 9.1M 2.0G  1% /run
tmpfs         tmpfs    2.0G   0 2.0G  0% /sys/fs/cgroup
tmpfs         tmpfs    398M   0 398M  0% /run/user/0
/dev/vdb      ext4     64G  55M  61G  1% /mnt/sdc
```

----End

Extending the XFS File System

Step 1 Run the following command to extend the file system:

xfs_growfs *Disk name*

In this example, run the following command:

xfs_growfs /dev/vdb

Information similar to the following is displayed:

```
[root@ecs-test ~]# xfs_growfs /dev/vdb
meta-data=/dev/vdb            isize=512    agcount=4, agsize=655360 blks
         =                    sectsz=512   attr=2, projid32bit=1
         =                    crc=1        finobt=0, spinodes=0
data     =                    bsize=4096   blocks=2621440, imaxpct=25
         =                    sunit=0      swidth=0 blks
naming   =version2            bsize=4096   ascii-ci=0 ftype=1
log      =internal            bsize=4096   blocks=2560, version=2
         =                    sectsz=512   sunit=0 blks, lazy-count=1
realtime =none                extsz=4096   blocks=0, rtextents=0
data blocks changed from 2621440 to 15728640.
```

Step 2 Run the following command to view the result:

df -TH

Information similar to the following is displayed:

```
[root@ecs-test ~]# df -TH
Filesystem    Type     Size Used Avail Use% Mounted on
/dev/vda1     ext4     40G 2.3G  35G  7% /
devtmpfs      devtmpfs 1.9G   0 1.9G  0% /dev
tmpfs         tmpfs    1.9G   0 1.9G  0% /dev/shm
tmpfs         tmpfs    1.9G 8.6M 1.9G  1% /run
tmpfs         tmpfs    1.9G   0 1.9G  0% /sys/fs/cgroup
tmpfs         tmpfs    379M   0 379M  0% /run/user/0
/dev/vdb      xfs      60G  34M  60G  1% /mnt/sdc
```

----End

4.13 How Do I Extend the Root Partition of a Quickly Provisioned BMS?

Scenarios

If the root partition of your quickly provisioned BMS is too small, extend the root partition by referring to the following procedure.

This example uses CentOS 7.3 and system disk, **/dev/sdf**. The way you allocate additional space depends on the OS. This example is used for reference only. For detailed operations and differences, see the corresponding OS documentations.

In this example, the initial size of the BMS system disk (**sdf**) is 40 GiB and needs to be expanded to 140 GiB. The initial disk partitions are as follows:

```
sdf       8:80    0      40G    0 disk
├─sdf1    8:81    0     500M    0 part /boot
├─sdf2    8:82    0       5G    0 part [SWAP]
├─sdf3    8:83    0    34.5G    0 part /
└─sdf4    8:84    0      64M    0 part
```

Procedure

Step 1 Log in to the EVS console and expand the system disk capacity to 140 GiB.

Step 2 Log in to the BMS and run the following command to view the system disk capacity:

lsblk

Information similar to the following is displayed:

```
sdf       8:80    0     140G    0 disk
├─sdf1    8:81    0     500M    0 part /boot
├─sdf2    8:82    0       5G    0 part [SWAP]
├─sdf3    8:83    0    34.5G    0 part /
└─sdf4    8:84    0      64M    0 part
```

The system disk (**sdf**) has been expanded from 40 GiB to 140 GiB. The **sdf4** partition (64 MiB) is the configdriver partition that stores the BMS configuration information.

Step 3 Run the following command to back up the configdriver partition:

dd if=/dev/sdf4 of=/root/configdriver.img

Information similar to the following is displayed:

```
[root@bms-6acd ~]# dd if=/dev/sdf4 of=/root/configdriver.img
131072+0 records in
131072+0 records out
67108864 bytes (67 MB) copied, 0.291739 s, 230 MB/s
```

Step 4 Run the following command and delete the configdriver partition:

fdisk /dev/sdf

```
[root@bms-6acd ~]# fdisk /dev/sdf
Welcome to fdisk (util-linux 2.23.2).

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


Command (m for help): d
Partition number (1-4, default 4): 4
Partition 4 is deleted

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

Calling ioctl() to re-read partition table.

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

Step 5 Run the **partprobe** command to update the partition information.

If partition configdriver has been deleted, information similar to the following is displayed:

```
sdf       8:80    0    140G   0 disk
├─sdf1    8:81    0    500M   0 part /boot
├─sdf2    8:82    0      5G   0 part [SWAP]
├─sdf3    8:83    0   34.5G   0 part /
└─sdf4    8:84    0   97.7M   0 part
```

Step 6 Recreate the configdriver partition with 100 MB.

If the available sectors range from 83755008 to 293601279, set 293401279 (293601279 – 200000) as the new partition's start sector and 293601279 (default value) as the end sector.

```
Command (m for help): n
Partition type:
   p   primary (3 primary, 0 extended, 1 free)
   e   extended
Select (default e): p
Selected partition 4
First sector (83755008-293601279, default 83755008): 293401279
Last sector, +sectors or +size{K,M,G} (293401279-293601279, default 293601279):
Using default value 293601279
Partition 4 of type Linux and of size 97.7 MiB is set

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

Calling ioctl() to re-read partition table.

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

Run the **partprobe** command to update the partition information.

```
sdf       8:80    0    140G   0 disk
├─sdf1    8:81    0    500M   0 part /boot
├─sdf2    8:82    0      5G   0 part [SWAP]
├─sdf3    8:83    0   34.5G   0 part /
└─sdf4    8:84    0   97.7M   0 part
```

Step 7 Run the following command to extend the root partition:

growpart /dev/sdf 3

Information similar to the following is displayed:

```
[root@bms-6acd ~]# growpart /dev/sdf 3
CHANGED: partition=3 start=11511808 old: size=72243200 end=83755008 new: size=28
1889471,end=293401279
```

Run the **lsblk** command to view the new root partition size.

```
sdf        8:80    0    140G    0 disk
├─sdf1     8:81    0    500M    0 part /boot
├─sdf2     8:82    0      5G    0 part [SWAP]
├─sdf3     8:83    0  134.4G    0 part /
└─sdf4     8:84    0   97.7M    0 part
```

Step 8 Run the following command to extend the file system of the root partition:

resize2fs /dev/sdf3

Information similar to the following is displayed:

```
[root@bms-6acd ~]# resize2fs /dev/sdf3
resize2fs 1.42.9 (28-Dec-2013)
Filesystem at /dev/sdf3 is mounted on /; on-line resizing required
old_desc_blocks = 5, new_desc_blocks = 17
The filesystem on /dev/sdf3 is now 35236183 blocks long.

You have new mail in /var/spool/mail/root
[root@bms-6acd ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/sdf3       133G  1.9G  125G   2% /
devtmpfs         63G     0   63G   0% /dev
tmpfs            63G     0   63G   0% /dev/shm
tmpfs            63G  9.0M   63G   1% /run
tmpfs            63G     0   63G   0% /sys/fs/cgroup
/dev/sdf1       477M  104M  344M  24% /boot
tmpfs            13G     0   13G   0% /run/user/0
```

Step 9 Run the following command to restore the content of the configdriver partition:

dd if=/root/configdriver.img of=/dev/sdf4

Information similar to the following is displayed:

```
[root@bms-6acd ~]# dd if=/root/configdriver.img of=/dev/sdf4
131072+0 records in
131072+0 records out
67108864 bytes (67 MB) copied, 0.372614 s, 180 MB/s
[root@bms-6acd ~]#
```

```
/dev/sdf1: UUID="b9c472f9-6737-4200-910a-efa3af16629a" TYPE="ext4"
/dev/sdf2: UUID="b07ff4d0-8b0b-4c43-a40a-0b27290ea215" TYPE="swap"
/dev/sdf3: UUID="1e57f71e-6adc-4e98-9407-0f7d678d4525" TYPE="ext4"
/dev/sdf4: UUID="2018-09-27-19-13-01-00" LABEL="config-2" TYPE="iso9660"
[root@bms-6acd ~]#
```

The root partition of the quickly provisioned BMS has been extended.

----End

5 Attachment

5.1 Why Can't I View the Attached Data Disk on the Server?

Troubleshooting

Table 5-1 Possible causes

OS	Possible Cause	Solution
Linux	• New data disks are not formatted and partitioned by default, and an unformatted disk will not be listed in the command output. You must manually initialize the disk. • If a data disk cannot be found after the server is restarted, automatic partition mounting at system start may not be configured.	**Linux Data Disk**
Windows	New data disks are not formatted and partitioned by default. Only formatted and partitioned drives show up in the resource manager. You must manually initialize the disk.	**Windows Data Disk**

Linux Data Disk

Symptom: A data disk has been attached to a Linux server on the management console, but the disk cannot be viewed on the server.

Run **df -TH** to view the disk information. CentOS 7.4 is used in this example. The normal command output is as follows:

```
[root@ecs-test-0001 ~]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/vda1       ext4       43G  1.9G   39G   5% /
devtmpfs        devtmpfs  2.0G     0  2.0G   0% /dev
```

```
tmpfs       tmpfs    2.0G    0  2.0G   0% /dev/shm
tmpfs       tmpfs    2.0G  9.1M 2.0G   1% /run
tmpfs       tmpfs    2.0G    0  2.0G   0% /sys/fs/cgroup
tmpfs       tmpfs    398M    0  398M   0% /run/user/0
/dev/vdb1   ext4     106G   63M  101G   1% /mnt/sdc
```

Unlike the normal command output, only system disk **/dev/vda1** is visible, but data disk **/dev/vdb1** is missing from the command output.

Cause Analysis:

- **Cause 1**: New data disks are not formatted and partitioned by default, and an unformatted disk will not be listed in the command output. You must manually initialize the disk.

  For details, see **Introduction to Data Disk Initialization Scenarios and Partition Styles**.

- **Cause 2**: If a data disk cannot be found after the server is restarted, automatic partition mounting at system start may not be configured. Perform the following steps:

  a. Run the following command to mount the partition again:

     mount *Disk partition Mount point*

     In this example, run the following command:

     mount /dev/vdb1 /mnt/sdc

     Perform the following steps to enable automatic partition mounting at system start:

  b. Run the following command to query the partition UUID:

     blkid *Disk partition*

     In this example, run the following command to query the UUID of the **/dev/vdb1** partition:

     blkid /dev/vdb1

     Information similar to the following is displayed:

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

     The UUID of the **/dev/vdb1** partition is displayed.

  c. Run the following command to open the **fstab** file using the vi editor:

     vi /etc/fstab

     Press **i** to enter editing mode.

  d. Move the cursor to the end of the file and press **Enter**. Then, add the following information:
     ```
     UUID=0b3040e2-1367-4abb-841d-ddb0b92693df /mnt/sdc          ext4    defaults      0 2
     ```

     The preceding content is used for reference only. Add the information that is used in the environment. The parameters are described as follows:

     ▪ The first column indicates the partition UUID obtained in **b**.

     ▪ The second column indicates the directory on which the partition is mounted. You can query the mount point using the **df -TH** command.

     ▪ The third column indicates the file system format of the partition. You can query the file system format using the **df -TH** command.

■ The fourth column indicates the partition mount option. Normally, this parameter is set to **defaults**.

■ The fifth column indicates the Linux dump backup option.

○ **0**: Linux dump backup is not used. Normally, dump backup is not used, and you can set this parameter to **0**.

○ **1**: Linux dump backup is used.

■ The sixth column indicates the fsck option, that is, whether to use fsck to check the attached disk during startup.

○ **0**: not use fsck.

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

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

e. Press **Esc**, enter **:wq**, and press **Enter**.

The system saves the configurations and exits the vi editor.

Perform the following operations to verify the automatic mounting function:

i. Run the following command to unmount the partition:

umount *Disk partition*

In this example, run the following command:

umount /dev/vdb1

ii. Run the following command to reload all the content in the **/etc/fstab** file:

mount -a

iii. Run the following command to query the file system mounting information:

mount | grep *Mount point*

In this example, run the following command:

mount | grep /mnt/sdc

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

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

Windows Data Disk

Symptom: A data disk has been attached to a Windows server on the management console, but the disk cannot be viewed on the server. For example, Volume (D:) was not shown in **This PC** of a Windows server running Windows Server 2012. Normally, Volume (D:) appears, as shown in **Figure 5-1**.

Figure 5-1 Volume (D:) appears
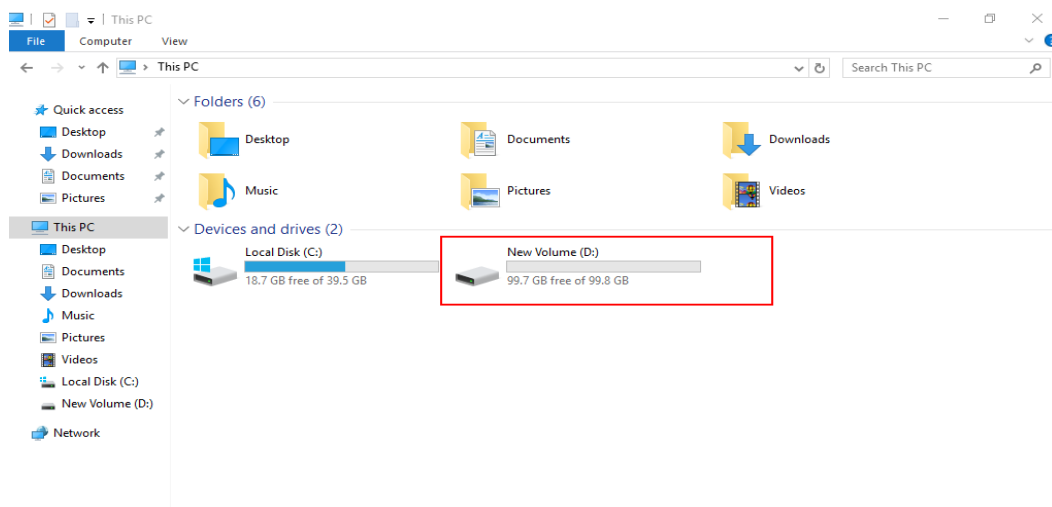
Solution: New data disks are not formatted and partitioned by default. Only formatted and partitioned drives show up in **This PC**. You must manually initialize the disk before it can be viewed here.

For details, see **Introduction to Data Disk Initialization Scenarios and Partition Styles**.

5.2 Why Can't I Attach My Disk to a Server?

Symptom

My disk cannot be attached to a server.

Troubleshooting

Possible causes are listed here in order of their probability.

If the fault persists after you have ruled out one cause, move on to the next one in the list.

Table 5-2 Troubleshooting

Symptom	Solution
The target server on the **Attach Disk** page could not be found.	- Go to **Check Whether the Disk and Server Are in the Same AZ**. - Cloud servers created from ISO images are only used for OS installation. They have limited functions and cannot have EVS disks attached.
The **Attach** button is grayed out.	- Go to **Maximum Number of Disks That Can Be Attached to the Server Has Been Reached**. - Go to **Check Whether the Disk Has Been Added to a Replication Pair**.

Symptom	Solution
An incorrect OS type warning is displayed when a shared disk is attached.	Go to **Check Whether the Servers Attached with the Shared Disk Are Running the Same Type of OS**.

Check Whether the Disk and Server Are in the Same AZ

Symptom: After you click **Attach**, the target server cannot be found on the **Attach Disk** page.

Solution: A disk can only be attached to a server in the same AZ and region. The **Attach Disk** page filters and shows all the servers that the disk can be attached to. Determine whether your disk data is required.

- If the disk data is no longer needed, delete or unsubscribe from the disk, and then purchase a new disk in the AZ where your target server is located.
- If the disk data is still required, create a new disk with the same data in the target AZ. The procedure is as follows:
  a. Create a backup for the disk.

     Figure 5-2 Create Backup
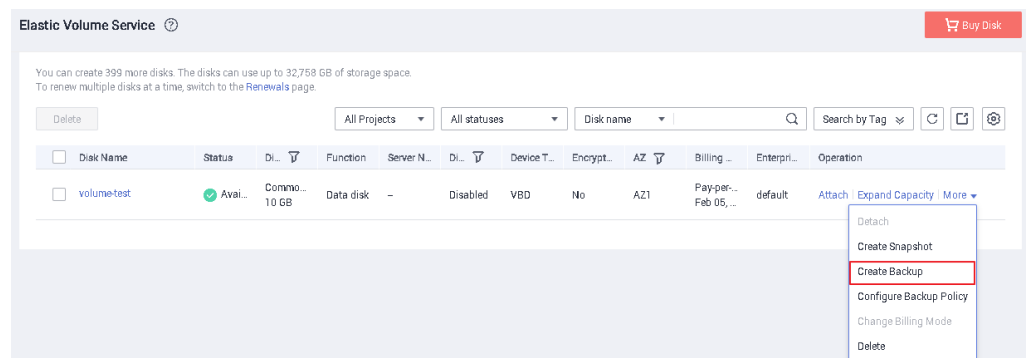
  b. Create a new disk from this backup. During the creation, select the target AZ. You can also change the settings of **Disk Type** and **Advanced Settings** if needed.

Figure 5-3 Create from backup
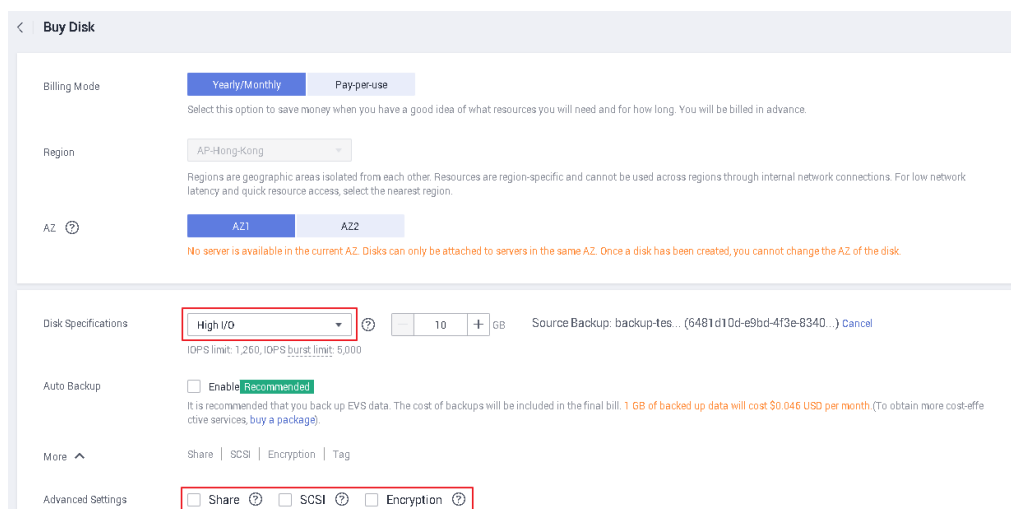
c. After the disk is created, click **Attach**. Your target server is displayed on the **Attach Disk** page.

Maximum Number of Disks That Can Be Attached to the Server Has Been Reached

Symptom: The **Attach** button is grayed out.

Solution:

● Non-shared disk: When you hover the mouse over the **Attach** button, message "This operation can be performed only when the disk is in the Available state" is displayed.

A non-shared disk can only be attached to one server. If the disk status is **In-use**, the disk has been attached. You can detach the disk, wait until the disk status changes to **Available**, and then attach the disk to the target server.

● Shared disk: When you hover the mouse over the **Attach** button, message "This operation cannot be performed because the maximum number of servers that a shared disk can be attached to has been reached" is displayed.

A shared disk can be attached to a maximum of 16 servers, but you can detach the shared disk from one server and attach it to a new one if needed.

Check Whether the Disk Has Been Added to a Replication Pair

Symptom: The **Attach** button is grayed out. When you hover the mouse over the **Attach** button, message "This operation cannot be performed on a disk in a replication pair" is displayed.

Solution: Delete the replication pair and attach the disk again.

Step 1 Choose **Storage** > **Storage Disaster Recovery Service**.

The **Storage Disaster Recovery Service** page is displayed.

Step 2 Locate the protection group containing the disk and click the protection group name.

The protection group details page is displayed.

Step 3 Click the **Replication Pairs** tab.

Check that the disk in the **Production Site Disk** column is the target disk.

Step 4 Confirm the information and click **Delete** in the **Operation** column.

Step 5 After the replication pair is deleted, return to the disk list, and the disk can be attached.

----End

Check Whether the Servers Attached with the Shared Disk Are Running the Same Type of OS

Symptom: After you click **Attach**, the target server cannot be selected on the **Attach Disk** page, and message "A shared disk must be attached to servers with the same OS type" is displayed.

Solution: This message indicates that the OS type of the target server is inconsistent with that of the servers attached with the shared disk. You can change the OS type based your service requirements.

5.3 Can I Attach a Disk to Multiple Servers?

- A non-shared disk can only be attached to one server.
- A shared disk can be attached to up to 16 servers. For more information, see **Shared EVS Disks and Usage Instructions**.

  > NOTE
  >
  > - Shared disks are a type of EVS disks that can be attached to multiple servers.
  > - To use shared disks, you must set up a shared file system or cluster management system. If you directly attach a disk to multiple servers, the disk sharing attribute does not work and data may be overwritten.

5.4 Can I Attach a Disk to a Server in Another AZ?

No.

Disks and the servers you attach the disks to must be in the same AZ. The same is true for shared disks. Shared disks can only be attached to the servers in the same AZ.

If your disk and server are in different AZs, there are two scenarios:

- The disk is empty, or data on the disk is no longer needed.

  Delete or unsubscribe from the disk. Then, purchase a new disk in the AZ where the server belongs. For how to delete disks, see **Deleting EVS Disks**.

- Data on the disk is required.

  Create a backup for the disk using the CBR service. Then, use this backup to create a new disk in the AZ where the server is. For details, see **Creating a Cloud Disk Backup** and **Using a Backup to Create a Disk**.

5.5 How Can I Add a Data Disk to an Existing Server?

Data disks can be purchased during or after the server purchase. If you purchase data disks during the server purchase, the system will automatically attach the data disks to the server. If you purchase data disks after the server purchase, you need to manually attach the data disks.

- On a Windows server:
  - If a data disk is purchased along with the server, you need to log in to the server and initialize the disk. The data disk will be visible after the initialization succeeds.
  - If no data disk is purchased along with the server, you need to purchase a data disk and attach it to the server. Then, you need to log in to the server and initialize the disk. The data disk will be visible after the initialization succeeds.
- On a Linux server:
  - If a data disk is purchased along with the server, you need to log in to the server and initialize the disk. The data disk will be visible after the initialization succeeds and the disk has been mounted via the **mount** command.
  - If no data disk is purchased along with the server, you need to purchase a data disk and attach it to the server. Then, you need to log in to the server and initialize the disk. The data disk will be visible after the initialization succeeds and the disk has been mounted via the **mount** command.

For details, see **Introduction to Data Disk Initialization Scenarios and Partition Styles**.

5.6 Can I Attach Different Types of Disks to the Same Server?

Yes. Different types of EVS disks can be attached to the same server. You only need to make sure that these disks and the server are in the same AZ.

5.7 Can I Switch Between System Disks and Data Disks?

☐ NOTE

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

Pay-per-Use Disks

Table 5-3 Detaching and reattaching of pay-per-use EVS disks

Original Use	Target Server	New Use	Supported
System disk (It becomes a bootable disk after being detached.)	Attach to the same server	System disk	√
		Data disk	√
	Attach to another server	System disk	√
		Data disk	√
Data disk	Attach to the same server	System disk	×
		Data disk	√
	Attach to another server	System disk	×
		Data disk	√

5.8 What Should I Do If a Linux EVS Disk Is Attached to a Windows Server?

You are not advised to attach a Linux EVS disk to a Windows server or attach a Windows EVS disk to a Linux server.

The disk information may fail to be displayed due to the inconsistent file systems. To solve this issue, you need to initialize and partition the disk again. Formatting the disk will destroy any data the disk has contained, so you should back up the data first.

5.9 Can I Change the Function of a System Disk or Data Disk Purchased Along with a Server?

This section describes how to change the function of a disk on the original server. To change the function of a disk and use it on a different server, see**Can I Switch Between System Disks and Data Disks?**.

- System disk purchased along with a server: You can detach the system disk and then re-attach it to the same server as a system disk or data disk.

  ◯ NOTE

  Each server can have only one system disk. If there is already a system disk attached to the server, then the detached system disk cannot be attached as a second system disk.

- Data disk purchased along with a server: You can detach the data disk and then re-attach it to the same server only as a data disk. It cannot be attached as a system disk.

5.10 How Do I Obtain My Disk Device Name in the ECS OS Using the Device Identifier Provided on the Console?

Scenarios

You find that the device name displayed in the ECS OS is different from that displayed on the management console and you cannot determine which disk name is correct. This section describes how to obtain the disk name used in an ECS OS according to the device identifier on the console.

For details about how to attach disks, see **Attaching an EVS Disk to an ECS**.

Obtaining the Disk ID of an ECS on the Console

1. Log in to the management console.
2. Under **Compute**, choose **Elastic Cloud Server**.
3. Click the target ECS name in the ECS list.

   The ECS details page is displayed.

4. Click the **Disks** tab and then click ⌄ to expand the disk information.
5. Check the device type and ID of the disk.

   📖 NOTE

   If **Device Identifier** is not displayed on the page, stop the ECS and restart it.

   – KVM ECS

     ■ If **Device Type** is **VBD**, use a serial number or BDF to obtain the disk device name.

       If you use a serial number (recommended) to obtain the disk name, see **Using a Serial Number to Obtain a Disk Device Name (Linux)**.

       If you use a BDF to obtain the disk device name, see **Using a VBD to Obtain a Disk Device Name (Linux)**.

     ■ If **Device Type** is **SCSI**, use a WWN to obtain the disk name. For details, see **Using a WWN to Obtain a Disk Device Name (Linux)**.

Using a Serial Number to Obtain a Disk Device Name (Linux)

If a serial number is displayed on the console, run either of the following commands to obtain the device name.

 # udevadm info --query=all --name=/dev/xxx | grep ID_SERIAL

 # ll /dev/disk/by-id/*

📖 NOTE

A serial number is the first 20 digits of a disk UUID.

For example, if the serial number of the VBD disk is 62f0d06b-808d-480d-8, run either of the following commands:

 # udevadm info --query=all --name=/dev/vdb | grep ID_SERIAL

 # ll /dev/disk/by-id/*

The following information is displayed:

```
[root@ecs-ab63 ~]# udevadm info --query=all --name=/dev/vdb | grep ID_SERIAL
E: ID_SERIAL=62f0d06b-808d-480d-8
[root@ecs-ab63 ~]# ll /dev/disk/by-id/*
lrwxrwxrwx 1 root root  9 Dec 30 15:56 /dev/disk/by-id/virtio-128d5bfd-f215-487f-9 -> ../../vda
lrwxrwxrwx 1 root root 10 Dec 30 15:56 /dev/disk/by-id/virtio-128d5bfd-f215-487f-9-part1 -> ../../vda1
lrwxrwxrwx 1 root root  9 Dec 30 15:56 /dev/disk/by-id/virtio-62f0d06b-808d-480d-8 -> ../../vdb
```

/dev/vdb is the disk device name.

Using a VBD to Obtain a Disk Device Name (Linux)

1. Run the following command to use a BDF to obtain the device name:

   ll /sys/bus/pci/devices/_BDF disk ID_**/virtio*/block**

   For example, if the BDF disk ID of the VBD disk is 0000:02:02.0, run the following command to obtain the device name:

   ll /sys/bus/pci/devices/0000:02:02.0/virtio*/block

   The following information is displayed:

```
[root@ecs-ab63 ~]# ll /sys/bus/pci/devices/0000:02:02.0/virtio*/block
total 0
drwxr-xr-x 8 root root 0 Dec 30 15:56 vdb
```

   /dev/vdb is the disk device name.

Using a WWN to Obtain a Disk Device Name (Linux)

1. Log in as user root.

2. Run the following command to view the disk device name:

   ll /dev/disk/by-id |grep _WWN_**|grep scsi-3**

   For example, if the WWN obtained on the console is 6888603000008b32fa16688d09368506, run the following command:

   ll /dev/disk/by-id |grep 6888603000008b32fa16688d09368506|grep scsi-3

   The following information is displayed:

```
[root@host-192-168-133-148 block]# ll /dev/disk/by-id/ |grep 6888603000008b32fa16688d09368506 |
grep scsi-3
lrwxrwxrwx 1 root root  9 May 21 20:22 scsi-36888603000008b32fa16688d09368506 -> ../../sda
```

6 Detachment

6.1 If I Detach a Disk, Will I Lose the Data on My Disk?

It depends on if the disk is encrypted or not.

- Encrypted
  - The custom key is disabled or scheduled for deletion.

    The disk can still be used, but there is no guarantee for how long it will be usable. If the disk is detached, it will be impossible to re-attach it later. In this case, do not detach the disk without a working custom key.

    The restoration method varies depending on the key status. For details, see **EVS Encryption**.
  - The custom key is available.

    The disk can be detached and re-attached, and data on the disk will not be lost.

    To ensure your data safety, you are advised to follow the instructions described in **Disk Detachment Process**.

- Non-encrypted

  Data on a disk will not be lost after the disk is detached, and the disk can be re-attached later if needed.

  To ensure your data safety, you are advised to follow the instructions described in **Disk Detachment Process**.

- If you detach a disk that has an ongoing backup task, the task will fail.

Disk Detachment Process

- For disks not supporting online detachment:
  a. Stop the server where the disk was attached.
  b. After the server has stopped, detach the disk.
- For disks supporting online detachment:

  Detach the disk from a running ECS. For details, see **Management** > **Detaching an EVS Disk from a Running ECS** in the *Elastic Cloud Server User Guide*.

6.2 Why Can't I Detach My Disk?

EVS disks can be used as system disks or data disks, but the way you detach each one is different.

- System disks: A system disk can only be detached offline. You must first stop the server that uses this system disk and then detach the disk.

  □ NOTE

  In Linux, a system disk is typically mounted on **/dev/vda**. In Windows, a system disk is normally **Volume (C:)**.

- Data disks: A data disk can be detached regardless of whether it is offline or online.

  □ NOTE

  In Linux, a data disk is typically mounted on a mount point other than **/dev/vda**. In Windows, a data disk is normally a volume other than **Volume (C:)**.

  - Offline detachment: The server must be in the **Stopped** state. If it is not, stop the server and then detach the data disk.

  - Online detachment: Some OSs support online detachment. In this case, you do not need to stop the server before detaching the data disk. For more information, see **Detaching an EVS Disk from a Running ECS**.

7 Deletion

7.1 How Can I Recover Data from a Disk That Was Accidentally Deleted?

Check whether the disk has any snapshots or backups created.

- If there are, use the snapshot or backup to restore the disk data to the state when the snapshot or backup was created. For details, see **Rolling Back Data from a Snapshot (OBT)** or **Restoring from a Cloud Disk Backup**.

  > **NOTICE**
  >
  > If the disk was deleted after the last snapshot or backup was created, that incremental data cannot be restored.

- If there are not, the disk data cannot be restored.

7.2 Why Can't I Delete My Disk?

Check whether the following conditions are met:

- Pay-per-use disk:

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

  – The disk status is **Available**, **Error**, **Expansion failed**, **Restoration failed**, or **Rollback failed**.

  – The disk is not added to any replication pair in the Storage Disaster Recovery Service (SDRS). For any disk already added to a replication pair, you need to first and then delete the disk.

  – The disk is not locked by any service.

  – The shared disk has been detached from all its servers.

8 Capacity

8.1 What Is the Maximum Capacity Supported for the System and Data Disks?

The maximum capacity supported for a system disk is 1024 GiB.

The maximum capacity supported for a data disk is 32768 GiB.

8.2 What Should I Do If My Disk Starts to Run Out of Space?

If your disk space starts to fill up, you can:

● Create a new disk and attach it to the server. For details, see **Purchase an EVS Disk**

● Expand the capacity of the existing disk. Both system disks and data disks can be expanded. For details, see **Disk Capacity Expansion**

● Free up the space on the disk by following the directions in:Clear up the disk space. For details, see **Handling Insufficient Disk Space on a Windows ECS**.

Differences Between Expanding an EVS Disk and Creating an EVS Disk

The differences are as follows:

● Expanding an EVS disk is when you expand the capacity of an existing EVS disk. Some systems let you expand the capacity of EVS disks in use. In this case, services are not interrupted.

● If you create a new EVS disk and attach it to a server that already has an existing EVS disk, the new EVS disk and the original EVS disk are attached to the same server but independent from each other.

8.3 What Can I Do If the Capacity of My Disk Reaches the Maximum But I Still Need More Space?

The capacity of a single disk can be expanded to up to 32 TiB. If this still fails to meet your needs, it is recommended that you **create RAID arrays with EVS disks** or **manage EVS disks using LVM**.

8.4 What Should I Do If I Use fdisk to Initialize a Disk Larger Than 2 TiB and Then the Space in Excess of 2 TiB Cannot Be Displayed?

If your disk capacity is greater than 2 TiB, do not use fdisk to partition the disk. Or any space in excess of 2 TiB will be unable to show up after the disk is partitioned.

In this case, use parted to repartition the disk and choose the GPT partition style because MBR does not support disks over 2 TiB.

For details, see **Introduction to Data Disk Initialization Scenarios and Partition Styles**.

8.5 How Can I View My Disk Usage?

You can view your disk usages in either of the following ways:

- View disk usages manually.

  The details depend on the OS. This FAQ uses Windows Server 2008, Windows Server 2016, and Linux as samples to describe how to view the disk usage.

  - **Viewing Disk Usage in Linux**
  - **Viewing Disk Usage in Windows Server 2008**
  - **Viewing Disk Usage in Windows Server 2016**

- **Installing Agent to View Disk Usage**

Viewing Disk Usage in Linux

In this section, CentOS 7.4 64bit is used as an example. The details depend on if you need to view the available space or not.

- To query the total capacity only, run **lsblk**.

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

  In the command output, the server has two disks, **/dev/vda** and **/dev/vdb**. System disk **/dev/vda** has 40 GiB of capacity, as does data disk **/dev/vdb**.

- To query the total capacity and display the space available as well, run **df -TH**. Ensure that the disk has been attached and initialized before running this command.

  Information similar to the following is displayed:

  ```
  [root@ecs-0001 ~]# df -TH
  Filesystem     Type      Size  Used Avail Use% Mounted on
  /dev/vda1      ext4       43G  2.0G  39G   5% /
  devtmpfs       devtmpfs  509M    0  509M  0% /dev
  tmpfs          tmpfs     520M    0  520M  0% /dev/shm
  tmpfs          tmpfs     520M 7.2M  513M  2% /run
  tmpfs          tmpfs     520M    0  520M  0% /sys/fs/cgroup
  tmpfs          tmpfs     104M    0  104M  0% /run/user/0
  /dev/vdb1      ext4       43G  51M   40G   1% /mnt/sdc
  ```

  In the command output, the server has two partitions, **/dev/vda1** and **/dev/vdb1**. Partition **/dev/vda1** is used to deploy the OS, and its total capacity, used capacity, and available capacity are 43 GiB, 2 GiB, and 39 GiB, respectively. Partition **/dev/vdb1**'s total capacity, used capacity, and available capacity are 43 GiB, 51 MiB, and 40 GiB, respectively.

Viewing Disk Usage in Windows Server 2008

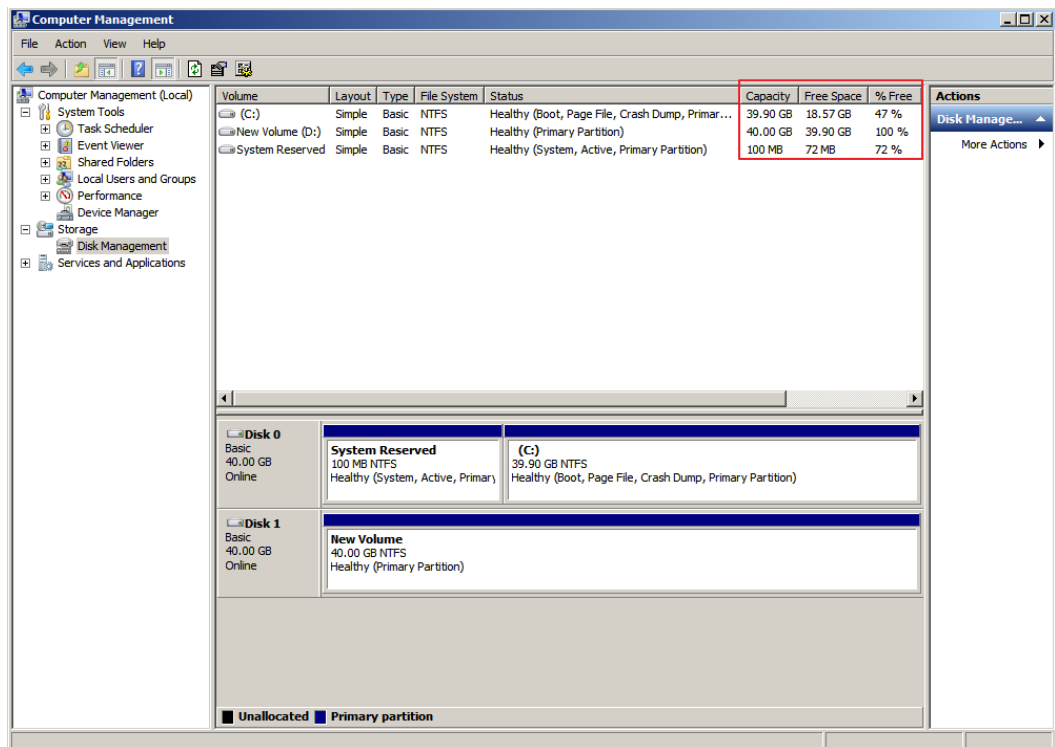In this section, Windows Server 2008 R2 Enterprise 64bit is used as an example.

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

The **Server Manager** window is displayed.

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

The sizes and available spaces of the volumes on the current disk are displayed in the middle pane.

Figure 8-1 Disk Management page

----End

Viewing Disk Usage in Windows Server 2016

In this section, Windows Server 2016 Standard 64bit is used as an example.

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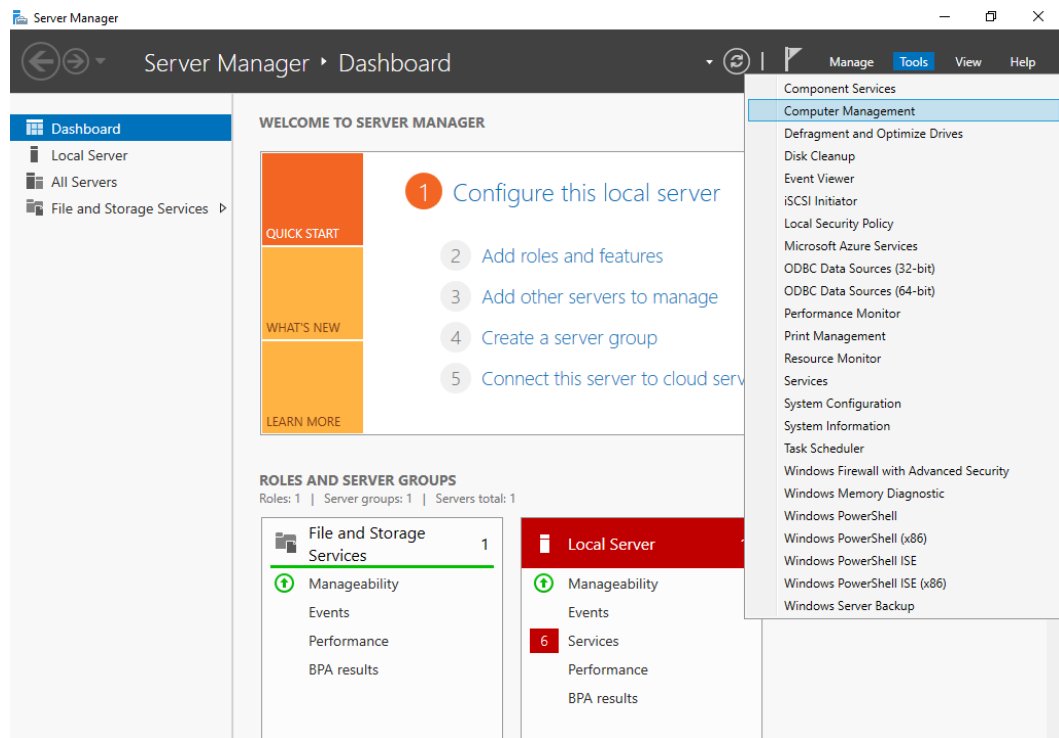
The **Windows Server** window is displayed.

Step 2 Click **Server Manager**.

The **Server Manager** window is displayed.
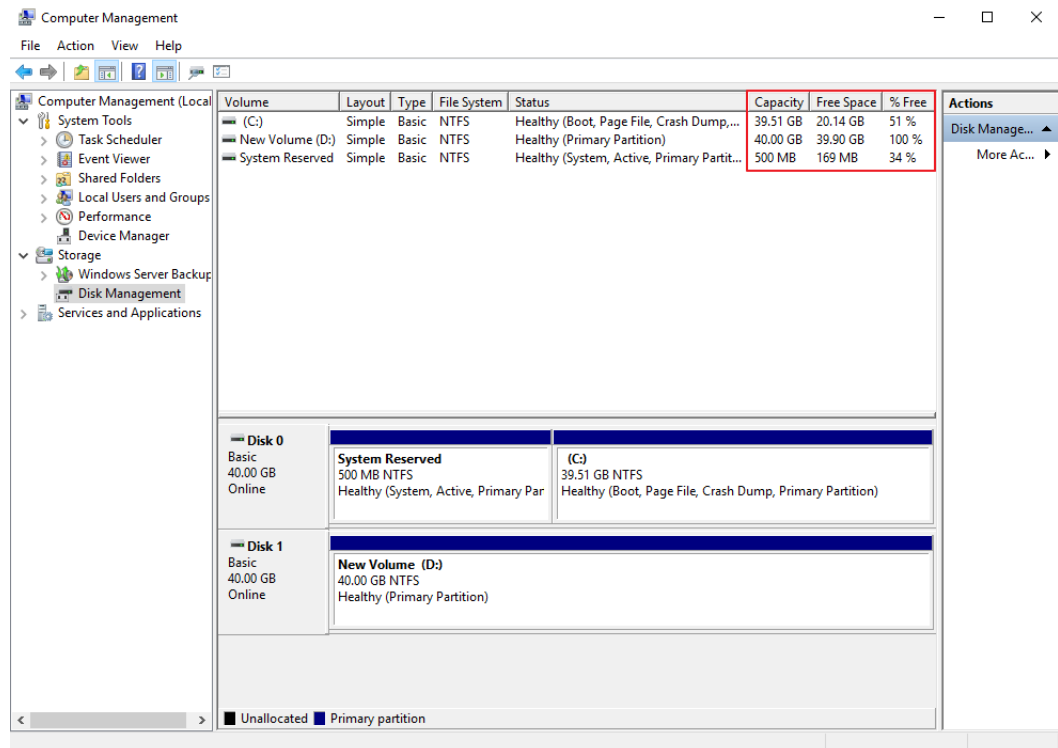
Figure 8-2 Server Manager page

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

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

In the middle pane, you can view the sizes and available spaces of the volumes on the disk.

Figure 8-3 Disk list page

----**End**

Installing Agent to View Disk Usage

Some disk monitoring metrics require that the agent to be installed.

● For instructions about how to install the agent on a Windows ECS, see "Installing and Configuring the Agent (Windows)" in the *Cloud Eye User Guide*.

● For instructions about how to install the agent on a Linux ECS, see "Installing and Configuring the Agent (Linux)" in the *Cloud Eye User Guide*.

Table 8-1 Disk metrics

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_free	(Agent) Available Disk Space	Free space on the disks Unit: GB ● Linux: Run the **df -h** command to check the value in the **Avail** column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). ● Windows: Use the WMI interface to call GetDiskFreeSpaceExW API to obtain disk space data. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	≥ 0	ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_total	(Agent) Disk Storage Capacity	Total space on the disks, including used and free Unit: GB - Linux: Run the **df -h** command to check the value in the **Size** column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). - Windows: Use the WMI interface to call GetDiskFreeSpaceExW API to obtain disk space data. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	≥ 0	ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_used	(Agent) Used Disk Space	Used space on the disks Unit: GB ● Linux: Run the **df -h** command to check the value in the **Used** column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). ● Windows: Use the WMI interface to call GetDiskFreeSpaceExW API to obtain disk space data. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	≥ 0	ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_usedPercent	(Agent) Disk Usage	Percentage of total disk space that is used, which is calculated as follows: **Disk Usage** = **Used Disk Space**/**Disk Storage Capacity** Unit: percent ● Linux: It is calculated as follows: Used/Size. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). ● Windows: Use the WMI interface to call GetDiskFreeSpaceExW API to obtain disk space data. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	0-100	ECS - Mount point	1 minute

8.6 How Can I Monitor My Disk Usages?

Some disk monitoring metrics require that the agent to be installed.

● For instructions about how to install the agent on a Windows ECS, see "Installing and Configuring the Agent (Windows)" in the *Cloud Eye User Guide*.

● For instructions about how to install the agent on a Linux ECS, see "Installing and Configuring the Agent (Linux)" in the *Cloud Eye User Guide*.

Table 8-2 Disk metrics

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_free	(Agent) Available Disk Space	Free space on the disks Unit: GB ● Linux: Run the **df -h** command to check the value in the **Avail** column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). ● Windows: Use the WMI interface to call GetDiskFreeSpaceExW API to obtain disk space data. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	≥ 0	ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_total	(Agent) Disk Storage Capacity	Total space on the disks, including used and free Unit: GB ● Linux: Run the **df -h** command to check the value in the **Size** column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). ● Windows: Use the WMI interface to call GetDiskFreeSpaceExW API to obtain disk space data. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	≥ 0	ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_used	(Agent) Used Disk Space	Used space on the disks Unit: GB ● Linux: Run the **df -h** command to check the value in the **Used** column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). ● Windows: Use the WMI interface to call GetDiskFreeSpaceExW API to obtain disk space data. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	≥ 0	ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_usedPercent	(Agent) Disk Usage	Percentage of total disk space that is used, which is calculated as follows: **Disk Usage** = **Used Disk Space**/**Disk Storage Capacity** Unit: percent • Linux: It is calculated as follows: Used/Size. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). • Windows: Use the WMI interface to call GetDiskFreeSpaceExW API to obtain disk space data. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	0-100	ECS - Mount point	1 minute

8.7 Can I Transfer the Data Disk Capacity to a System Disk?

Sorry, you cannot.

Currently, the capacity of an EVS disk cannot be transferred to another disk. Multiple EVS disks cannot be combined into a single, larger disk, either.

Common Scenarios

1. You want to expand the system disk capacity, but purchased a new data disk.

2. You want to expand the system disk capacity, but expanded the data disk capacity.

Recommended Solution

- If you do not need the data on the data disk, you can just delete or unsubscribe from the data disk, and then **expand the system disk**.

- If you need the data on the data disk, you can purchase a small-capacity data disk, copy data from the original data disk to the new data disk, and then expand the system disk capacity.

  a. Back up the data disk using the CBR service or by creating a snapshot.

     For details about backups, see **Managing EVS Backups**. For details about snapshots, see **Creating a Snapshot (OBT)**.

  b. Buy a new data disk with the desired capacity and attach it to the server. After initializing the disk, copy the data from the original data disk to the new data disk.

  c. Confirm that the services on the new data disk are available. Then, delete or unsubscribe from the original data disk, and delete the created backup.

  d. Expand the system disk capacity by referring to **Disk Capacity Expansion**

8.8 Why the Space of My New Disk Is Full After I Uploaded Only 500 MB of Files to the Disk?

Troubleshoot this issue by performing the following steps:

1. Check whether the disk partition usage is 100% or almost 100%.
   df -h

   Figure 8-4 Checking the partition usage

   ```
   [root@ecs-a058 /]# df -h
   Filesystem      Size  Used Avail Use% Mounted on
   devtmpfs        7.8G     0  7.8G   0% /dev
   tmpfs           7.8G   20K  7.8G   1% /dev/shm
   tmpfs           7.8G  8.7M  7.8G   1% /run
   tmpfs           7.8G     0  7.8G   0% /sys/fs/cgroup
   /dev/vda1        99G   99G     0 100% /
   tmpfs           1.6G     0  1.6G   0% /run/user/0
   ```

   In this example, the **/dev/vda1** partition usage is 100%.

2. Check the disk space usage.
   df -i

   Figure 8-5 Checking the disk space usage

   ```
   [root@ecs-a058 /]# df -i
   Filesystem      Inodes   IUsed    IFree IUse% Mounted on
   devtmpfs       2030569     369  2030200    1% /dev
   tmpfs          2033158       2  2033156    1% /dev/shm
   tmpfs          2033158     507  2032651    1% /run
   tmpfs          2033158      16  2033142    1% /sys/fs/cgroup
   /dev/vda1      6553600  139249  6414351    3% /
   tmpfs          2033158       1  2033157    1% /run/user/0
   ```

In this example, the disk space usage is low.

3. Check the deleted process files in the system.

lsof | grep deleted

Figure 8-6 Checking the deleted process files in the system

Roughly calculate the total size of the deleted files based on the returned command output. If it is almost the same as the used space of the disk, the disk space may be used up by the deleted processes that have not been released.

4. Go to the location of a deleted file to check whether the file is still there.

ll */tmp/*

Note that variable */tmp/* in the command indicates the path of the deleted file.

5. If the file is not there, run the following command to terminate the process, or restart the server to release the used space.

kill -9 *PID*

Note that variable *PID* in the command indicates the process ID.

6. Check that the process is terminated.

lsof | grep deleted

7. Check that the disk partition usage is no longer 100%.

df -h

9 Performance

9.1 How Can I Test My Disk Performance?

Precautions

In the disk performance test, if the start sector number is not 4-KiB aligned, the disk performance will be greatly affected. Ensure that the start sector number is 4-KiB aligned before you start the test.

📖 NOTE

To test the performance of a shared disk, the following requirements must be met:

- The shared disk must be attached to multiple servers (ECSs or BMSs).
- If the shared disk is attached to multiple ECSs, these ECSs must belong to the same anti-affinity ECS group.

  If these ECSs fail to meet the anti-affinity requirement, the shared disk cannot reach the optimal performance.

The testing process for Windows and Linux is different.

- **Windows**
- **Linux**

If the test results do not meet your expectation, see **Why Does My Disk Performance Test Using Fio Have Incorrect Results?**

Windows

The way you test disk performance depends on the server OS. This section uses Windows 7 Professional 64-bit as an example. For other Windows OSs, see the corresponding OS documentations.

Install the performance measurement tool Iometer before the test. You can obtain the tool at **http://www.iometer.org/**.

Step 1 Log in to the server.

Step 2 Press **win+R** to open the **Run** window. Enter **msinfo32** and click **OK**.

The system information window is displayed.

Step 3 Choose **components** > **storage** > **disks**. In the right pane, view the partition offset.

- If 4096 can be divided by the parameter value, the partition is 4-KiB aligned. Go to **Step 4**.

- If 4096 cannot be divided by the parameter value, the partition is not 4-KiB aligned. Ensure 4-KiB alignment for the partition before continuing the test.

NOTICE

If you delete the partition and select another start sector number for 4-KiB alignment, you will lose all the data on that partition. Exercise caution when performing this operation.

Step 4 Use Iometer to test the disk performance. For details, see the Iometer product documentation.

When the disk IOPS and throughput are tested, the parameter settings for Iometer and fio are the same. For details, see **Table 9-1**.

The following example uses Iometer to test the disk performance.

1. Set the workflow.
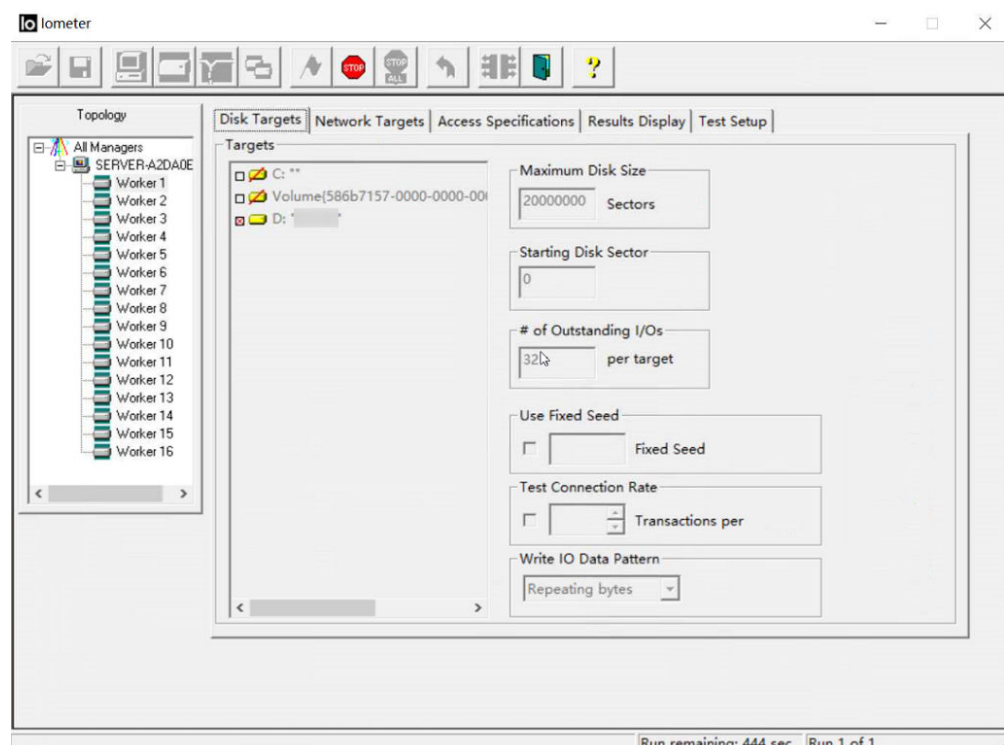
2. Set the test run time.

In this example, the test run time is set to 10 minutes, with 60 seconds ramp up time. Disk performance is tested after the writes are stable.
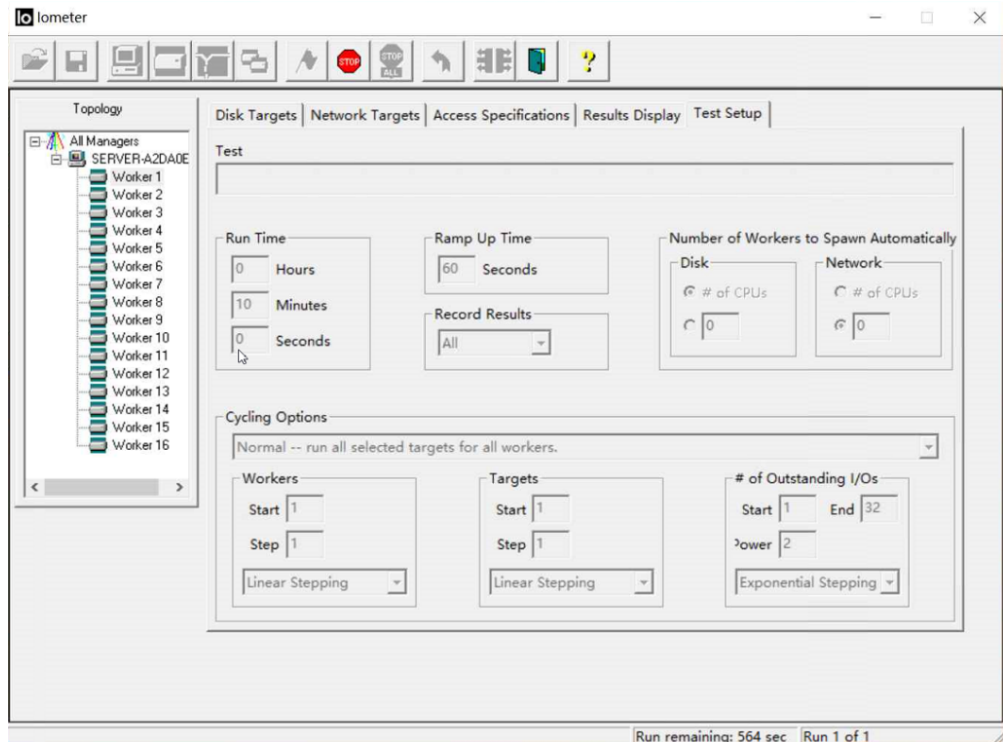
3. Set the data block size and read/write policy. In this example, the disk size is set to 64 KiB, the policy is 100% sequential write.
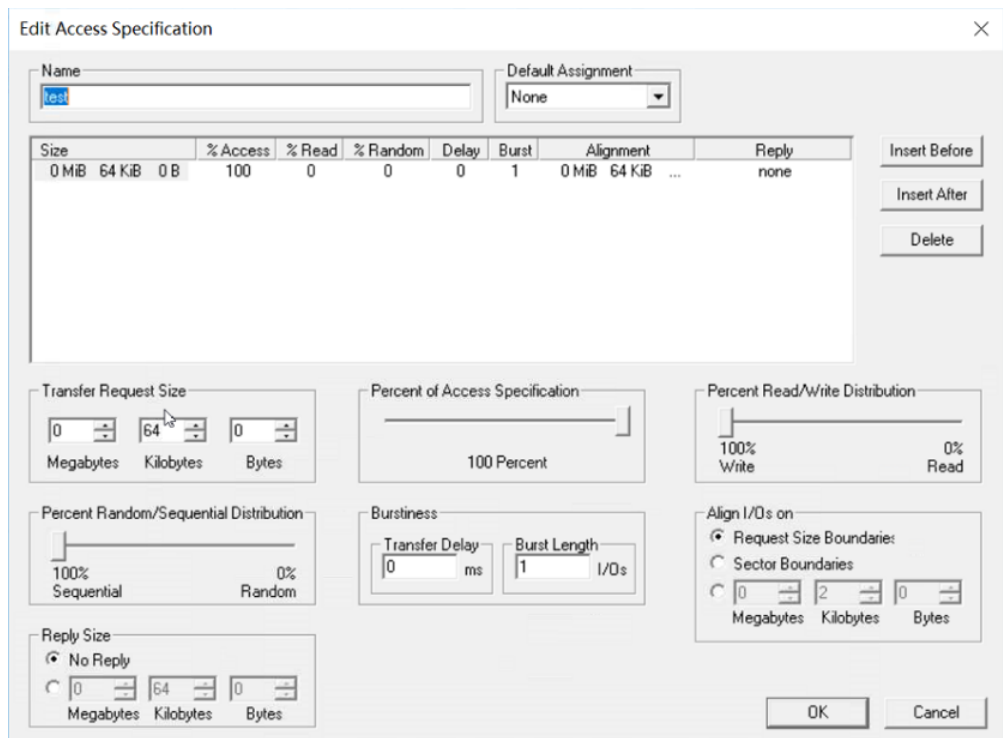
4. View the test results.
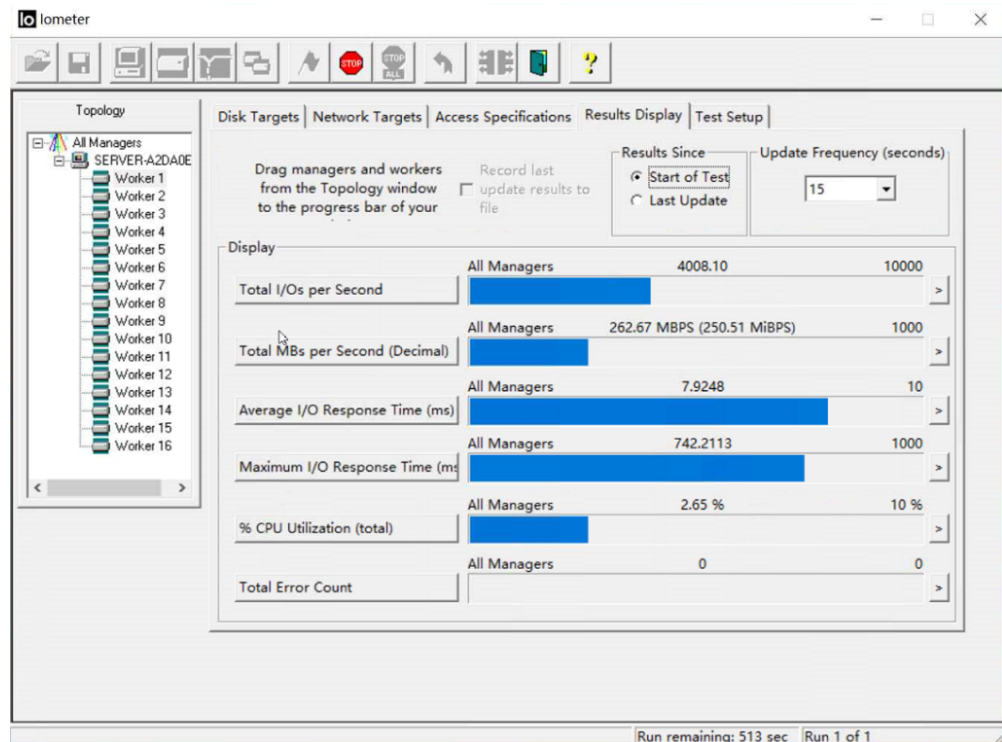
----End

Linux

If you use an old version Linux OS, for example CentOS 6.5, and run **fdisk** to create partitions, the default start sector number will not be 4-KiB aligned, which will greatly affect the test performance. For this reason, if such an OS is used, you are advised to select a new start sector number, one that is 4-KiB aligned, when creating partitions.

The way you test disk performance depends on the server OS. This section uses CentOS 7.2 64-bit as an example. For other Linux OSs, see the corresponding OS documentations.

Install the performance measurement tool, fio, before the test.

Step 1 Log in to the server and switch to user **root**.

Step 2 Before you start the test, run the following command to check whether the start sector number is 4-KiB aligned:

fdisk -lu

Information similar to the following is displayed:
```
[root@ecs-centos sdc]# fdisk -lu

Disk /dev/xvda: 10.7 GiB, 10737418240 bytes, 20971520 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x7db77aa5

  Device Boot      Start         End      Blocks   Id  System
/dev/xvda1   *      2048    20968919    10483436   83  Linux
```

```
Disk /dev/xvdb: 10.7 GiB, 10737418240 bytes, 20971520 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes


Disk /dev/xvdc: 53.7 GiB, 53687091200 bytes, 104857600 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x3cf3265c

  Device Boot      Start        End      Blocks   Id  System
/dev/xvdc1          2048    41943039    20970496   83  Linux
```

- If 8 can be divided by the start sector number, the number is 4-KiB aligned. Go to **Step 3**.

- If 8 cannot be divided by the start sector number, the number is not 4-KiB aligned. Delete the partition and select a 4-KiB aligned start sector number for the new partition before continuing the test.

NOTICE

If you delete the partition and select another start sector number for 4-KiB alignment, you will lose all the data on that partition. Exercise caution when performing this operation.

Step 3 Run the following commands and use fio to test the disk performance:

- To test random write IOPS, run the following command: **fio -direct=**_1_ **-iodepth=**_128_ **-rw=**_randwrite_ **-ioengine=**_libaio_ **-bs=**_4k_ **-size=**_10G_ **-numjobs=**_1_ **-runtime=**_600_ **-group_reporting -filename=**_/opt/fiotest/fiotest.txt_ **-name=**_Rand_Write_IOPS_Test_

- To test random read IOPS, run the following command: **fio -direct=**_1_ **-iodepth=**_128_ **-rw=**_randread_ **-ioengine=**_libaio_ **-bs=**_4k_ **-size=**_10G_ **-numjobs=**_1_ **-runtime=**_600_ **-group_reporting -filename=**_/opt/fiotest/fiotest.txt_ **-name=**_Rand_Read_IOPS_Test_

- To test write throughput, run the following command: **fio -direct=**_1_ **-iodepth=**_32_ **-rw=**_write_ **-ioengine=**_libaio_ **-bs=**_1024k_ **-size=**_10G_ **-numjobs=**_1_ **-runtime=**_600_ **-group_reporting -filename=**_/opt/fiotest/fiotest.txt_ **-name=**_Write_BandWidth_Test_

- To test read throughput, run the following command: **fio -direct=**_1_ **-iodepth=**_32_ **-rw=**_read_ **-ioengine=**_libaio_ **-bs=**_1024k_ **-size=**_10G_ **-numjobs=**_1_ **-runtime=**_600_ **-group_reporting -filename=**_/opt/fiotest/fiotest.txt_ **-name=**_Read_BandWidth_Test_

- To test random read latency, run the following command: **fio -direct=**_1_ **-iodepth=**_1_ **-rw=**_randread_ **-ioengine**=_libaio_ **-bs=**_4k_ **-size=**_10G_ **-numjobs=**_1_ **-runtime=**_60_ **-group_reporting -filename=**_/opt/fiotest/fiotest.txt_ **-name=**_Rand_Read_LATE_Test_

> **NOTICE**
>
> - When using fio to perform a raw disk performance test, ensure that no partitions and file systems have been created on the disk and there is no data stored on the disk. Or, the raw disk test will damage the file system, and data on the disk will become read-only. In this case, your only option will be to delete the disk and buy a new one to continue the test.
>
> - Do not perform the test on a disk with service data on it. If such test is a must, you are advised to perform the test as follows:
>
>   ■ Back up the disk data before the test as you may damage the data on the disk.
>
>   ■ Specify a file, for example **-filename=/opt/fiotest/fiotest.txt**, to test the performance of the file system.

Table 9-1 lists the fio test parameters.

Table 9-1 Parameter description

Parameter	Description
direct	Defines whether direct I/O is used. – Set to **0**: buffered I/O is used. – Set to **1**: direct I/O is used.
iodepth	Defines the I/O queue depth. This queue depth refers to the queue depth of each thread regardless of whether a single or multiple threads are used in the test. Total concurrent I/Os of fio = iodepth x numjobs Examples: – If there is a single thread and **-iodepth=32**, the I/O queue depth of this thread is 32 and the total concurrent I/Os of fio is 32 (32 x 1). – If there are three threads and **-iodepth=32**, the I/O queue depth of each thread is 32 and the total concurrent I/Os of fio is 96 (32 x 3).
rw	Defines the test read/write policy. – **randread**: random read – **randwrite**: random write – **read**: sequential read – **write**: sequential write – **randrw**: mixed random read/write

Parameter	Description
ioengine	Defines how fio delivers the I/O request (synchronously or asynchronously). – Synchronous I/O: Only one I/O request is delivered at a time, and the response is returned after the kernel has processed the request. That said, the single-thread I/O queue depth is always less than 1, and multi-thread concurrent processing can be used to handle such issues. Normally, 16 to 32 concurrent working threads fully occupy the I/O queue depth. – Asynchronous I/O: Multiple I/O requests are delivered using libaio at a time. Wait for the process to complete and reduce the interaction times to improve efficiency.
bs	Defines the I/O block size. The unit can be KiB, Kb, MiB, and Mb, and the default value is 4 KiB.
size	Defines the amount of data processed by the test I/Os. If parameters, such as **runtime**, are not specified, the test ends when fio has processed all the specified data amount. The value can be a number with a unit or percentage. A number with a unit indicates the read/write data amount, for example **size=10G**, indicating a 10-GiB read/write data amount. A percentage indicates the ratio of read/write data amount to the total size of files, for example **size=20%**, indicating the read/write data amount takes 20% of the total file space.
numjobs	Defines the number of concurrent threads.
runtime	Defines the test time. If this parameter is not specified, the test ends until the specified data amount is processed by the block size defined using parameter **size**.
group_reporting	Defines the test result display mode. The parameter value displays the statistics on a single thread instead of that on all jobs.

filename	Defines the name of the test file or device.
	
	– If a file is specified, the performance of the file system is tested. Example: **-filename=/opt/fiotest/fiotest.txt**
	
	– If a device name is specified, the performance of the raw disk is tested. Example: **-filename=/dev/vdb**
	
	  NOTICE
	  If the test is performed on a disk already has partitions and file systems created as well as data on it, user parameter **filename** to specify a file so that the original file system is not damaged and the data is not overwritten.

name	Defines the test task name.

----End

9.2 Why Does My Disk Performance Test Using Fio Have Incorrect Results?

Symptom

You have followed the test performance method, but the test results do not meet expectations.

Troubleshooting

During a disk performance test, the disk and stress test conditions play an important role.

Possible causes are listed here in order of their probability.

If the fault persists after you have ruled out one cause, move on to the next one in the list.
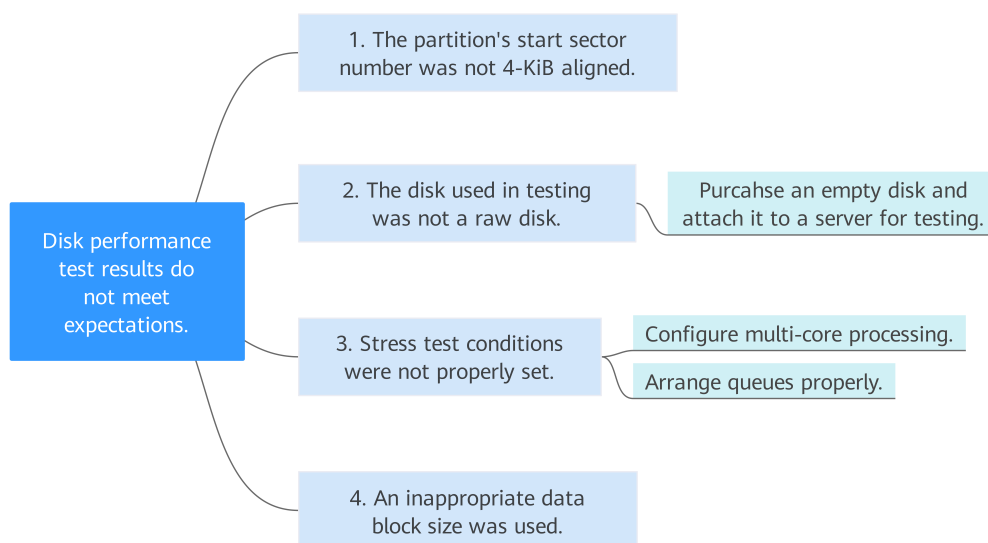
> **NOTICE**
>
> Some operations may result in data loss. It is recommended that you use raw disks for performance test.

Figure 9-1 Troubleshooting

Table 9-2 Troubleshooting

Possible Cause	Solution
The partition's start sector number is not 4-KiB aligned.	Go **Check Whether Partition's Start Sector Number Is 4-KiB Aligned**. Delete the partition and select a 4-KiB aligned start sector number for the new partition.
The disk used in testing was not a raw disk.	Purchase an empty disk and attach it to a server for testing.
Stress test conditions were not properly set.	Configure multi-core processing and arrange queues properly to maximize the concurrent performance.
An inappropriate data block size was used.	Set a suitable data block size. ● When testing the disk IOPS, set the data block size to a small value, for example, 4 KiB. ● When testing the disk throughput, set the data block size to a large value, for example, 1024 KiB.

Check Whether Partition's Start Sector Number Is 4-KiB Aligned

Step 1 Log in to the server and switch to user **root**.

Step 2 Before you start the test, run the following command to check whether the start sector number is 4-KiB aligned:

fdisk -lu

Information similar to the following is displayed:

```
[root@ecs-centos sdc]# fdisk -lu

Disk /dev/xvda: 10.7 GiB, 10737418240 bytes, 20971520 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x7db77aa5

   Device Boot      Start         End      Blocks   Id  System
/dev/xvda1   *        2048    20968919    10483436   83  Linux

Disk /dev/xvdb: 10.7 GiB, 10737418240 bytes, 20971520 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes


Disk /dev/xvdc: 53.7 GiB, 53687091200 bytes, 104857600 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x3cf3265c

   Device Boot      Start         End      Blocks   Id  System
/dev/xvdc1            2048    41943039    20970496   83  Linux
```

- If 8 can be divided by the start sector number, the number is 4-KiB aligned.
- If 8 cannot be divided by the start sector number, the number is not 4-KiB aligned. Delete the partition and select a 4-KiB aligned start sector number for the new partition before continuing the test.

> **NOTICE**
>
> If you delete the partition and select another start sector number for 4-KiB alignment, you will lose all the data on that partition. Exercise caution when performing this operation.

----End

9.3 How Can I Handle a Slowdown in Disk Read/Write Speed and Increased I/O?

Symptom

If you are aware of a service slowdown, depending on if you are examining a Windows or Linux server, you can take the following actions:

- Windows: Open **Task Manager** and view the average response time.
- Linux: Run **iostat -dx** to view the I/O performance.

If the disk read/write speed is slowed down, disk I/O increases, or the await time increases, the disk is likely encounters a performance bottleneck.

Solution

It is recommended that you choose a disk type with higher performance.

If your disk contains important data, create a new disk from the backup, so you do not lose any of the disk data. The procedure is as follows:

1. Create a backup for the disk.

   Figure 9-2 Create Backup
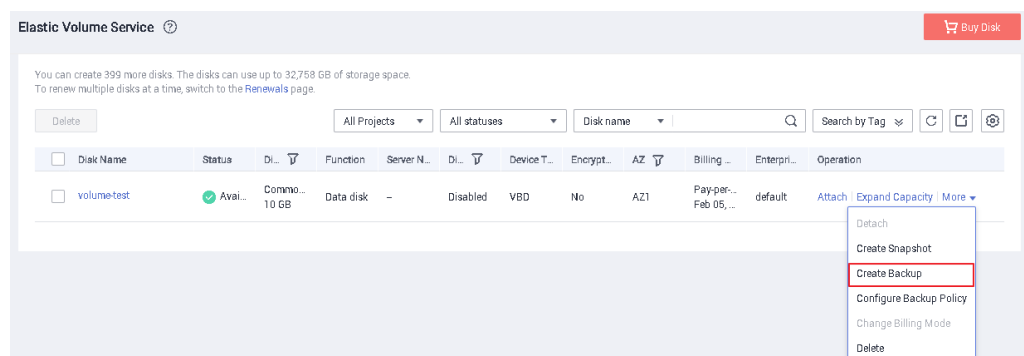
2. Create a new disk from this backup. During the creation, select a new disk type and configure advanced settings (sharing, SCSI, and encryption) based on your service requirements.

Figure 9-3 Create from backup

9.4 How Can I Improve My Disk Performance?

You can use the following methods to improve your disk performance:

- **RAID Array Creation with EVS Disks**
- **Using LVM to Manage EVS Disks**

9.5 Why My Disk's Read IOPS Can't Reach the Theoretical Maximum IOPS When the Disk I/O Usage Is Almost 100%?

Symptom

A 500 GiB ultra-high I/O disk had an I/O usage of 99.94%, but it only had 12,000 IOPS.

Description

- **100% disk I/O usage does not mean that the disk IOPS reaches the maximum.**

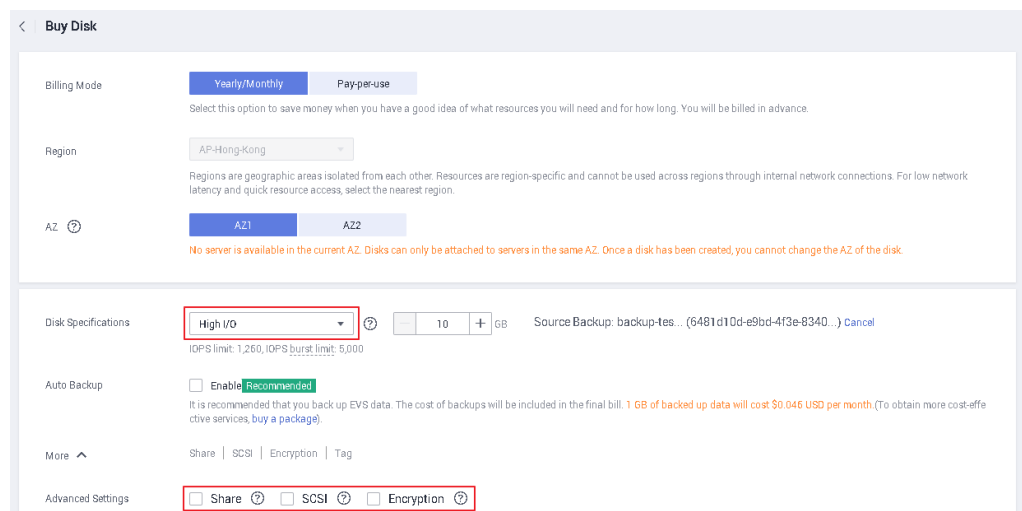 Disk I/O usage calculates the read or write operations performed by a disk in a measurement period. It describes how busy a disk is, not the disk I/O performance.

  EVS disks can process I/O requests concurrently, so 100% disk I/O usage does not mean that the disk encounters the performance bottleneck. For example, an EVS disk takes 0.1 second to process an I/O request and can process 10 I/O requests concurrently. If 10 I/O requests are submitted serially, the disk takes 1 second to process all I/O requests. In this 1-second measurement period, the disk I/O usage reaches 100%. However, if 10 I/O requests are submitted concurrently, the disk takes just 0.1 second to process all the requests. This way, the disk I/O usage in a 1-second measurement period is only 10%. This means that a disk can still process I/O requests even if its I/O usage reaches 100%.

- **Why does the disk not reach the theoretical maximum IOPS?**

  The actual maximum IOPS that a disk can reach is calculated as follows: Disk IOPS = Min. (Max. IOPS, Min. IOPS + IOPS per GiB x Disk size). For a 500 GiB disk, its IOPS is calculated as follows: Disk IOPS = Min. (50,000, 1,800 + 50 × 500) = 26,800

  The disk read IOPS is the number of read operations performed by the disk per second. IOPS is also affected by latency. In a single-queue access scenario with 4 KiB data blocks, the access latency of an ultra-high I/O disk is 1 ms, which means the disk can process 1,000 requests (IOPS) in a second. 12,000 IOPS indicates that the queue depth is 12. To reach the theoretical maximum IOPS (26,800), the queue depth should reach 26.

10 Sharing

10.1 Do I Have to Deploy a Cluster to Use Shared Disks?

Yes.

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

Shared EVS disks do not have cluster management capabilities. You need to build a clustered system for data sharing, such as Windows MSCS, Veritas VCS, and Veritas CFS clusters.

10.2 How Many Servers Can I Attach a Shared Disk to?

A shared disk can be attached to up to 16 servers.

10.3 How Can I Attach a Shared Disk to Multiple Servers?

A shared disk can be attached to multiple servers on the management console. You can choose to attach it to servers one by one or in a batch.

For details, see **Attaching a Shared Disk**.

10.4 Can a Shared Disk Be Attached to Servers That Belong to Different Accounts?

No. A Shared disk can only be attached to servers that belong to the same account and are in the same AZ.

10.5 Can I Attach a Shared Disk to Servers Running Different OSs?

- It is recommended that you do not simultaneously attach a shared disk to servers running Linux and Windows.
- If a shared disk is attached to servers running different versions of the same OS type, it can be used normally. For example, a shared disk attached to one server running CentOS 6 and another server running CentOS 7 can work fine.

11 Snapshot

11.1 What Are the Typical Causes of a Snapshot Creation Failure?

A snapshot creation will fail if its source disk is in an intermediate state, such as **Attaching** and **Expanding**, or an abnormal state, such as **Error** and **Restoration failed**.

Ensure that a disk is in the **In-use** or **Available** state before creating a snapshot.

11.2 Does EVS Support Automatic Snapshot Creation?

No, but you can develop a program to periodically create snapshots through APIs. Alternatively, you may use the CBR service and enable auto backup to create backups periodically. For details, see **Creating a Backup Policy**.

11.3 Can I Create Snapshots for Multiple Disks at a Time?

No.

Snapshots cannot be created for multiple disks at the same time. You must create snapshots for disks one by one.

11.4 How Is a Snapshot Created for My Disk?

There are two types of snapshots: manual snapshots you create on-demand and automatic snapshots created by the system.

- Manual snapshots: You may manually create snapshots to rapidly save the disk data at specific points of time. This feature is now in open beta testing and is free for trial use. For details about how to create snapshots, see **Creating a Snapshot (OBT)**.

- Automatic snapshots: During the creation of a cloud server backup, or a disk backup with the CBR service, the system automatically creates a snapshot and saves the latest snapshot for each disk. If the disk already has a backup and a new backup is created, the system will automatically delete the old snapshot and save the latest one generated. This snapshot is free of charge. You can view the snapshot details only but cannot perform any operations on it.

11.5 Why Can't I Roll Back My Disk Data from a Snapshot?

Possible causes are as follows:

- Snapshot data can be rolled back only when the status of the snapshot's source disk is **Available** or **Rollback failed**. If the snapshot's source disk is in the **In-use** state, detach the disk first and then roll back the snapshot data. After the rollback succeeds, re-attach the disk.
- A snapshot whose name starts with **autobk_snapshot_vbs_**, **manualbk_snapshot_vbs_**, **autobk_snapshot_csbs_**, or **manualbk_snapshot_csbs_** is automatically generated during backup. Such a snapshot can only be viewed. It cannot be used to roll back the disk data.

11.6 Can I Roll Back the Data from a Snapshot After Reinstalling the OS or Formatting the Disk?

- If you have reinstalled or changed the server OS, snapshots of the system disk are automatically deleted. Snapshots of the data disks can be used as usual.
- If you have formatted a disk, snapshots can still be used to roll back the disk data.

11.7 How Is the Snapshot Size Calculated?

When you create a snapshot for a disk for the first time, the snapshot is a full snapshot, and its size is the same as the disk size. Subsequent snapshots created for the same disk are incremental snapshots, so they are only as large as the amount of data that has been added.

11.8 Do Snapshots Take Space on the Disk?

No.

Snapshots are stored on the physical disks that provide storage resources for EVS disks. Therefore, snapshots do not use the EVS disk space.

11.9 Can I Perform Multiple Rollback Operations for a Snapshot?

Yes. You can roll back a snapshot to the source disk multiple times as needed.

11.10 Can I Replicate Snapshots to Other Regions or Accounts?

No. Snapshots cannot be replicated to other regions or accounts.

11.11 Why Can't I Find My Snapshot?

Possible causes are:

- You have manually deleted the snapshot.
- You have deleted the disk based on which the snapshot was created. If a disk is deleted, all the snapshots created for this disk will be automatically deleted.
- You have reinstalled or changed the server OS. System disk snapshots will be automatically deleted after a server OS has been reinstalled or changed.

11.12 Can I Use a Snapshot to Create a Disk and Roll Back the Snapshot at Almost the Same Time?

There are two conditions:

- If you first use the snapshot to create a new disk, you can then roll back the snapshot to the original disk.
- If you first roll back the snapshot to the original disk, you cannot use the snapshot to create a new disk then.

12 Backup

12.1 Do I Need to Stop the Server Before Performing a Backup?

No. You can back up servers that are in use. When a server is running, data is written into disks on the server, and some newly generated data is cached in the server memory. During a backup task, data in the memory will not be automatically written into disks, so the disk data and their backups may be inconsistent.

To ensure data integrity, you are advised to perform the backup during off-peak hours when no data is written to the disks. For applications that require strict consistency, such as databases and email systems, you are advised to enable application-consistent backup.

12.2 Can I Back Up and Restore My EVS Disk to a Different Region?

EVS disks already attached to ECSs support cross-region backup and restoration. Such disks can be backed up together with their ECSs using CBR cloud server backup, which supports cross-region replication. You can replicate server backups to your desired region and use the replicated backups to create images and provision servers.

EVS disks that have not been attached to ECSs do not support cross-region backup and restoration.

12.3 How Do I View My Backup Data?

You can check your backup data in the following ways:

📖 **NOTE**

Backup data cannot be viewed on the CBR console.

Server Backups

1. Create an image from a server backup. For details, see **Using a Backup to Create an Image**.

2. Use the image to create a server. For details, see **Creating an ECS from an Image**.

3. Log in to the server to view the data.

Disk Backups

1. Create a new disk from a disk backup. For details, see **Using a Backup to Create a Disk**.

2. Attach the disk to a server. For details, see **Attaching a Non-Shared Disk** or **Attaching a Shared Disk**.

3. Log in to the server to view the data.

SFS Turbo Backups

1. Create a new SFS Turbo file system from an SFS turbo backup. For details, see **Using a Backup to Create a File System**.

2. Mount the file system to a server.

   – To mount the file system to a Linux server, see **Mounting an NFS File System to ECSs (Linux)**.

   – To mount the file system to a Windows server, see **Mounting an NFS File System to ECSs (Windows)**.

3. Log in to the server to view the data.