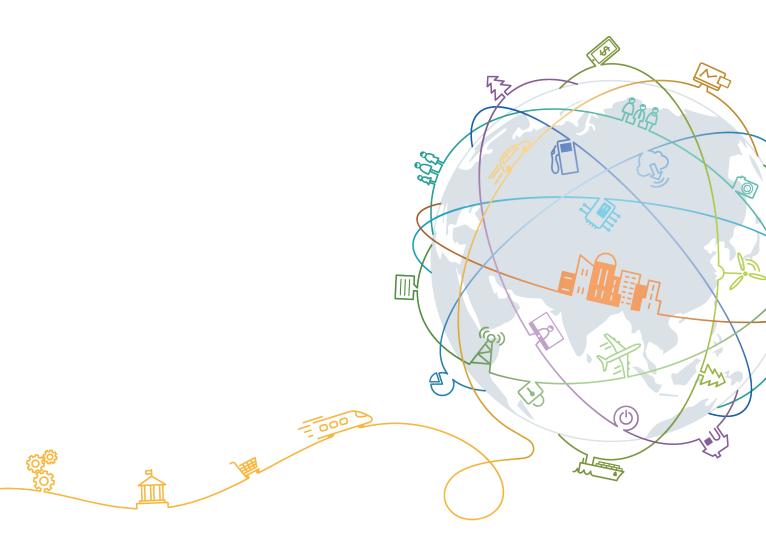
### **Dedicated Distributed Storage Service**

## **Getting Started**

**Issue** 01

**Date** 2018-04-30





#### Copyright © Huawei Technologies Co., Ltd. 2020. All rights reserved.

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

#### **Trademarks and Permissions**

HUAWEI and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd. All other trademarks and trade names mentioned in this document are the property of their respective holders.

#### **Notice**

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

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

## **Contents**

1 Operation Procedure	1
2 Step 1: Apply for a Storage Pool	3
3 Step 2: Create a Disk	5
4 Step 3: Attach a Disk	10
4.1 Attaching a Non-Shared Disk	10
4.2 Attaching a Shared Disk	11
5 Step 4: Initialize a Data Disk	13
5.1 Introduction to Data Disk Initialization Scenarios and Partition Styles	13
5.2 Initializing a Data Disk in Windows (Windows Server 2008)	15
5.3 Initializing a Data Disk in Windows (Windows Server 2016)	
5.4 Initializing a Data Disk in Linux (fdisk)	30
5.5 Initializing a Data Disk in Linux (parted)	36
5.6 Initializing a Data Disk Greater Than 2 TB in Windows (Windows Server 2008)	42
5.7 Initializing a Data Disk Greater Than 2 TB in Windows (Windows Server 2012)	50
5.8 Initializing a Data Disk Greater Than 2 TB in Linux (parted)	58
6 Change History	65

# **1** Operation Procedure

Figure 1-1 shows the basic DSS operation procedure.

Figure 1-1 Basic operation procedure



- 1. Before using DSS, apply for a storage pool first. For more information, see **Step 1: Apply for a Storage Pool**.
- 2. After the storage pool you requested is available, create disks in the storage pool to make use of the storage pool space. For more information, see **Step 2**: **Create a Disk**.
- 3. Attach the disks to servers. For more information, see the following topics:
  - Attaching a Non-Shared Disk
  - Attaching a Shared Disk
- 4. A disk cannot be used right away after being attached to a server. You must log in to the server and initialize the disk. For more information, see the following topics:
  - Introduction to Data Disk Initialization Scenarios and Partition Styles
  - Windows
    - Initializing a Data Disk in Windows (Windows Server 2008)
    - Initializing a Data Disk in Windows (Windows Server 2016)
    - Initializing a Data Disk Greater Than 2 TB in Windows (Windows Server 2008)
    - Initializing a Data Disk Greater Than 2 TB in Windows (Windows Server 2012)
  - Linux

- Initializing a Data Disk in Linux (fdisk)
- Initializing a Data Disk in Linux (parted)
- Initializing a Data Disk Greater Than 2 TB in Linux (parted)

# 2 Step 1: Apply for a Storage Pool

#### **Scenarios**

DSS provides you with dedicated, physical storage resources, which can be flexibly interconnected with various compute services, such as ECS, BMS, and DCC, and are suitable to a wide-range of scenarios, including HPC, OLAP, and a mixed of loads. Before using DSS, you need to apply for a storage pool first.

This topic describes how to apply for a storage pool.

#### **Procedure**

- **Step 1** Log in to the management console.
- **Step 2** Click on the upper left corner and select a region and project.
- **Step 3** Under **Storage**, click **Dedicated Distributed Storage Service** to switch to the **Dedicated Distributed Storage Service** page.
- **Step 4** Click **Apply for Storage Pool** to switch to the page for applying for a storage pool. Configure the basic information about the storage pool by referring to **Table 2-1**.

Table 2-1 Parameter description

Parameter	Description
Region	Specifies the region where the storage pool belongs.
AZ	Specifies the AZ where the storage pool belongs.
Туре	High I/O Ultra-high I/O
Capacity (TB)	Specifies the storage pool capacity you requested.

Parameter	Description
Storage Pool Name	Specifies the storage pool name, which can contain a maximum of 64 characters.
	The name can contain only letters, digits, underscores (_), and hyphens (-).
Quantity	You can choose 1 year to 5 years for the purchased duration.

After the configuration is complete, the details about the storage pool specifications are displayed.

**Step 5** Click **Apply Now**, and the page for you to confirm the specifications is displayed.

You can confirm the details about the storage pool specifications.

- If you do not need to modify the specifications, click **Submit Application** and go to **Step 6**.
- If you need to modify the specifications, click **Previous** to modify parameters.
- **Step 6** Check the application result.
  - On the displayed page, you can click **View Review Status** to view the order details.
  - Or click **Back to Storage Pool List** to return to the storage pool list.
- **Step 7** After the order is approved and paid, view the storage pool status.
  - If the storage pool status is **Deploying**, your order is being processed and the storage resources are being deployed.
  - When the storage pool status changes to **In-use**, the storage pool has been deployed and your order is completed.
- **Step 8** You can perform the following operations for your storage pool in the **Operation** column:
  - Expand the storage pool capacity. For details, see Expanding the Capacity of a Storage Pool.
  - Create disks. For details, see Creating a DSS Disk.

----End

# 3 Step 2: Create a Disk

#### **Scenarios**

DSS storage capabilities are implemented through DSS disks. You need to create disks to make use of the storage resources you requested.

This topic describes how to create a disk in a DSS storage pool.

#### **Precautions**

When disks are created in a storage pool, the type of DSS disks will be the same as that of the storage pool.

#### **Procedure**

- **Step 1** Log in to the management console.
- **Step 2** Choose **Storage > Dedicated Distributed Storage Service** to switch to the **Dedicated Distributed Storage Service** page.
- **Step 3** Choose **Dedicated Distributed Storage Service** > **Disks**.

The disk list page is displayed.

- **Step 4** On the displayed page, click **Create Disk**, or in the storage pool list, locate the storage pool in which the disk will be created and click **Create Disk** in the **Operation** column.
- **Step 5** Configure basic information about the disk by referring to **Table 3-1**.

Table 3-1 Parameter description

Parameter	Description	Example Value
Region	Specifies the region where the tenant belongs.	-

Parameter	Description	Example Value
AZ	Specifies the AZ to which the disk belongs.  NOTE  The disk, storage pool, and the server where the disk is attached must be in the same AZ.  Once a disk has been created, you cannot change the AZ of the disk.	az1.dc1
Storage Pool	Select a storage pool from the drop-down list.	-
Capacity (GB)	The disk capacity must be as follows:  10 GB to 32768 GB  NOTE  If you create a disk from a backup, the disk capacity must be greater than or equal to the backup size. In the condition that you do not specify the disk capacity, if the backup size is smaller than 10 GB, the default capacity 10 GB will be used as the disk capacity; if the backup size is greater than 10 GB, the disk capacity will be consistent with the backup size.  The system shows the remaining disk quotas.	-
Select Data Source	<ul> <li>Optional</li> <li>You can create a new disk or create a disk from a data source.</li> <li>If you do not specify any data source, the disks you created are brand new and contain no data.</li> <li>If you specify a backup, the disk will be created from the backup data. Click Select Data Source and choose Create from backup. On the displayed page, select the target backup and click OK.</li> <li>If you specify an image, the disk will be created from the image data. Click Select Data Source and choose Create from image. On the displayed page, select the target image and click OK.</li> </ul>	_

Parameter	Description	Example Value	
Auto Backup	Optional Auto backup is a function provided by CBR. If auto backup is enabled, the system automatically creates backups for the disk data at specified points in time and deletes outdated backups according to the configured backup policy. When <b>Enable</b> is selected, a backup policy must be configured. You can either use the default backup policy or customize the policy based on your service requirements.	-	
Advanced Settings	<ul> <li>Optional</li> <li>If you select Not required, the advanced disk functions are not used.</li> <li>If you select Configure, the following advanced disk functions will be displayed:         <ul> <li>Share</li> <li>SCSI</li> <li>Encryption</li> </ul> </li> </ul>	-	
Share	<ul> <li>Optional</li> <li>If Share is not selected, a non-shared disk is created.</li> <li>If Share is selected, a shared disk is created and can be attached to multiple servers.</li> </ul>	-	
SCSI	<ul> <li>If SCSI is not selected, a VBD disk is created. VBD is the default device type of disks.</li> <li>If SCSI is not selected, a SCSI disk is created. Such disks allow the server OS to directly access the underlying storage media and send SCSI commands to the disks.</li> <li>If SCSI and Share are both selected, a shared SCSI disk is created.</li> </ul>	-	

		-
encrypt	ncryption is used for data disk tion only. The system disk tion depends on images. For details, to Image Management Service User	
When I followi  Created An a between tenares of tenares	Encryption is selected, the ang parameters will be displayed. The Agency agency is a trust relationship ween two tenants or services. A sent can create an agency to grant curce access rights to another and or service. If the KMS access its are not granted to EVS, the layed. Otherwise, it will not be layed. Otherwise, it will not be layed. A compared to EVS and the KMS access rights be access rights have been atted, EVS can obtain KMS keys to rept or decrypt disks. The KMS access rights have been atted, follow-up operations do not aire the rights to be granted again. We key Name is displayed only after the MS access rights have been granted. For tails, see Create Agency above. We Key Name is the identifier of the and you can use KMS Key Name pecify the KMS key that is to be done for encryption. One of the owing keys can be used:  Default Master Key: After the KMS access rights have been granted to EVS, the system automatically reates a Default Master Key and names it evs/default.  EMKs: Existing or newly created EMKs. For details, see Key Management Service > Creating a EMK in the Data Encryption	

Parameter	Description	Example Value
	NOTE  Before you use the disk encryption function, the KMS access rights need to be granted to EVS. If you have the right granting permission, grant KMS access rights directly. If you do not have this permission, contact a user with the security administrator rights to grant the KMS access rights to EVS, then repeat the preceding operations.	
Disk Name	<ul> <li>If you create one disk, this value is used as the disk name and can contain a maximum of 64 bytes.</li> <li>If you create disks in a batch, this value is used as the prefix of disk names, and one disk name will be composed of this parameter value and a four-digit number. The value can contain a maximum of 64 bytes.</li> <li>NOTE         <ul> <li>The name can contain only letters, digits, underscores (_), and hyphens (-).</li> </ul> </li> </ul>	For example, if you create two disks and set volume for Disk Name, the disk names will be volume-0001 and volume-0002.
Quantity	The default value is set to 1, which means one disk will be created. Currently, you can create up to 100 disks at a time.  NOTE  When you create a disk from a backup, batch creation is not possible, and this parameter value must be set to 1.  When you create a disk from an image, batch creation is not possible, and this parameter value must be set to 1.  The system shows you the maximum number of disks as well as the maximum disk capacity allowed to create. To ensure effective resource usage, if the number of disks you need exceeds the upper limit, click Increase Quota to obtain a higher quota. You can create the disks you need after the request is approved.	1

#### Step 6 Click Next.

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

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

----End

# 4 Step 3: Attach a Disk

### 4.1 Attaching a Non-Shared Disk

#### **Scenarios**

Independently purchased disks are data disks. In the disk list, the function of such disks is displayed as **Data disk**, and the status is displayed as **Available**. In this case, you need to attach the data disks to servers for use.

A system disk must be purchased during a server creation and is automatically attached. In the disk list, the function of such disks is displayed as **System disk**, and the status is displayed as **In-use**. After a system disk is detached from the server, the disk function changes to **Bootable disk**, and the status changes to **Available**.

#### **□** NOTE

Bootable disks are the system disks detached from servers. A bootable disk can be reattached to a server and be used as system disk or data disk depending on the device name selected.

This topic describes how to attach a non-shared DSS disk to a server. A non-shared DSS disk can be attached to one server only.

#### Procedure

- **Step 1** Log in to the management console.
- Step 2 Under Storage, click Dedicated Distributed Storage Service.

The **Dedicated Distributed Storage Service** details page is displayed.

**Step 3** Choose **Dedicated Distributed Storage Service** > **Disks**.

The disk list page is displayed.

**Step 4** Locate the target disk in the list and click **Attach** in the **Operation** column.

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

**Step 5** Select the server you want to attach the disk and choose a device name from the drop-down list. Ensure that the selected server and the target disk are in the same AZ.

One device name can be attached with one disk only. For the mapping between device names displayed on the management console and those on the server, see **What Is the Mapping Between Device Names and Disks?** in the *Elastic Cloud Server User Guide*.

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

**Step 6** Initialize the disk.

After the disk has been attached to a server, the disk can be used only after you have initialized it. For details, see **Step 4: Initialize a Data Disk**.

----End

### 4.2 Attaching a Shared Disk

#### **Scenarios**

Independently purchased disks are data disks. In the disk list, the function of such disks is displayed as **Data disk**, and the status is displayed as **Available**. In this case, you need to attach the data disks to servers for use.

This topic describes how to attach a shared DSS disk to servers. A shared DSS disk can be attached to a maximum of 16 servers.

#### Procedure

- **Step 1** Log in to the management console.
- **Step 2** Choose **Storage** > **Dedicated Distributed Storage Service**.

The storage pool list is displayed.

Step 3 Choose Dedicated Distributed Storage Service > Disks.

The disk list page is displayed.

**Step 4** Locate the target disk in the list and click **Attach** in the **Operation** column.

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

**Step 5** Select servers and select a device name from the drop-down list for each server you selected. Ensure that the disk and servers are in the same AZ.

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

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

#### NOTICE

If you simply attach a shared disk to multiple servers, files cannot be shared between the servers as shared DSS disks do not have the cluster capability. Therefore, build a shared file system or deploy a cluster management system if you need to share files between servers.

----End

# 5 Step 4: Initialize a Data Disk

# 5.1 Introduction to Data Disk Initialization Scenarios and Partition Styles

#### **Scenarios**

After a disk is attached to a server, you need to log in to the server to initialize the disk, that is, format the disk. You must initialize a disk before accessing it.

System disk

A system disk does not require manual initialization because it is automatically created and initialized upon server creation. The default disk partition style is master boot record (MBR).

- Data disk
  - If a data disk is created along with a server, it will be automatically attached to the server.
  - If a data disk is created separately, you need to manually attach it to a server.

In both cases, you must initialize the data disk before using it. Choose a proper disk partition style based on your service plan.

#### **Constraints**

Disks created from backups, snapshots, or images do not need to be initialized. Otherwise, data on the disks may be lost.

#### **Disk Partition Styles**

**Table 5-1** lists the common disk partition styles. In Linux, different disk partition styles require different partitioning tools.

**Table 5-1** Disk partition styles

Disk Partition Style	Maximum Disk Capacity Supported	Maximum Number of Partitions Supported	Linux Partitioning Tool
Master Boot Record (MBR)	2 TB	<ul> <li>4 primary partitions</li> <li>3 primary partitions and 1 extended partition</li> <li>With MBR, one may create several primary partitions and an extended partition. An extended partition must be divided into several logical partitions before use. For example, if 6 partitions need to be created, you can create the partitions in the following two ways:</li> <li>3 primary partitions and 1 extended partition, with the extended partition divided into 3 logical partitions</li> <li>1 primary partition and 1 extended partition, with the extended partition divided</li> </ul>	• fdisk • parted
Guid Partition Table (GPT)	18 EB 1 EB = 1048576 TB	into 5 logical partitions  Unlimited  Disk partitions created using  GPT are not categorized.	parted

#### **NOTICE**

The maximum disk capacity supported by MBR is 2 TB, and that supported by GPT is 18 EB. Because a data disk currently supports up to 32 TB, use the GPT partition style if your disk capacity is larger than 2 TB.

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

### **Partitioning Operation Guide**

For a disk smaller than 2 TB, see the following topics:

- Initializing a Data Disk in Windows (Windows Server 2008)
- Initializing a Data Disk in Windows (Windows Server 2016)

- Initializing a Data Disk in Linux (fdisk)
- Initializing a Data Disk in Linux (parted)

For a disk larger than 2 TB, see the following topics:

- Initializing a Data Disk Greater Than 2 TB in Windows (Windows Server 2008)
- Initializing a Data Disk Greater Than 2 TB in Windows (Windows Server 2012)
- Initializing a Data Disk Greater Than 2 TB in Linux (parted)

## 5.2 Initializing a Data Disk in Windows (Windows Server 2008)

#### **Scenarios**

This topic uses Windows Server 2008 R2 Enterprise 64bit to describe how to initialize a data disk attached to a server running Windows.

The maximum disk capacity supported by MBR is 2 TB, and that supported by GPT is 18 EB. Therefore, use the GPT partition style if your disk capacity is larger than 2 TB. For details, see Initializing a Data Disk Greater Than 2 TB in Windows (Windows Server 2008). For details about disk partition styles, see Introduction to Data Disk Initialization Scenarios and Partition Styles.

The method for initializing a disk varies depending on the OS running on the server. This document is used for reference only. For the detailed operations and differences, see the product documents of the corresponding OS.

#### **NOTICE**

When using a disk for the first time, if you have not initialized it, including creating partitions and file systems, the additional space added to this disk in an expansion later may not be normally used.

#### **Prerequisites**

- A data disk has been attached to a server and has not been initialized.
- You have logged in to the server.
  - For how to log in to an ECS, see the *Elastic Cloud Server User Guide*.
  - For how to log in to a BMS, see the Bare Metal Server User Guide.

#### Procedure

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

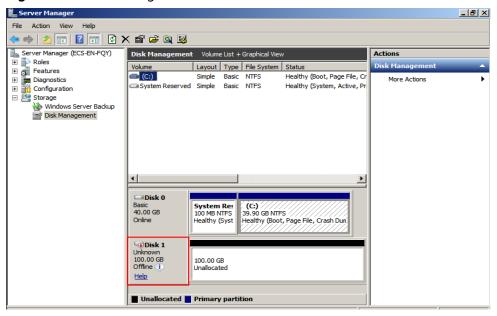
The **Server Manager** window is displayed.

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

The **Disk Management** window is displayed.

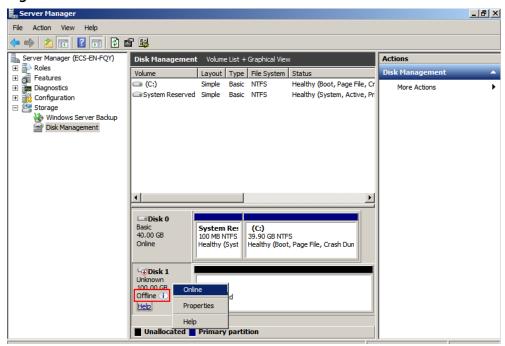
- If Figure 5-1 is displayed, the new disk is offline. Go to Step 3.
- If Figure 5-4 is displayed, the Initialize Disk window is prompted. Go to Step 5.

Figure 5-1 Disk Management



**Step 3** Disks are displayed in the right pane. In the **Disk 1** area, right-click **Offline** and choose **Online** from the shortcut menu to online the disk.



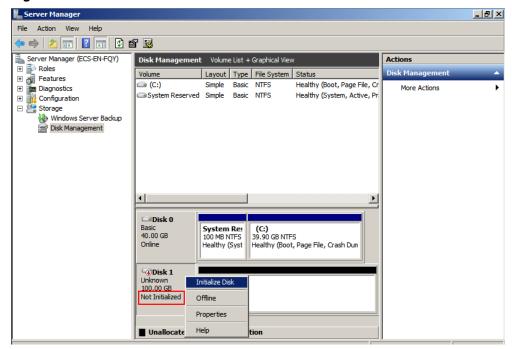


#### **□** NOTE

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

**Step 4** After making the disk online, the disk status changes from **Offline** to **Not Initialized**. Right-click the disk status and choose **Initialize Disk** from the shortcut menu, as shown in **Figure 5-3**.

Figure 5-3 Initialize Disk



Step 5 In the Initialize Disk dialog box, select the target disk, click MBR (Master Boot Record) or GPT (GUID Partition Table), and click OK, as shown in Figure 5-4.

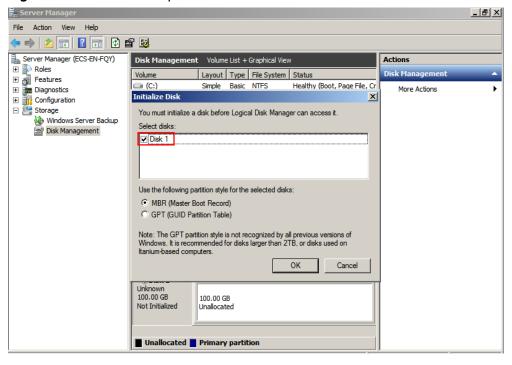


Figure 5-4 Unallocated space

#### **NOTICE**

The maximum disk capacity supported by MBR is 2 TB, and that supported by GPT is 18 EB. Because a data disk currently supports up to 32 TB, use the GPT partition style if your disk capacity is larger than 2 TB.

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

**Step 6** Right-click at the unallocated space and choose **New Simple Volume** from the shortcut menu, as shown in **Figure 5-5**.

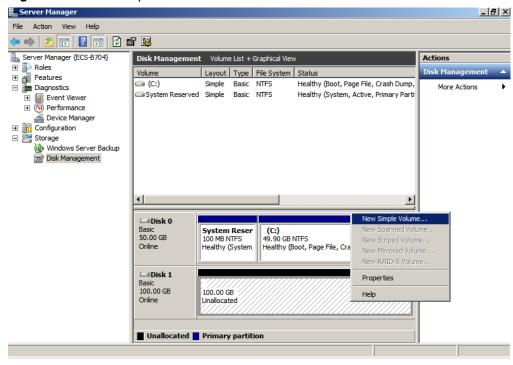


Figure 5-5 New Simple Volume

Step 7 On the displayed New Simple Volume Wizard window, click Next.





**Step 8** Specify the volume size and click **Next**. The default value is the maximum size.

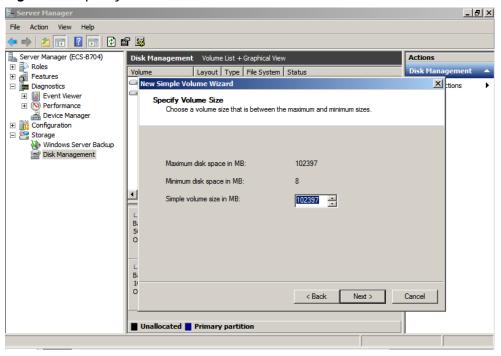
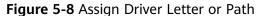
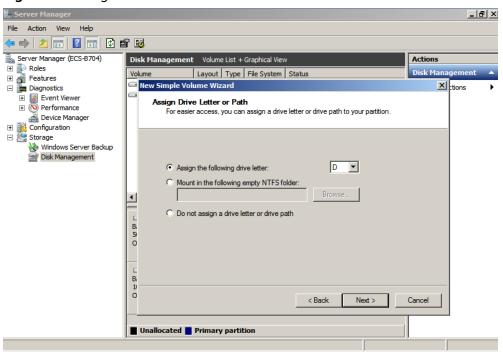


Figure 5-7 Specify Volume Size

Step 9 Assign the driver letter and click Next.





Step 10 Select Format this volume with the following settings, set parameters based on the actual requirements, and select Perform a quick format. Then, click Next.

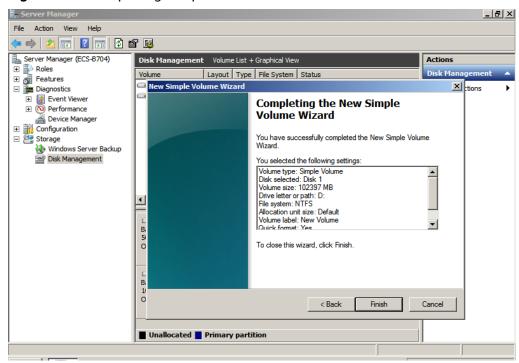
File Action View Help (= 📦 | 🚈 📊 | 🛭 ந | 🗗 🗳 Server Manager (ECS-B704)

Roles Disk Management Volume List + Graphical View Actions Layout Type File System Status Volume tions New Simple Volume Wizard Event Viewer
 Performance To store data on this partition, you must format it first. Device Manager

Configuration ☐ Market Storage Choose whether you want to format this volume, and if so, what settings you want to use. Windows Server Backup Disk Management O Do not format this volume • Format this volume with the following settings: File system: NTFS ▼ Default ▾ Allocation unit size: New Volume Perform a quick format Enable file and folder compression < Back Next > Cancel ■ Unallocated ■ Primary partition

Figure 5-9 Format Partition

Figure 5-10 Completing the partition creation



#### NOTICE

The partition sizes supported by file systems vary. Therefore, you are advised to choose an appropriate file system based on your service requirements.

**Step 11** Click **Finish**. Wait for the initialization to complete. When the volume status changes to **Healthy**, the initialization has finished successfully, as shown in **Figure 5-11**.

🛴 Server Manager \_ B × File Action View Help 🚣 Server Manager (ECS-B704) Actions Volume List + Graphical Viev Layout Type File System Status Volume (C:) Simple Basic NTFS Healthy (Boot, Page File, Crash Dump More Actions New Volume (D:) Simple Basic NTFS Healthy (Primary Partition) System Reserved Simple Basic NTFS Healthy (System, Active, Primary Parti ☐ Storage 🐌 Windows Server Backup Disk Management □Disk 0 System Reser 100 MB NTFS (C:) 49.90 GB NTFS Healthy (Boot, Page File, Crash Dump, Prima 50.00 GB Healthy (System □Disk 1 New Volume (D:) 100.00 GB NTFS 100.00 GB Healthy (Primary Partition) ■ Unallocated ■ Primary partition

Figure 5-11 Disk initialization succeeded

----End

# 5.3 Initializing a Data Disk in Windows (Windows Server 2016)

#### **Scenarios**

This topic uses Windows Server 2016 Standard 64bit to describe how to initialize a data disk attached to a server running Windows.

The method for initializing a disk varies depending on the OS running on the server. This document is used for reference only. For the detailed operations and differences, see the product documents of the corresponding OS.

#### NOTICE

When using a disk for the first time, if you have not initialized it, including creating partitions and file systems, the additional space added to this disk in an expansion later may not be normally used.

#### **Prerequisites**

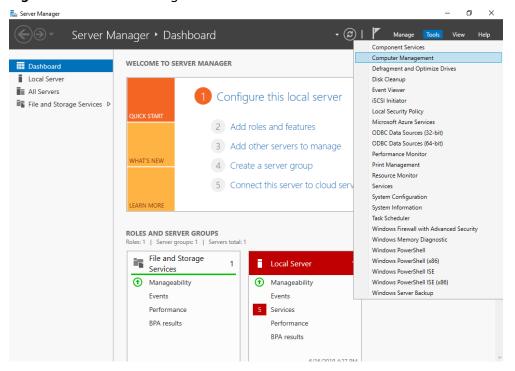
- A data disk has been attached to a server and has not been initialized.
- You have logged in to the server.
  - For how to log in to an ECS, see the Elastic Cloud Server User Guide.
  - For how to log in to a BMS, see the Bare Metal Server User Guide.

#### **Procedure**

- **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 5-12 Server Manager



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

The Computer Management window is displayed.

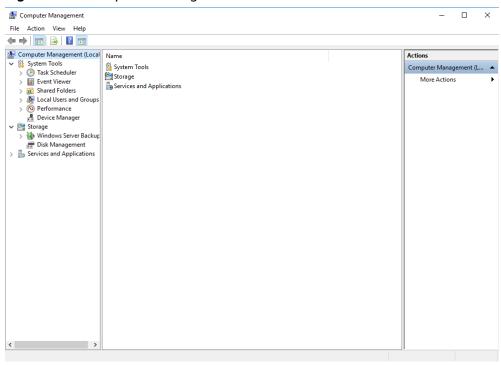


Figure 5-13 Computer Management

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

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

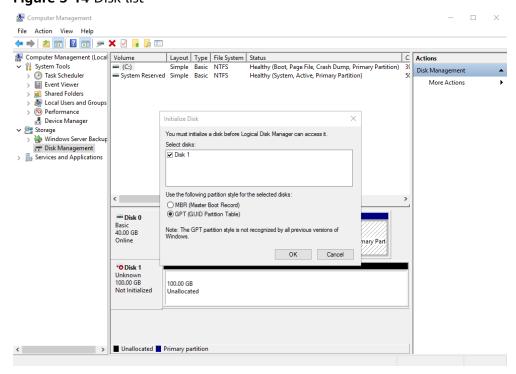
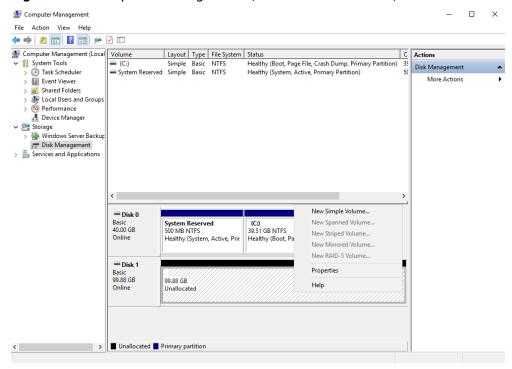


Figure 5-14 Disk list

**Step 5** In the **Initialize Disk** dialog box, the to-be-initialized disk is selected. Select a disk partition style and click **OK**. In this example, **GPT (GUID Partition Table)** is selected.

The **Computer Management** window is displayed.

Figure 5-15 Computer Management (Windows Server 2016)



#### **NOTICE**

The maximum disk capacity supported by MBR is 2 TB, and that supported by GPT is 18 EB. Because a data disk currently supports up to 32 TB, use the GPT partition style if your disk capacity is larger than 2 TB.

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

**Step 6** Right-click at the unallocated disk space and choose **New Simple Volume** from the shortcut menu.

The New Simple Volume Wizard window is displayed.

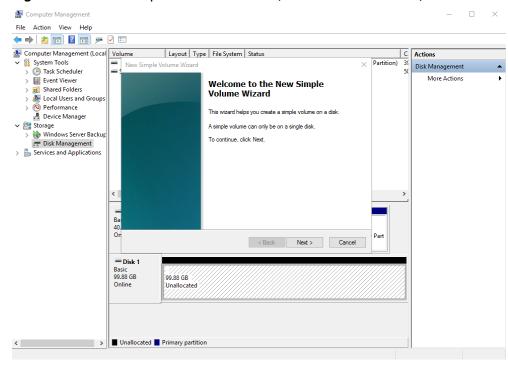
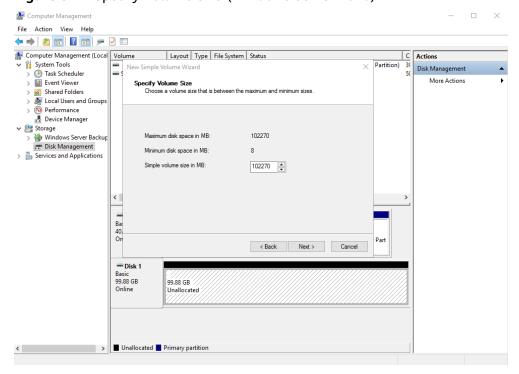


Figure 5-16 New Simple Volume Wizard (Windows Server 2016)

**Step 7** Follow the prompts and click **Next**.

The **Specify Volume Size** page is displayed.

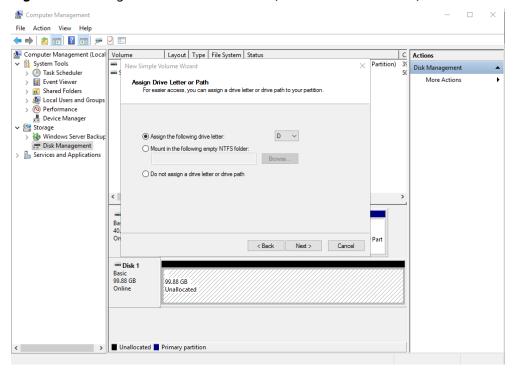
Figure 5-17 Specify Volume Size (Windows Server 2016)



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

The **Assign Drive Letter or Path** page is displayed.

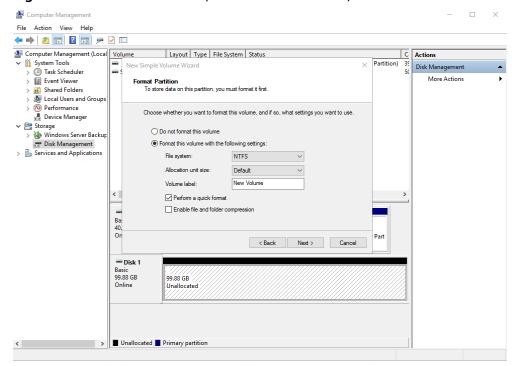
Figure 5-18 Assign Driver Letter or Path (Windows Server 2016)



**Step 9** Assign a drive letter or path to your partition and click **Next**. The system assigns drive letter D by default. In this example, the default setting is used.

The Format Partition page is displayed.

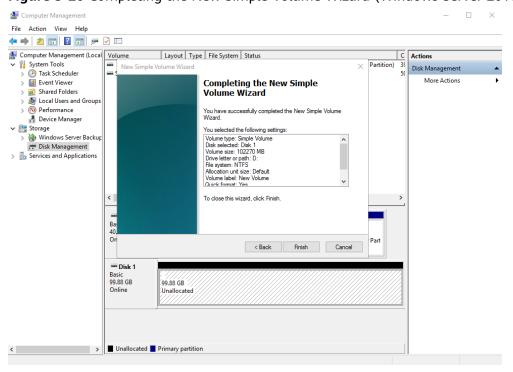
Figure 5-19 Format Partition (Windows Server 2016)



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

The Completing the New Simple Volume Wizard page is displayed.

Figure 5-20 Completing the New Simple Volume Wizard (Windows Server 2016)



#### NOTICE

The partition sizes supported by file systems vary. Therefore, you are advised to choose an appropriate file system based on your service requirements.

#### Step 11 Click Finish.

Wait for the initialization to complete. When the volume status changes to **Healthy**, the initialization has finished successfully, as shown in **Figure 5-21**.

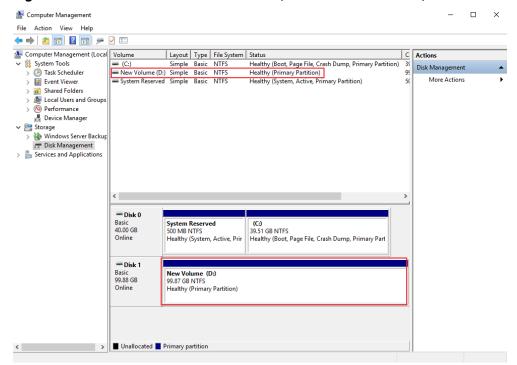


Figure 5-21 Disk initialization succeeded (Windows Server 2016)

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

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

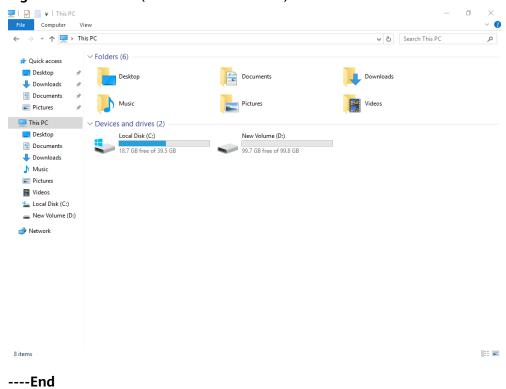


Figure 5-22 This PC (Windows Server 2016)

### 5.4 Initializing a Data Disk in Linux (fdisk)

#### **Scenarios**

This topic uses CentOS 7.4 64bit to describe how to initialize a data disk attached to a server running Linux and use fdisk to partition the data disk.

The maximum disk capacity supported by MBR is 2 TB, and that supported by GPT is 18 EB. Therefore, use the GPT partition style if your disk capacity is larger than 2 TB. In Linux, if you choose to use the GPT partition style, the fdisk partitioning tool cannot be used. Use the parted partitioning tool instead. For details about disk partition styles, see Introduction to Data Disk Initialization Scenarios and Partition Styles.

The method for initializing a disk varies depending on the OS running on the server. This document is used for reference only. For the detailed operations and differences, see the product documents of the corresponding OS.

#### **NOTICE**

When using a disk for the first time, if you have not initialized it, including creating partitions and file systems, the additional space added to this disk in an expansion later may not be normally used.

#### **Prerequisites**

- A data disk has been attached to a server and has not been initialized.
- You have logged in to the server.
  - For how to log in to an ECS, see the Elastic Cloud Server User Guide.
  - For how to log in to a BMS, see the *Bare Metal Server User Guide*.

#### **Creating and Mounting a Partition**

The following example shows you how a new primary partition can be created on a new data disk that has been attached to a server. The primary partition will be created using fdisk, and MBR is the default partition style. Furthermore, the partition will be formatted using the ext4 file system, mounted on /mnt/sdc, and configured automatic mounting at system start.

**Step 1** Run the following command to query information about the new data disk:

#### fdisk -l

Information similar to the following is displayed:

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

Disk /dev/vda: 42.9 GB, 42949672960 bytes, 83886080 sectors

Units = sectors of 1 \* 512 = 512 bytes

Sector size (logical/physical): 512 bytes / 512 bytes I/O size (minimum/optimal): 512 bytes / 512 bytes

Disk label type: dos Disk identifier: 0x000bcb4e

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

Disk /dev/vdb: 107.4 GB, 107374182400 bytes, 209715200 sectors

Units = sectors of 1 \* 512 = 512 bytes

Sector size (logical/physical): 512 bytes / 512 bytes I/O size (minimum/optimal): 512 bytes / 512 bytes

In the command output, the server contains two disks. /dev/vda is the system disk, and /dev/vdb is the new data disk.

Step 2 Run the following command to enter fdisk to partition the new data disk:

fdisk New data disk

In this example, run the following command:

#### fdisk /dev/vdb

Information similar to the following is displayed:

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

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

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

Command (m for help):

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

Information similar to the following is displayed:

```
Command (m for help): n
Partition type:
p primary (0 primary, 0 extended, 4 free)
e extended
```

There are two types of disk partitions:

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

#### **□** NOTE

If the MBR partition style is used, a maximum of 4 primary partitions, or 3 primary partitions and 1 extended partition can be created. The extended partition cannot be used directly and must be divided into logical partitions before use.

Disk partitions created using GPT are not categorized.

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

Information similar to the following is displayed:

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

**Partition number** indicates the serial number of the primary partition. The value ranges from 1 to 4.

**Step 5** Enter the serial number of the primary partition and press **Enter**. Primary partition number **1** is used in this example. One usually starts with partition number **1** when partitioning an empty disk.

Information similar to the following is displayed:

```
Partition number (1-4, default 1): 1
First sector (2048-209715199, default 2048):
```

**First sector** indicates the start sector. The value ranges from **2048** to **209715199**, and the default value is **2048**.

**Step 6** Select the default start sector **2048** and press **Enter**.

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

Information similar to the following is displayed:

```
First sector (2048-209715199, default 2048):
Using default value 2048
Last sector, +sectors or +size{K,M,G} (2048-209715199, default 209715199):
```

**Last sector** indicates the end sector. The value ranges from **2048** to **209715199**, and the default value is **209715199**.

**Step 7** Select the default end sector **209715199** and press **Enter**.

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

Information similar to the following is displayed:

```
Last sector, +sectors or +size{K,M,G} (2048-209715199, default 209715199):
Using default value 209715199
Partition 1 of type Linux and of size 100 GiB is set
```

Command (m for help):

A primary partition has been created for the new data disk.

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

Information similar to the following is displayed:

Command (m for help): p

Disk /dev/vdb: 107.4 GB, 107374182400 bytes, 209715200 sectors

Units = sectors of 1 \* 512 = 512 bytes

Sector size (logical/physical): 512 bytes / 512 bytes I/O size (minimum/optimal): 512 bytes / 512 bytes

Disk label type: dos Disk identifier: 0x38717fc1

Device Boot Start End Blocks Id System /dev/vdb1 2048 209715199 104856576 83 Linux

Command (m for help):

Details about the /dev/vdb1 partition are displayed.

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

Information similar to the following is displayed:

Command (m for help): w

The partition table has been altered!

Calling ioctl() to re-read partition table.

Syncing disks.

The partition is created.

□ NOTE

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

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

#### partprobe

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

#### mkfs -t File system format /dev/vdb1

In this example, run the following command to set the **ext4** file system for the new partition:

#### mkfs -t ext4 /dev/vdb1

Information similar to the following is displayed:

[root@ecs-test-0001 ~]# mkfs -t ext4 /dev/vdb1

mke2fs 1.42.9 (28-Dec-2013)

Filesystem label=

OS type: Linux

Block size=4096 (log=2)

Fragment size=4096 (log=2)

Stride=0 blocks, Stripe width=0 blocks

6553600 inodes, 26214144 blocks

1310707 blocks (5.00%) reserved for the super user

First data block=0

Maximum filesystem blocks=2174746624

800 block groups

32768 blocks per group, 32768 fragments per group

8192 inodes per group

The formatting takes a period of time. Observe the system running status and do not exit.

#### **NOTICE**

The partition sizes supported by file systems vary. Therefore, you are advised to choose an appropriate file system based on your service requirements.

**Step 12** Run the following command to create a mount point:

mkdir Mount point

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

mkdir /mnt/sdc

**Step 13** Run the following command to mount the new partition on the created mount point:

mount Disk partition Mount point

In this example, run the following command to mount the new partition /dev/vdb1 on /mnt/sdc:

mount /dev/vdb1 /mnt/sdc

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

#### df -TH

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# df -TH
Filesystem
           Type
                   Size Used Avail Use% Mounted on
                   43G 1.9G 39G 5%/
/dev/vda1
            ext4
           devtmpfs 2.0G 0 2.0G 0% /dev
devtmpfs
          tmpfs 2.0G 0 2.0G 0% /dev/shm
tmpfs 2.0G 9.1M 2.0G 1% /run
tmpfs
tmpfs
          tmpfs
                  2.0G 0 2.0G 0% /sys/fs/cgroup
tmnfs
          tmpfs
                  398M 0 398M 0% /run/user/0
tmpfs
/dev/vdb1 ext4 106G 63M 101G 1% /mnt/sdc
```

New partition /dev/vdb1 is mounted on /mnt/sdc.

# **□** NOTE

If the server is restarted, the mounting will become invalid. You can set automatic mounting for partitions at system start by modifying the /etc/fstab file. For details, see Setting Automatic Mounting at System Start.

----End

# **Setting Automatic Mounting at System Start**

To automatically mount disk partitions at system start, do not specify partitions, for example /dev/vdb1, in /etc/fstab because the sequence of cloud devices, and therefore their names may change during the server stop and start. You are advised to use the universally unique identifier (UUID) in /etc/fstab to set automatic mounting at system start.

#### **□** NOTE

UUID is the unique character string for disk partitions in a Linux system.

**Step 1** 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.

**Step 2** Run the following command to open the **fstab** file using the vi editor:

# vi /etc/fstab

- **Step 3** Press **i** to enter the editing mode.
- **Step 4** Move the cursor to the end of the file and press **Enter**. Then, add the following information:

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

The preceding 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 **Step 1**.
- The second column indicates the directory on which the partition is mounted. You can guery 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**: not use Linux dump backup. Normally, dump backup is not used, and you can set this parameter to **0**.
  - 1: use Linux dump backup.
- 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** so that the system checks the partitions in the ascending order of the values.

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

The system saves the configurations and exits the vi editor.

- **Step 6** Perform the following operations to verify the automatic mounting function:
  - 1. Run the following command to unmount the partition:

umount Disk partition

In this example, run the following command:

umount /dev/vdb1

- 2. Run the following command to reload all the content in the /etc/fstab file: mount -a
- 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, the automatic mounting function takes effect:

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

----End

# 5.5 Initializing a Data Disk in Linux (parted)

# **Scenarios**

This topic uses CentOS 7.4 64bit to describe how to initialize a data disk attached to a server running Linux and use parted to partition the data disk.

The maximum disk capacity supported by MBR is 2 TB, and that supported by GPT is 18 EB. Therefore, use the GPT partition style if your disk capacity is larger than 2 TB. In Linux, if you choose to use the GPT partition style, the fdisk partitioning tool cannot be used. Use the parted partitioning tool instead. For details about disk partition styles, see Introduction to Data Disk Initialization Scenarios and Partition Styles.

The method for initializing a disk varies depending on the OS running on the server. This document is used for reference only. For the detailed operations and differences, see the product documents of the corresponding OS.

# **NOTICE**

When using a disk for the first time, if you have not initialized it, including creating partitions and file systems, the additional space added to this disk in an expansion later may not be normally used.

# **Prerequisites**

- A data disk has been attached to a server and has not been initialized.
- You have logged in to the server.
  - For how to log in to an ECS, see the *Elastic Cloud Server User Guide*.
  - For how to log in to a BMS, see the *Bare Metal Server User Guide*.

# **Creating and Mounting a Partition**

The following example shows you how a new partition can be created on a new data disk that has been attached to a server. The partition will be created using parted, and GPT is used as the partition style. Furthermore, the partition will be formatted using the ext4 file system, mounted on /mnt/sdc, and configured automatic mounting at system start.

**Step 1** Run the following command to query information about the new data disk:

# 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 100G 0 disk
```

In the command output, the server contains two disks. **/dev/vda** is the system disk, and **/dev/vdb** is the new data disk.

**Step 2** Run the following command to enter parted to partition the new data disk:

```
parted New data disk
```

In this example, run the following command:

#### parted /dev/vdb

Information similar to the following is displayed:

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

**Step 3** Enter **p** and press **Enter** to view the current disk partition style.

# Information similar to the following is displayed:

```
(parted) p
Error: /dev/vdb: unrecognised disk label
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 107GB
Sector size (logical/physical): 512B/512B
Partition Table: unknown
Disk Flags:
(parted)
```

In the command output, the **Partition Table** value is **unknown**, indicating that no partition style is set for the new disk.

**Step 4** Run the following command to set the disk partition style:

mklabel Disk partition style

In this example, run the following command to set the partition style to GPT: (Disk partition styles can be MBR or GPT.)

# mklabel gpt

#### **NOTICE**

The maximum disk capacity supported by MBR is 2 TB, and that supported by GPT is 18 EB. Because a data disk currently supports up to 32 TB, use the GPT partition style if your disk capacity is larger than 2 TB.

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

**Step 5** Enter **p** and press **Enter** to view the disk partition style.

Information similar to the following is displayed:

(parted) mklabel gpt

(parted) p

Model: Virtio Block Device (virtblk)

Disk /dev/vdb: 107GB

Sector size (logical/physical): 512B/512B

Partition Table: gpt

Disk Flags:

Number Start End Size File system Name Flags

(parted)

In the command output, the **Partition Table** value is **gpt**, indicating that the disk partition style is GPT.

- **Step 6** Enter **unit s** and press **Enter** to set the measurement unit of the disk to sector.
- **Step 7** Run the following command and press **Enter**:

mkpart Partition name Start sector End sector

In this example, run the following command:

# mkpart test 2048s 100%

In this example, one partition is created for the new data disk. Variable *2048s* indicates the disk start sector, and variable *100%* indicates the disk end sector. The two values are used for reference only. You can determine the number of partitions and the partition size based on your service requirements.

Information similar to the following is displayed:

(parted) mkpart opt 2048s 100%

(parted)

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

Information similar to the following is displayed:

(parted) p

Model: Virtio Block Device (virtblk)

Disk /dev/vdb: 209715200s

Sector size (logical/physical): 512B/512B

Partition Table: gpt

Disk Flags:

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

(parted)
```

# **Step 9** Enter **q** and press **Enter** to exit parted.

Information similar to the following is displayed:

(parted) q Information: You may need to update /etc/fstab.

You can set automatic disk mounting by updating the /etc/fstab file. Before updating the file, set the file system format for the partition and mount the partition on the mount point.

**Step 10** Run the following command to view the disk partition information:

#### lsblk

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vda 253:0 0 40G 0 disk
__vda1 253:1 0 40G 0 part /
vdb 253:16 0 100G 0 disk
__vdb1 253:17 0 100G 0 part
```

In the command output, /dev/vdb1 is the partition you created.

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

# mkfs -t File system format /dev/vdb1

In this example, run the following command to set the **ext4** file system for the new partition:

# mkfs -t ext4 /dev/vdb1

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# mkfs -t ext4 /dev/vdb1
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
6553600 inodes, 26213888 blocks
1310694 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2174746624
800 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
     32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
     4096000, 7962624, 11239424, 20480000, 23887872
Allocating group tables: done
Writing inode tables: done
```

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

The formatting takes a period of time. Observe the system running status and do not exit.

#### **NOTICE**

The partition sizes supported by file systems vary. Therefore, you are advised to choose an appropriate file system based on your service requirements.

**Step 12** Run the following command to create a mount point:

mkdir Mount point

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

mkdir /mnt/sdc

**Step 13** Run the following command to mount the new partition on the created mount point:

mount Disk partition Mount point

In this example, run the following command to mount the new partition /dev/vdb1 on /mnt/sdc:

mount /dev/vdb1 /mnt/sdc

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

#### df -TH

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# df -TH
                  Size Used Avail Use% Mounted on
Filesystem
          Type
/dev/vda1
           ext4
                  43G 1.9G 39G 5% /
           devtmpfs 2.0G 0 2.0G 0% /dev
devtmpfs
          tmpfs 2.0G 0 2.0G 0% /dev/shm
tmpfs
          tmpfs 2.0G 9.0M 2.0G 1% /run
tmpfs 2.0G 0 2.0G 0% /svs/fs
tmpfs
                 tmpfs
          tmpfs
tmpfs
/dev/vdb1 ext4 106G 63M 101G 1% /mnt/sdc
```

New partition /dev/vdb1 is mounted on /mnt/sdc.

# **Ⅲ** NOTE

If the server is restarted, the mounting will become invalid. You can set automatic mounting for partitions at system start by modifying the /etc/fstab file. For details, see Setting Automatic Mounting at System Start.

----End

# **Setting Automatic Mounting at System Start**

To automatically mount disk partitions at system start, do not specify partitions, for example /dev/vdb1, in /etc/fstab because the sequence of cloud devices, and therefore their names may change during the server stop and start. You are advised to use the universally unique identifier (UUID) in /etc/fstab to set automatic mounting at system start.

#### □ NOTE

UUID is the unique character string for disk partitions in a Linux system.

**Step 1** Run the following command to guery 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.

**Step 2** Run the following command to open the **fstab** file using the vi editor:

# vi /etc/fstab

- **Step 3** Press **i** to enter the editing mode.
- **Step 4** Move the cursor to the end of the file and press **Enter**. Then, add the following information:

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

The preceding 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 **Step 1**.
- 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**: not use Linux dump backup. Normally, dump backup is not used, and you can set this parameter to **0**.
  - 1: use Linux dump backup.
- 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** so that the system checks the partitions in the ascending order of the values.

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

The system saves the configurations and exits the vi editor.

- **Step 6** Perform the following operations to verify the automatic mounting function:
  - 1. Run the following command to unmount the partition: **umount** *Disk partition*

In this example, run the following command:

# umount /dev/vdb1

- Run the following command to reload all the content in the /etc/fstab file: mount -a
- 3. 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, the automatic mounting function takes effect:

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

----End

# 5.6 Initializing a Data Disk Greater Than 2 TB in Windows (Windows Server 2008)

# **Scenarios**

This topic uses Windows Server 2008 R2 Standard 64bit to describe how to initialize a data disk whose capacity is larger than 2 TB. In the following operations, the capacity of the example disk is 3 TB.

The maximum disk capacity supported by MBR is 2 TB, and that supported by GPT is 18 EB. Therefore, use the GPT partition style if your disk capacity is larger than 2 TB. For details, see Initializing a Data Disk Greater Than 2 TB in Windows (Windows Server 2008). For details about disk partition styles, see Introduction to Data Disk Initialization Scenarios and Partition Styles.

The method for initializing a disk varies depending on the OS running on the server. This document is used for reference only. For the detailed operations and differences, see the product documents of the corresponding OS.

#### NOTICE

When using a disk for the first time, if you have not initialized it, including creating partitions and file systems, the additional space added to this disk in an expansion later may not be normally used.

# **Prerequisites**

- A data disk has been attached to a server and has not been initialized.
- You have logged in to the server.
  - For how to log in to an ECS, see the *Elastic Cloud Server User Guide*.
  - For how to log in to a BMS, see the *Bare Metal Server User Guide*.

# **Procedure**

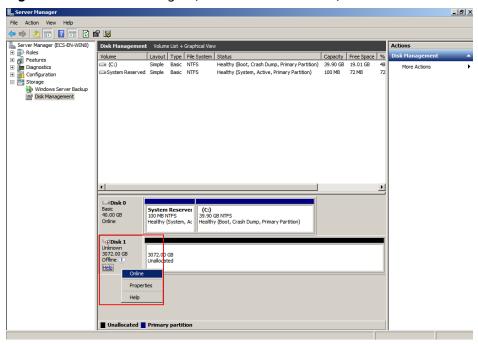
**Step 1** On the desktop of the server, click **Start**.

The **Start** window is displayed.

**Step 2** Right-click **Computer** and choose **Manage** from the short-cut menu.

The **Server Manager** window is displayed.

Figure 5-23 Server Manager (Windows Server 2008)



**Step 3** Disks are listed in the right pane. If the new disk is in the offline state, bring it online before initialize it.

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

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

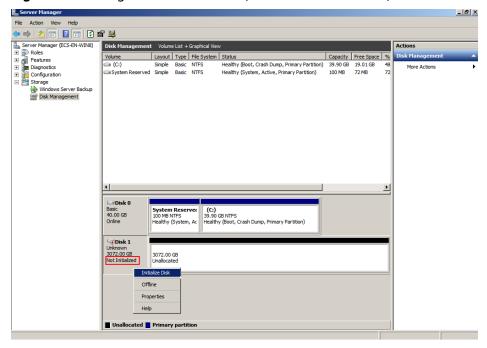


Figure 5-24 Bring online succeeded (Windows Server 2008)

Step 4 In the Disk 1 area, right-click and choose Initialize Disk from the shortcut menu.

The Initialize Disk dialog box is displayed.

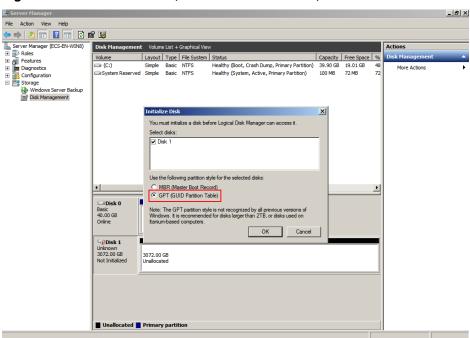


Figure 5-25 Initialize Disk (Windows Server 2008)

**Step 5** In the **Initialize Disk** dialog box, the to-be-initialized disk is selected. In this example, the disk capacity is larger than 2 TB. Therefore, select **GPT (GUID Partition Table)** and click **OK**.

The **Server Manager** window is displayed.

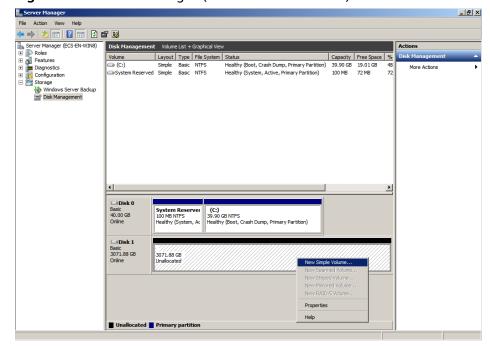


Figure 5-26 Server Manager (Windows Server 2008)

# **NOTICE**

The maximum disk capacity supported by MBR is 2 TB, and that supported by GPT is 18 EB. Because a data disk currently supports up to 32 TB, use the GPT partition style if your disk capacity is larger than 2 TB.

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

**Step 6** Right-click at the unallocated disk space and choose **New Simple Volume** from the shortcut menu.

The New Simple Volume Wizard window is displayed.

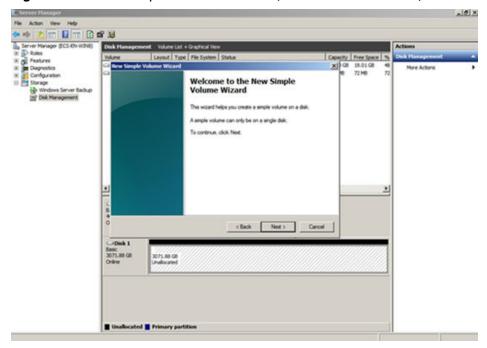


Figure 5-27 New Simple Volume Wizard (Windows Server 2008)

Step 7 Follow the prompts and click Next.

The **Specify Volume Size** page is displayed.

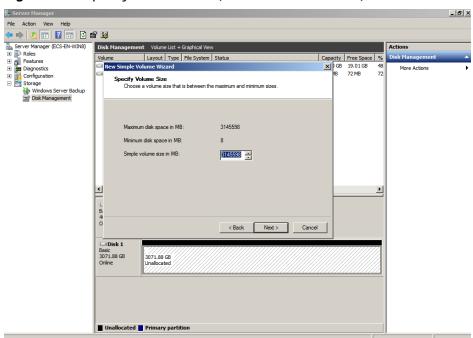


Figure 5-28 Specify Volume Size (Windows Server 2008)

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

The **Assign Drive Letter or Path** page is displayed.

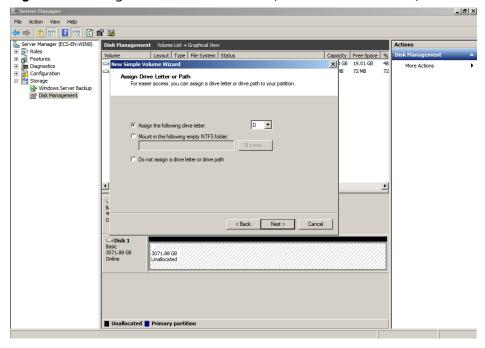
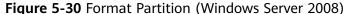
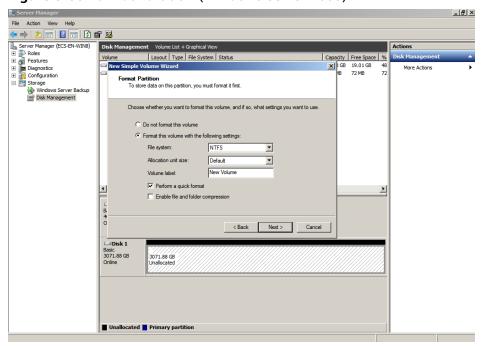


Figure 5-29 Assign Driver Letter or Path (Windows Server 2008)

**Step 9** Assign a drive letter or path to your partition and click **Next**. The system assigns drive letter D by default. In this example, the default setting is used.

The Format Partition page is displayed.





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

The Completing the New Simple Volume Wizard page is displayed.



Figure 5-31 Completing the New Simple Volume Wizard

# **NOTICE**

The partition sizes supported by file systems vary. Therefore, you are advised to choose an appropriate file system based on your service requirements.

# Step 11 Click Finish.

Wait for the initialization to complete. When the volume status changes to **Healthy**, the initialization has finished successfully, as shown in **Figure 5-32**.

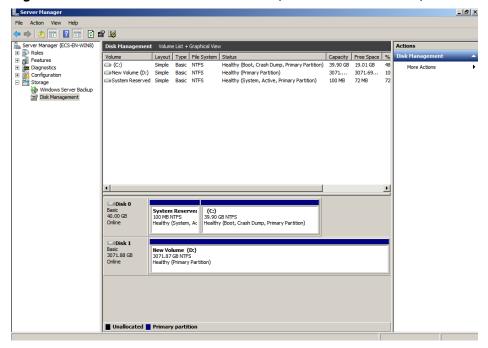


Figure 5-32 Disk initialization succeeded (Windows Server 2008)

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

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

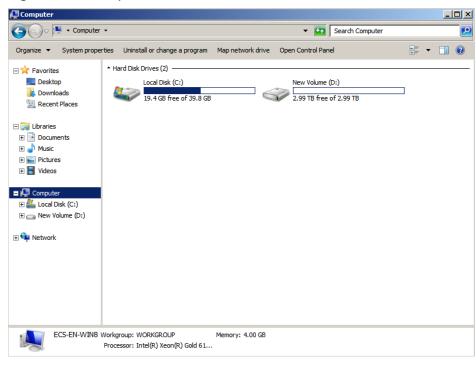


Figure 5-33 Computer (Windows Server 2008)

----End

# 5.7 Initializing a Data Disk Greater Than 2 TB in Windows (Windows Server 2012)

# **Scenarios**

This topic uses Windows Server 2012 R2 Standard 64bit to describe how to initialize a data disk whose capacity is larger than 2 TB. In the following operations, the capacity of the sample disk is 3 TB.

The maximum disk capacity supported by MBR is 2 TB, and that supported by GPT is 18 EB. Therefore, use the GPT partition style if your disk capacity is larger than 2 TB. For details, see Initializing a Data Disk Greater Than 2 TB in Windows (Windows Server 2008). For details about disk partition styles, see Introduction to Data Disk Initialization Scenarios and Partition Styles.

The method for initializing a disk varies depending on the OS running on the server. This document is used for reference only. For the detailed operations and differences, see the product documents of the corresponding OS.

# **NOTICE**

When using a disk for the first time, if you have not initialized it, including creating partitions and file systems, the additional space added to this disk in an expansion later may not be normally used.

# **Prerequisites**

- A data disk has been attached to a server and has not been initialized.
- You have logged in to the server.
  - For how to log in to an ECS, see the *Elastic Cloud Server User Guide*.
  - For how to log in to a BMS, see the Bare Metal Server User Guide.

# **Procedure**

**Step 1** On the desktop of the server, click in the lower area.

The **Server Manager** window is displayed.

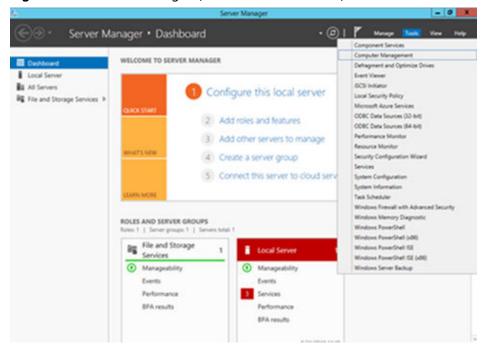
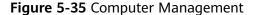
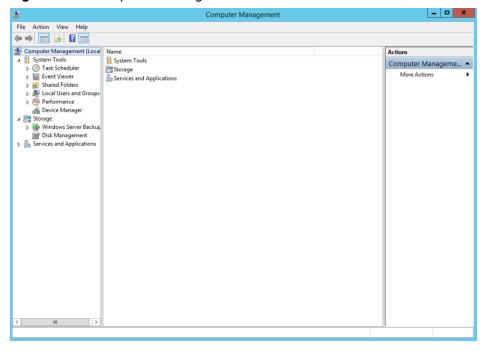


Figure 5-34 Server Manager (Windows Server 2012)

**Step 2** In the upper right corner, choose **Tools** > **Computer Management**.

The **Computer Management** window is displayed.





Step 3 Choose Storage > Disk Management.

Disks are displayed in the right pane.

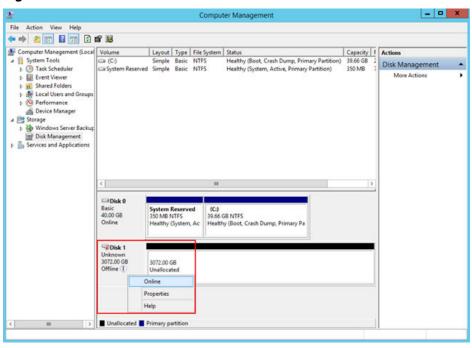


Figure 5-36 Disk list

**Step 4** (Optional) If the new disk is in the offline state, bring it online before initializing it.

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

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

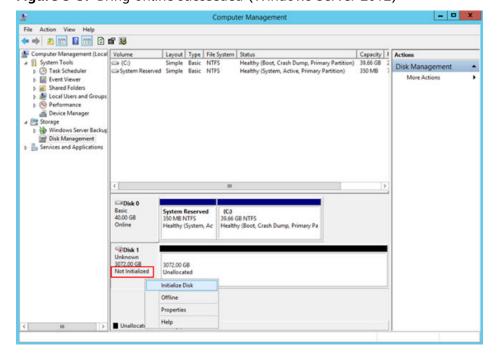
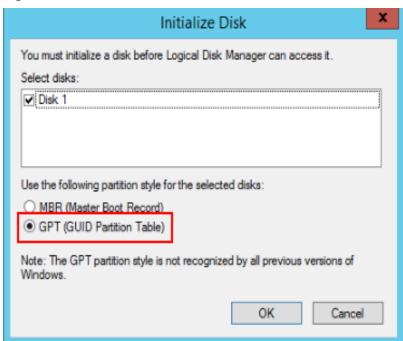


Figure 5-37 Bring online succeeded (Windows Server 2012)

**Step 5** In the **Disk 1** area, right-click and choose **Initialize Disk** from the shortcut menu.

The Initialize Disk dialog box is displayed.

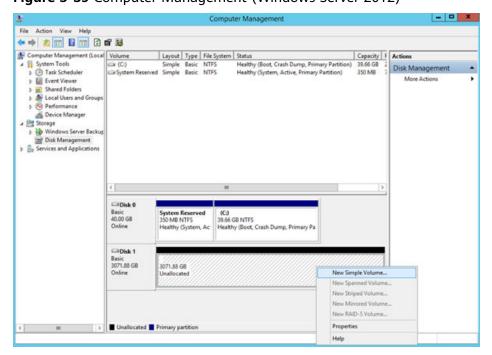
Figure 5-38 Initialize Disk (Windows Server 2012)



**Step 6** In the **Initialize Disk** dialog box, the to-be-initialized disk is selected. In this example, the disk capacity is larger than 2 TB. Therefore, select **GPT (GUID Partition Table)** and click **OK**.

The **Computer Management** window is displayed.

Figure 5-39 Computer Management (Windows Server 2012)



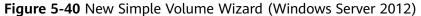
# **NOTICE**

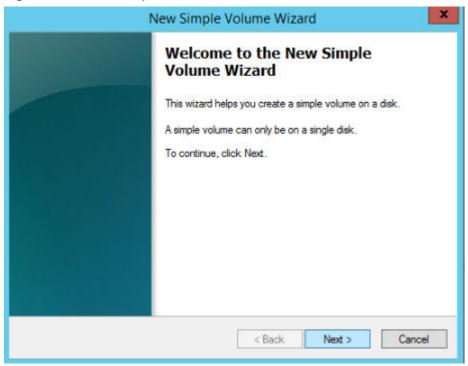
The maximum disk capacity supported by MBR is 2 TB, and that supported by GPT is 18 EB. Because a data disk currently supports up to 32 TB, use the GPT partition style if your disk capacity is larger than 2 TB.

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

**Step 7** Right-click at the unallocated disk space and choose **New Simple Volume** from the shortcut menu.

The New Simple Volume Wizard window is displayed.





**Step 8** Follow the prompts and click **Next**.

The **Specify Volume Size** page is displayed.

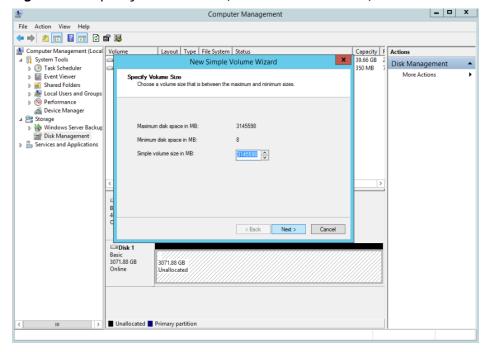
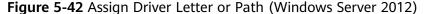


Figure 5-41 Specify Volume Size (Windows Server 2012)

**Step 9** Specify the volume size and click **Next**. The system selects the maximum volume size by default. You can specify the volume size as required. In this example, the default setting is used.

The **Assign Drive Letter or Path** page is displayed.

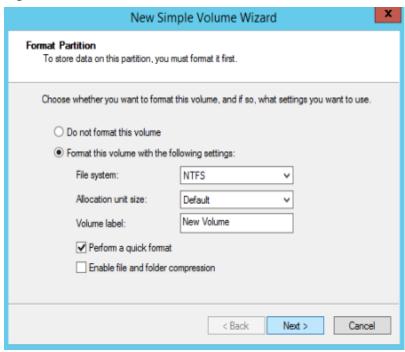




**Step 10** Assign a drive letter or path to your partition and click **Next**. The system assigns drive letter D by default. In this example, the default setting is used.

The Format Partition page is displayed.

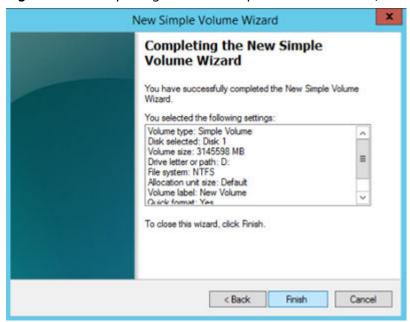
Figure 5-43 Format Partition (Windows Server 2012)



**Step 11** Specify format settings and click **Next**. The system selects the NTFS file system by default. You can specify the file system type as required. In this example, the default setting is used.

The Completing the New Simple Volume Wizard page is displayed.

Figure 5-44 Completing the New Simple Volume Wizard (Windows Server 2012)



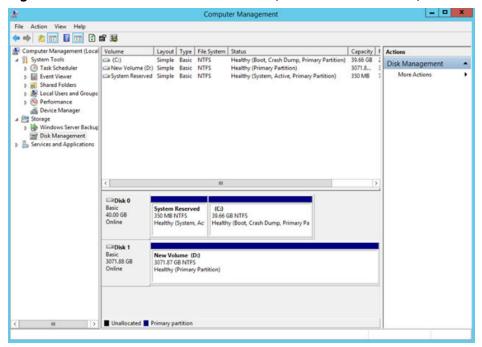
# **NOTICE**

The partition sizes supported by file systems vary. Therefore, you are advised to choose an appropriate file system based on your service requirements.

# Step 12 Click Finish.

Wait for the initialization to complete. When the volume status changes to **Healthy**, the initialization has finished successfully, as shown in **Figure 5-45**.

Figure 5-45 Disk initialization succeeded (Windows Server 2012)



Step 13 After the volume is created, click and check whether a new volume appears in This PC. In this example, New Volume (D:) is the new volume.

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

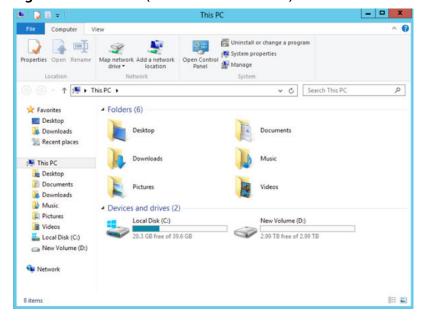


Figure 5-46 This PC (Windows Server 2012)

----End

# 5.8 Initializing a Data Disk Greater Than 2 TB in Linux (parted)

# **Scenarios**

This topic uses CentOS 7.4 64bit to describe how to use parted to initialize a data disk whose capacity is larger than 2 TB. In the following operations, the capacity of the sample disk is 3 TB.

The maximum disk capacity supported by MBR is 2 TB, and that supported by GPT is 18 EB. Therefore, use the GPT partition style if your disk capacity is larger than 2 TB. In Linux, if you choose to use the GPT partition style, the fdisk partitioning tool cannot be used. Use the parted partitioning tool instead. For details about disk partition styles, see Introduction to Data Disk Initialization Scenarios and Partition Styles.

The method for initializing a disk varies depending on the OS running on the server. This document is used for reference only. For the detailed operations and differences, see the product documents of the corresponding OS.

#### **NOTICE**

When using a disk for the first time, if you have not initialized it, including creating partitions and file systems, the additional space added to this disk in an expansion later may not be normally used.

# **Prerequisites**

- A data disk has been attached to a server and has not been initialized.
- You have logged in to the server.
  - For how to log in to an ECS, see the *Elastic Cloud Server User Guide*.
  - For how to log in to a BMS, see the Bare Metal Server User Guide.

# **Creating and Mounting a Partition**

The following example shows you how a new partition can be created on a new data disk that has been attached to a server. The partition will be created using parted, and GPT is used as the partition style. Furthermore, the partition will be formatted using the ext4 file system, mounted on /mnt/sdc, and configured automatic mounting at system start.

**Step 1** Run the following command to query information about the new data disk:

#### lsblk

Information similar to the following is displayed:

```
[root@ecs-centos74 ~]# lsblk

NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT

vda 253:0 0 40G 0 disk

—vda1 253:1 0 1G 0 part /boot

—vda2 253:2 0 39G 0 part /

vdb 253:16 0 3T 0 disk
```

In the command output, the server contains two disks. **/dev/vda** is the system disk, and **/dev/vdb** is the new data disk.

**Step 2** Run the following command to enter parted to partition the new data disk:

```
parted New data disk
```

In this example, run the following command:

# parted /dev/vdb

Information similar to the following is displayed:

```
[root@ecs-centos74 ~]# parted /dev/vdb
GNU Parted 3.1
Using /dev/vdb
Welcome to GNU Parted! Type 'help' to view a list of commands.

(narted)
```

**Step 3** Enter **p** and press **Enter** to view the current disk partition style.

Information similar to the following is displayed:

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

In the command output, the **Partition Table** value is **unknown**, indicating that no partition style is set for the new disk.

# **Step 4** Run the following command to set the disk partition style:

# mklabel Disk partition style

In this example, run the following command to set the disk partition style to GPT: (Disk partition styles can be MBR or GPT.)

# mklabel gpt

#### **NOTICE**

The maximum disk capacity supported by MBR is 2 TB, and that supported by GPT is 18 EB. Because a data disk currently supports up to 32 TB, use the GPT partition style if your disk capacity is larger than 2 TB.

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

# **Step 5** Enter **p** and press **Enter** to view the disk partition style.

Information similar to the following is displayed:

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

Number Start End Size File system Name Flags

**Step 7** Run the following command and press **Enter**:

mkpart Partition name Start sector End sector

In this example, run the following command:

# mkpart opt 2048s 100%

In this example, one partition is created for the new data disk. Value **2048s** indicates the disk start sector, and **100%** indicates the disk end sector. The two values are used for reference only. You can determine the number of partitions and the partition size based on your service requirements.

**Step 6** Enter **unit s** and press **Enter** to set the measurement unit of the disk to sector.

Information similar to the following is displayed:

(parted) mkpart opt 2048s 100%

Warning: The resulting partition is not properly aligned for best performance. Ignore/Cancel? Ignore

If the preceding warning message is displayed, enter **Ignore** to ignore the performance warning.

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

Information similar to the following is displayed:

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

Number Start End Size File system Name Flags
1 2048s 6442448895s 6442446848s opt
```

Details about the **dev/vdb1** partition are displayed.

- **Step 9** Enter **q** and press **Enter** to exit parted.
- **Step 10** Run the following command to view the disk partition information:

#### lsblk

Information similar to the following is displayed:

```
[root@ecs-centos74 ~]# lsblk

NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT

vda 253:0 0 40G 0 disk

—vda1 253:1 0 1G 0 part /boot

—vda2 253:2 0 39G 0 part /

vdb 253:16 0 3T 0 disk

—vdb1 253:17 0 3T 0 part
```

In the command output, /dev/vdb1 is the partition you created.

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

```
mkfs -t File system format /dev/vdb1
```

In this example, run the following command to set the **ext4** file system for the new partition:

# mkfs -t ext4 /dev/vdb1

Information similar to the following is displayed:

```
[root@ecs-centos74 ~]# mkfs -t ext4 /dev/vdb1
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
201326592 inodes, 805305856 blocks
40265292 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2952790016
24576 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
     32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
     4096000, 7962624, 11239424, 20480000, 23887872, 71663616, 78675968,
     102400000, 214990848, 512000000, 550731776, 644972544
Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

The formatting takes a period of time. Observe the system running status and do not exit.

#### **NOTICE**

The partition sizes supported by file systems vary. Therefore, you are advised to choose an appropriate file system based on your service requirements.

**Step 12** Run the following command to create a mount point:

mkdir Mount point

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

mkdir /mnt/sdc

**Step 13** Run the following command to mount the new partition on the created mount point:

mount Disk partition Mount point

In this example, run the following command to mount the new partition /dev/vdb1 on /mnt/sdc:

mount /dev/vdb1 /mnt/sdc

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

# df -TH

Information similar to the following is displayed:

```
[root@ecs-centos74 ~]# df -TH
Filesystem Type Size Used Avail Use% Mounted on
/dev/vda2 ext4 42G 1.5G 38G 4% /
devtmpfs devtmpfs 2.0G 0 2.0G 0% /dev
tmpfs tmpfs 2.0G 0 2.0G 0% /dev/shm
tmpfs tmpfs 2.0G 8.9M 2.0G 1% /run
tmpfs tmpfs 2.0G 0 2.0G 0% /sys/fs/cgroup
/dev/vda1 ext4 1.1G 153M 801M 17% /boot
tmpfs tmpfs 398M 0 398M 0% /run/user/0
/dev/vdb1 ext4 3.3T 93M 3.1T 1% /mnt/sdc
```

New partition **dev/vdb1** is mounted on **/mnt/sdc**.

----End

# **Setting Automatic Mounting at System Start**

To automatically mount disk partitions at system start, do not specify partitions, for example <code>/dev/vdb1</code>, in <code>/etc/fstab</code> because the sequence of cloud devices, and therefore their names may change during the server stop and start. You are advised to use the universally unique identifier (UUID) in <code>/etc/fstab</code> to set automatic mounting at system start.

□ NOTE

UUID is the unique character string for disk partitions in a Linux system.

**Step 1** 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.

**Step 2** Run the following command to open the **fstab** file using the vi editor:

# vi /etc/fstab

- **Step 3** Press **i** to enter the editing mode.
- **Step 4** Move the cursor to the end of the file and press **Enter**. Then, add the following information:

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

The preceding 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 **Step 1**.
- The second column indicates the directory on which the partition is mounted. You can guery the mount point using the **df-TH** command.
- The third column indicates the file system format of the partition. You can guery 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**: not use Linux dump backup. Normally, dump backup is not used, and you can set this parameter to **0**.
  - 1: use Linux dump backup.
- 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** so that the system checks the partitions in the ascending order of the values.

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

The system saves the configurations and exits the vi editor.

- **Step 6** Perform the following operations to verify the automatic mounting function:
  - 1. Run the following command to unmount the partition:

umount Disk partition

In this example, run the following command:

umount /dev/vdb1

- 2. Run the following command to reload all the content in the /etc/fstab file: mount -a
- 3. 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, the automatic mounting function takes effect:

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

----End

# 6 Change History

Released On	Description
2018-04-30	This issue is the first official release.