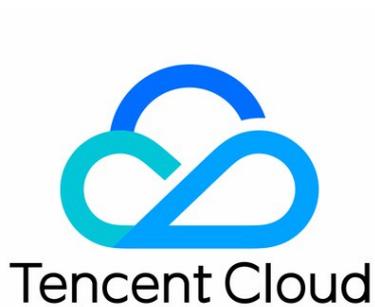


Cloud Block Storage Operation Guide



Copyright Notice

©2013–2023 Tencent Cloud. All rights reserved.

The complete copyright of this document, including all text, data, images, and other content, is solely and exclusively owned by Tencent Cloud Computing (Beijing) Co., Ltd. ("Tencent Cloud"); Without prior explicit written permission from Tencent Cloud, no entity shall reproduce, modify, use, plagiarize, or disseminate the entire or partial content of this document in any form. Such actions constitute an infringement of Tencent Cloud's copyright, and Tencent Cloud will take legal measures to pursue liability under the applicable laws.

Trademark Notice

 Tencent Cloud

This trademark and its related service trademarks are owned by Tencent Cloud Computing (Beijing) Co., Ltd. and its affiliated companies ("Tencent Cloud"). The trademarks of third parties mentioned in this document are the property of their respective owners under the applicable laws. Without the written permission of Tencent Cloud and the relevant trademark rights owners, no entity shall use, reproduce, modify, disseminate, or copy the trademarks as mentioned above in any way. Any such actions will constitute an infringement of Tencent Cloud's and the relevant owners' trademark rights, and Tencent Cloud will take legal measures to pursue liability under the applicable laws.

Service Notice

This document provides an overview of the as-is details of Tencent Cloud's products and services in their entirety or part. The descriptions of certain products and services may be subject to adjustments from time to time.

The commercial contract concluded by you and Tencent Cloud will provide the specific types of Tencent Cloud products and services you purchase and the service standards. Unless otherwise agreed upon by both parties, Tencent Cloud does not make any explicit or implied commitments or warranties regarding the content of this document.

Contact Us

We are committed to providing personalized pre-sales consultation and technical after-sale support. Don't hesitate to contact us at 4009100100 or 95716 for any inquiries or concerns.

Contents

Operation Guide

Operation Overview

Limits

Creating Cloud Disks

Attaching Cloud Disks

Initializing Cloud Disks

Initialization Scenarios

Initializing Cloud Disks (<2 TB)

Initializing Cloud Disks (≥2 TB)

Expanding cloud disks

Overview

Expanding cloud disks

Extending System Disk Partitions and File Systems Online

Extending Data Disk Partitions and File Systems (Windows)

Extending Data Disk Partitions and File Systems (Linux)

Extending Data Disk Partitions and File Systems Online

Extending Data Disk Partitions and File Systems Offline

Determining the Expansion Method

Extending MBR Partitions and File Systems Offline (Smaller than 2 TB)

Extending GPT Partitions and File Systems (Larger than 2 TB)

Extending Bare Data Disk File Systems

Adjusting cloud disk types

Detaching cloud disks

Viewing Cloud Disk Information

Terminating Cloud Disks

Adjusting Cloud Disk Performance

Cloud Disk Data Protection Configuration

Cloud Disk Data Backup Point

Adjusting Data Backup Point Quota

Using Data Backup Point to Restore Cloud Disk Data

Converting Data Backup Point into Snapshot

Manual Creation of Data Backup Points

Deleting Data Backup Point

Managing Snapshots

Creating Snapshot

Rolling back Data from Snapshot

Creating Cloud Disks Using Snapshots

Scheduled Snapshots

Cross-region snapshot replication

Deleting Snapshots

Shared Snapshots

Multi-mounting of Cloud Disk

Monitoring and Alarming

Operation Guide

Operation Overview

Last updated: 2023-09-19 18:24:45

While utilizing Tencent Cloud CBS, you may encounter issues such as managing cloud disks, managing snapshots, and monitoring and alerts. This document outlines the common operations for using Tencent Cloud CBS for your reference.

Cloud Disk

- [Creating Cloud Disk](#)
- [Attaching Cloud Disk](#)
- [Initializing Cloud Disks](#)
 - [Initialization Scenarios](#)
 - [Initialize Cloud Disks \(less than 2 TB\)](#)
 - [Initializing Cloud Disks \(\$\geq\$ 2 TB\)](#)
- [Expanding cloud disks](#)
 - [Overview of Expansion Scenarios](#)
 - [Expanding Cloud Disks](#)
 - [Extending Partitions and File Systems \(Windows\)](#)
 - [Extending Partitions and File Systems \(Linux\)](#)
- [Unmounting Cloud Disk](#)
- [Terminating Cloud Disks](#)

Snapshot

- [Create Snapshot](#)
- [Rolling Back Data from Snapshots](#)
- [Create Cloud Disk from Snapshot](#)
- [Scheduled Snapshots](#)
- [Cross-Region Snapshot Replication](#)
- [Deleting Snapshots](#)

Monitoring and alarms

- [Monitoring and Alerts](#)

Limits

Last updated: 2023-09-22 14:38:43

Use Limits on Cloud Disk

Item	Description
Enhanced SSD usage	<ul style="list-style-type: none"> Currently, Enhanced SSD is only available in certain AZs and will be supported in more AZs in the future. You can check the information on the purchase page. Enhanced SSDs can only be mounted on new generation instances such as S5, M5, SA2, and above, created after August 1, 2020 (for example, for standard instances, series not lower than the fifth generation like S5, S6 are supported). When mounted on older generation instances, the promised performance cannot be guaranteed. Enhanced SSD cannot be used as the system disk. Enhanced SSD cannot be encrypted. Enhanced SSD cannot be upgraded from other disk types.
ulTra SSD usage	<ul style="list-style-type: none"> Currently, ulTra SSD is only available in certain AZs and will be supported in more AZs in the future. You can check the information on the purchase page. Ultra-fast SSD cloud disks can only be purchased in conjunction with the S5se instance type. ulTra SSD cannot be purchased independently. ulTra SSD cannot be used as the system disk. ulTra SSD cannot be encrypted. ulTra SSD cannot be upgraded from other disk types.
Elastic cloud disk capability	Since May 2018, all data disks purchased with CVMs are elastic cloud disks, which can be detached from and reattached to CVMs. This feature is supported in all availability zones.
Cloud disk performances	<p>I/O performance is effective simultaneously.</p> <p>For instance, a 1TB SSD cloud disk can reach a maximum random IOPS of 26,000, which means both read IOPS and write IOPS can achieve this value. However, due to multiple performance constraints, in this case, I/O with a block size of 4KB/8KB can reach the maximum IOPS, but I/O with a block size of 16KB cannot reach the maximum IOPS (as the throughput has already reached the limit of 260MB/s).</p>
Max attachable elastic cloud disks per CVM	Up to 20.
Number of cloud disks created at a time	Up to 50.
Limits on the number of instances to which a cloud disk can be attached	<ul style="list-style-type: none"> Cloud disks and instances (CVM or Lighthouse instances) to which they are attached must be in the same AZ. The Cloud Disks of CVM instances and the Cloud Disks of Light Cloud instances are independent of each other. A Cloud Disk of a CVM instance cannot be attached to a Light Cloud instance, and vice versa.
Releasing overdue cloud disks	<p>If a monthly subscription elastic cloud disk is not renewed within 7 days after its expiration, the system will forcibly unmount the disk from the cloud server and recycle it to the recycle bin. For specific recycling mechanisms, please refer to the Arrears Description.</p> <p>At present, when a monthly subscribed elastic cloud disk is attached to a monthly subscribed cloud server, you can choose the following renewal methods based on your actual needs:</p> <ul style="list-style-type: none"> Unified expiry time with the CVM Enable monthly auto-renewal of the cloud disks

- Attach directly without enabling auto-renewal

Use Limits on Snapshot

Item	Description
Supported disk type	You can only use the data disk snapshot to create elastic cloud disks, while using the system disk snapshot to create a custom image.
Capacity of cloud disk created using a snapshot	The capacity of the cloud disk created using a snapshot should be greater than or equal to that of the snapshot.
Snapshot rollback	Snapshots can only be rolled back to the original cloud disk from which they were created. If you wish to create a new cloud disk using an existing snapshot, please refer to Create Cloud Disk from Snapshot .
Total snapshot size	No limit.
Snapshot quota in one region	64 + the number of cloud disks in the region * 64.

Use Limits on Scheduled Snapshot Policy

Item	Description
Scheduled snapshot policy quota in one region	A single Tencent Cloud account can set a maximum of 32 scheduled snapshot policies in each region.
Number of scheduled snapshot policies being associated with one cloud disk	A cloud disk can associate with a maximum of 10 scheduled snapshot policies in the same region.
Number of cloud disks associated with one scheduled snapshot policy	A scheduled snapshot policy can be associated with a maximum of 200 cloud disks in the same region.

Creating Cloud Disks

Last updated: 2023-09-19 18:28:09

Scenario

You can create a cloud disk and attach it to any cloud server within the same availability zone. Through block storage device mapping, the cloud disk is recognized and utilized by the cloud server, and it can reach its peak performance immediately after creation.

You can choose to create different types of cloud disks based on your business needs. For more information about cloud disk types, please refer to [Cloud Disk Types](#).

Preparations

Before creating a cloud disk, you need to sign up for a [Tencent Cloud](#) account and complete [identity verification](#).

Instructions

Creating a Cloud Disk via the Console

1. Log in to the [Cloud Disk Console](#) and select a region at the top of the page.
2. Navigate to the "Cloud Disk" page and click on **Create**.
3. Configure the following parameters in the pop-up window:

Configuration Item	Description
Availability Zones	Required. The availability zone of the cloud disk cannot be modified after the cloud disk has been created.
Cloud Disk Type	Required. The types of cloud disks include: <ul style="list-style-type: none">• Premium Cloud Disk• Balanced SSD• SSD cloud disk• Enhanced SSD <p>The ulTra SSD can only be purchased in conjunction with the storage-enhanced cloud server and cannot be purchased separately.</p> <p>For more detailed information about cloud disk types, please refer to Cloud Disk Types.</p>
Quick Disk Creation	Optional. When creating a cloud disk using a snapshot, you need to check the box "Create Cloud Disk with Snapshot" and select the snapshot you wish to use. <ul style="list-style-type: none">• You can adjust the capacity of a cloud disk created from a snapshot to be greater than the default value (snapshot size).• When you create a cloud disk using a snapshot, the disk type is the same as that of the snapshot's source disk. You can change the disk type.
Capacity	Required. The capacity and specifications of the cloud disk are as follows: <ul style="list-style-type: none">• Premium Cloud Disk: 10 GB – 32,000 GB• Balanced SSD Cloud Disk: 20 GB – 32,000 GB• SSD Cloud Disk: 20 GB – 32,000 GB• Enhanced SSD: 20 GB – 32,000 GB

	When creating a cloud disk from a snapshot, the capacity cannot be less than the size of the snapshot. If you do not specify the capacity of the cloud disk, it will default to the size of the snapshot.
Scheduled Snapshot	Optional. You can associate a regular snapshot policy when creating a cloud disk to manage the snapshots of the cloud disk regularly. Currently, a free quota of 80GB is provided in each region within the Chinese mainland. For more details, please refer to Snapshot Billing Overview .
Disk Name	Optional. The field can contain up to 128 characters, starting with a letter (upper or lower case) or a Chinese character. It can include letters (upper or lower case), Chinese characters, digits, and special symbols <code>._-:</code> . Modifications are permitted after the cloud disk has been created. <ul style="list-style-type: none"> Creating a Single Cloud Disk: The disk name is the name of the cloud disk. Batch creation of cloud disks: When creating multiple cloud disks at once, the disk name serves as the prefix of the cloud disk name, with the final cloud disk name being composed as <code>Disk Name_Number</code>.
Set Backup Point	Optional. You can set the data backup point quota when creating a cloud disk. Once set, Tencent Cloud will automatically create and manage the data backup points for your cloud disk. For more information about data backup points, please refer to Cloud Disk Data Backup Point .
Project	Required. When creating a cloud disk, you can assign it to a specific project. The default project is the default project .
Tag	Optional. You can bind tags to the cloud disk when creating it. Tags are used to identify cloud resources, and you can categorize and search for cloud resources through tags. For more detailed information about tags, please refer to the Tag Product Documentation .
Billing	Required. The cloud disk supports the following two billing types: <ul style="list-style-type: none"> For a Monthly Subscription, if the billing mode is set to Monthly Subscription, then the "purchase duration" must be set. Pay as you go.
Purchase quantity	Optional. The default quantity is 1, indicating the creation of a single cloud disk. Currently, up to 50 cloud disks can be created in bulk.
Duration	<ul style="list-style-type: none"> If the "Billing Mode" is set to Monthly Subscription, this parameter is mandatory. The selectable time range is from 1 month to 5 years. If the "Billing Mode" is set to pay-as-you-go, this parameter is not applicable.
Auto-Renewal	<ul style="list-style-type: none"> If the "Billing Mode" is set to Monthly Subscription, this is an optional parameter. When auto-renewal is selected, the device will be automatically renewed on a monthly basis after expiration, provided that the account balance is sufficient. If the "Billing Mode" is set to pay-as-you-go, this parameter is not applicable.
Expiration/Arrears Protection	Optional parameters. If your cloud disk is reclaimed due to reasons such as expiration of the annual/monthly package or account arrears, the system will automatically create a snapshot to protect your data.

4. Click **OK**.

- If the **Pay-as-You-Go** option is selected for "Billing Mode", then the creation is complete.
- If you select **Monthly Subscription** for "Billing Mode", you will be directed to the "Review Information" page.

- a. After confirming the specifications, select whether to use a voucher based on your actual situation and click **Confirm Purchase**.
 - b. Complete the payment.
5. You can view the created cloud disks on the [Cloud Disk List](#) page. The newly created elastic cloud disk is in the "To Be Mounted" state. You can refer to [Mounting Cloud Disks](#) to attach the cloud disk to a cloud server within the same availability zone.

Create Cloud Disk with Snapshot

If you wish to create a new cloud disk that directly includes relevant data, you can choose to [Create a Cloud Disk from a Snapshot](#).

Creating a Cloud Disk using API

You can use the CreateDisks interface to create a cloud disk. For specific operations, please refer to [Creating Cloud Disks](#).

Attaching Cloud Disks

Last updated: 2023-09-26 14:15:58

Scenario

You can attach elastic cloud disks to a CVM instance in the same availability zone as a data disk. Up to 20 data disks can be attached to a CVM instance.

After creating a cloud disk, it needs to be attached to an existing CVM instance in the same availability zone via the console or API. Please refer to this document to complete the cloud disk attachment process.

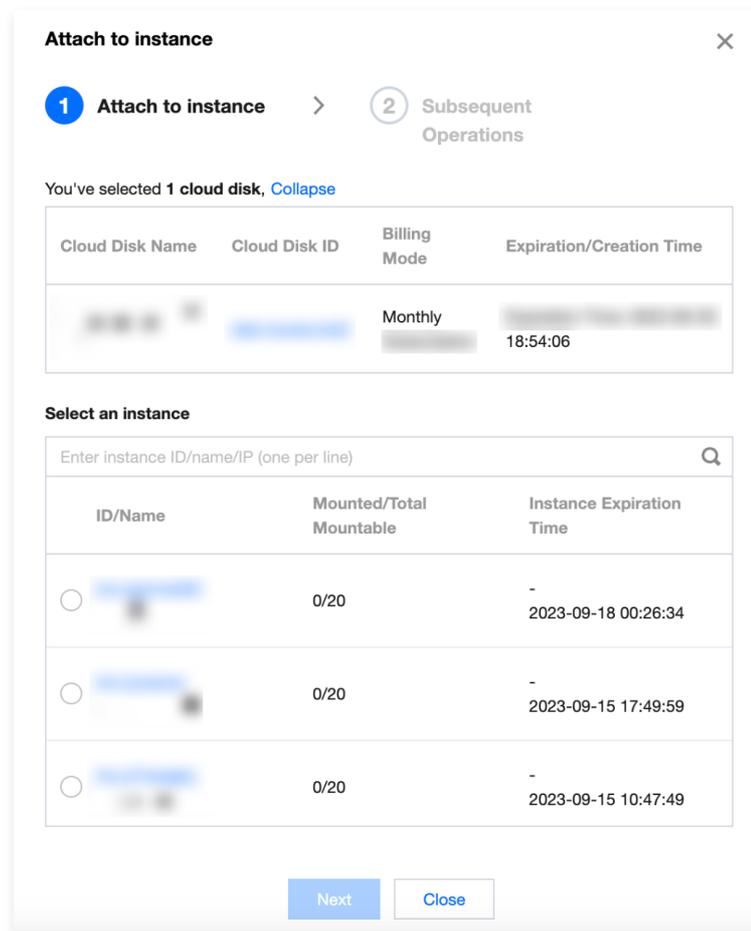
Note

Some Linux CVM instances may not recognize the elastic cloud disk. You can first enable the hot swapping feature in the CVM instance. For more information, please refer to [Enabling Disk Hot Swapping](#).

Instructions

Attaching cloud disks in the CBS console

1. Log in to the [Cloud Disk Console](#).
2. On the cloud disk list page, you can use the following methods to mount cloud disks.
 - **Attach a Single Cloud Disk:** Select **More** > **Attach** on the right side of the row where the cloud disk with the status "To be mounted" is located.
 - **Batch Attach Cloud Disks:** Select the cloud disks with the status "To be mounted", and click **Attach** above the cloud disk list to batch attach.
3. The **Attach to Instance** window will pop up, and you will enter the **Attach to Instance** step, as shown in the figure below:



Select the instance to which the disk needs to be attached and refer to the following information to choose the "Attach Option". Depending on whether your cloud disk has auto-renewal enabled and its expiration time, there will be three scenarios for the attach option. Please choose as needed:

- Unified expiry time with the instance (XXX)
- Auto-renew the disk monthly upon expiration (**Recommended**)
- Attach directly

4. Click **Next**. In the "Subsequent Operation Tips" step, understand the following:

After manually attaching a cloud disk, the disk is offline. You need to log in to the instance and perform initialization operations to make the cloud disk usable.

5. Click **Start Mounting**.

If the status of the cloud disk changes to "Mounted", it indicates that the mounting was successful.

6. Perform the subsequent operations corresponding to different cloud disk capacities.

Creation Mode	Cloud disk capacity	See Also
Manually created	Cloud disk capacity < 2 TB	Initialize Cloud Disks (less than 2 TB)
	Cloud disk capacity ≥ 2 TB	Initializing cloud disks (≥ 2 TB)
Created from a snapshot	Cloud disk capacity = Snapshot capacity	<ul style="list-style-type: none"> • To attach to a Windows CVM: After logging in to the instance, go to Server Management > Storage > Disk Management. The disk can be used once it is online. • To attach to a Linux CVM instance: After logging in to the instance, run the <code>mount <disk partition> <mount point></code> command, for example, <code>mount /dev/vdb /mnt</code> to start using it.
	Snapshot capacity < Cloud disk capacity ≤ 2 TB Or 2 TB < Snapshot capacity < Cloud disk capacity	<ul style="list-style-type: none"> • Attach to Windows CVM: Extending Partitions and File Systems (Windows) • Mount to Linux CVM instance: Extending Partitions and File Systems (Linux)
	Snapshot capacity ≤ 2 TB < Cloud disk capacity	<ul style="list-style-type: none"> • If the snapshot uses a MBR partition, You need to refer to Initializing Cloud Disk (Greater than or Equal to 2TB) to repartition using GPT. This operation will delete existing data. • If the snapshot uses a GPT partition, <ul style="list-style-type: none"> ○ Attach to Windows CVM: Extending Partitions and File Systems (Windows) ○ Mount to Linux CVM instance: Extending Partitions and File Systems (Linux)

Using the API to attach cloud disks

You can use the AttachDisks interface to attach cloud disks. For specific operations, please refer to [Attaching Cloud Disks](#).

Related Actions

Enabling the disk hot swapping feature

All images currently provided support the mounting/unmounting operations of elastic cloud disks. Before unmounting a cloud disk, you need to first execute **umount (Linux)** or **offline (Windows)** operations, otherwise, it may cause the CVM instance to fail to recognize the elastic cloud disk when it is mounted again.

Note

Hot swapping is only recommended for CVMs with the following operating systems.

CVM Operating System Type	Version
CentOS	5.11 64-bit
	5.11 32-bit
	5.8 64-bit
	5.8 32-bit
Debian	6.0.3 32-bit
Ubuntu	10.04 64-bit
	10.04 32-bit
openSUSE	12.3 64-bit
	12.3 32-bit

1. Log in to the Linux CVM as the root user. See [Logging in to Linux CVM](#).
2. Run the following command to add the driver.

```
modprobe acpihp
```

Note

If you need to load the `acpihp` driver module after shutting down or restarting the CVM instance, it is recommended to follow [Step 3](#) to set the `acpihp` module to load automatically at startup.

3. (Optional) Depending on the operating system, choose the corresponding method to set the `acpihp` module to load automatically at startup:

CentOS 5 Series

1. Run the following command to create and open the `acpihp.modules` file.

```
vi /etc/sysconfig/modules/acpihp.modules
```

2. Add the following content to the file, and save it.

```
#!/bin/bash
modprobe acpihp >& /dev/null
```

3. Run the following command to grant execute permissions on the file.

```
chmod a+x /etc/sysconfig/modules/acpihp.modules
```

Debian 6 series, Ubuntu 10.04 series

1. Run the following command to modify the file.

```
vi /etc/modules
```

2. Add the following content to the file, and save it.

```
acpiphp
```

openSUSE 12.3 Series

1. Run the following command to modify the file.

```
vi /etc/sysconfig/kernel
```

2. Add the following content to the file, and save it.

```
MODULES_LOADED_ON_BOOT="acpiphp"
```

Automatically Attach Cloud Disk When Creating an Instance

If you specify a custom image and data disk snapshot when creating a CVM instance, the concurrently created cloud disk will automatically be attached to the CVM instance. There is no need for partitioning, formatting, or other disk initialization operations to read and write data directly. However, your custom image and data disk snapshot must meet the following requirements:

Windows Instances

If you create a Windows CVM instance using a custom image, the platform will automatically attach the cloud disk created from the corresponding data disk snapshot of the image. The custom image and corresponding data disk snapshot must meet the following requirements:

- The data disk **must** be formatted to `NTFS` or `FAT32` before creating a snapshot.
- The SAN policy in the custom image is set to `onlineAll`.

Checking configurations

Although the public Windows images provided by Tencent Cloud have default settings, it is still recommended to check these configurations before creating a custom image. Please execute the following commands in sequence and check the returned results.

```
diskpart
```

```
san
```

Linux Instances

If you create a Linux CVM instance using a custom image, the platform will automatically mount the cloud disk created from the data disk snapshot corresponding to the image. The custom image and the corresponding data disk snapshot must meet the following requirements:

- The data disk **must** be formatted before creating a snapshot, that is, it has been successfully mounted on the source CVM instance.
- Before creating a custom image from the system disk, you need to add the following command to the `/etc/rc.local` file

to write the data disk mount point into the file.

```
mkdir -p <mount-point>  
mount <device-id> <mount-point>
```

Note

- `<mount-point>` should be set as the mount point of the file system, such as `/mydata` .
- `<device-id>` should be set to the actual file partition location. For instance, enter `/dev/vdb` when there is no partition with a file system, and `/dev/vdb1` when there is a partition with a file system.

Initializing Cloud Disks

Initialization Scenarios

Last updated: 2023-09-22 14:41:49

Upon manually attaching a cloud disk created via the console, it serves as a data disk for the cloud server and is offline by default. You need to initialize the data disk by formatting, partitioning, and creating a file system. Please select the initialization method based on your actual usage scenario:

- If the entire disk is presented as a single independent partition (i.e., there are no multiple logical disks, such as D drive /vdb1 and E drive /vdb2), it is recommended that you do not use partitioning, but instead directly build a file system on the raw device.
- If the entire disk needs to be presented as multiple logical partitions (that is, there are multiple logical disks), you need to perform the partitioning first, then create the file system on a partition.

The commonly used disk partition formats are the Main Boot Record (MBR) and the GUID Partition Table (GPT). If the partition format is changed after the disk is put into use, the original data on the disk will be erased. Therefore, please choose the partition format wisely based on actual requirements.

The brief introductions of the two partition formats are as follows:

Partition format	Max Disk Capacity Supported	Supported partition count	Partition tool
MBR	2TB	<ul style="list-style-type: none"> • Four primary partitions • Three primary partitions and one extended partition 	Windows Operating System: Disk Management Operating system: Linux <ul style="list-style-type: none"> • fdisk tool • Parted tool
GPT	18EB The maximum capacity currently supported by the cloud disk is 32 TB.	Unlimited partition quantity	Windows Operating System: Disk Management Linux Operating System: parted tool

Select the appropriate operations guide according to the disk capacity and the CVM operating system:

- For a disk smaller than 2 TB:
 - [Initializing Cloud Disks \(Windows\)](#)
 - [Initializing Cloud Disks \(Linux\)](#)
- For a disk equal to or larger than 2 TB:
 - [Initializing Cloud Disks \(Windows\)](#)
 - [Initializing Cloud Disks \(Linux\)](#)

Initializing Cloud Disks (<2 TB)

Last updated: 2023-09-25 17:21:03

Scenario

This guide provides instructions for initializing cloud disks with a capacity of less than 2 TB. For more information on disk initialization scenarios, please refer to [Introduction to Initialization Scenarios](#).

Preparations

The [cloud disk](#) has been attached to the cloud server.

Supports and Limits

- Before performing any operations on the cloud disk, it is advisable to familiarize yourself with the [Precautions for Using Cloud Disks](#) to prevent damage to important data.
- Formatting a data disk will erase all data. Make sure that the disk does not contain data, or important data has been backed up.
- To avoid service exceptions, ensure before formatting that the CVM has stopped external services.

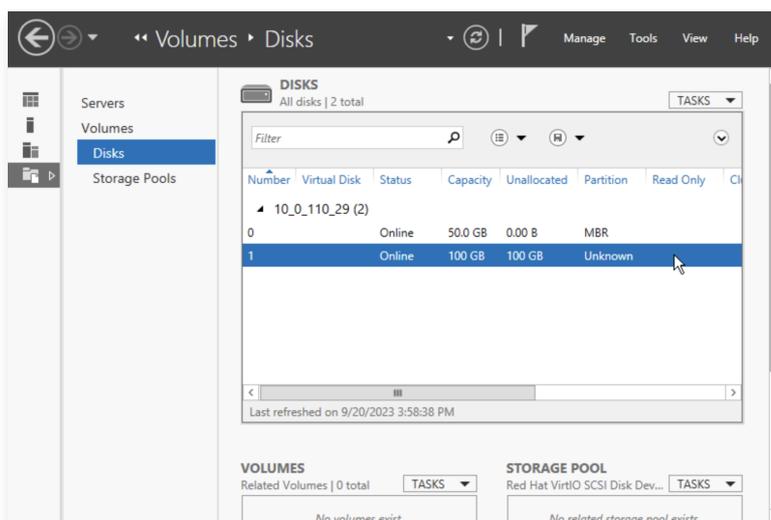
Instructions

Initializing Cloud Disk (Windows)

Note

This document uses a CVM with Windows Server 2012 R2 installed. Note that the steps may vary according to the operating system version.

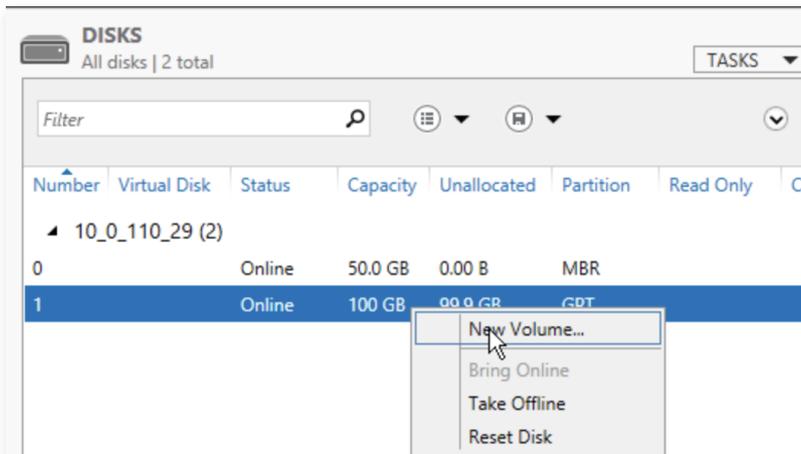
1. [Log in to the Windows CVM](#).
2. On the cloud server desktop, right-click on the  in the lower left corner.
3. From the pop-up menu, select **Disk Management** to open the **Disk Management** window.



Note

If the newly added disk is in an offline state (as shown above), you need to first execute [Step 4](#) to bring it online before executing [Step 5](#) for initialization. Otherwise, you can directly execute [Step 5](#) for initialization.

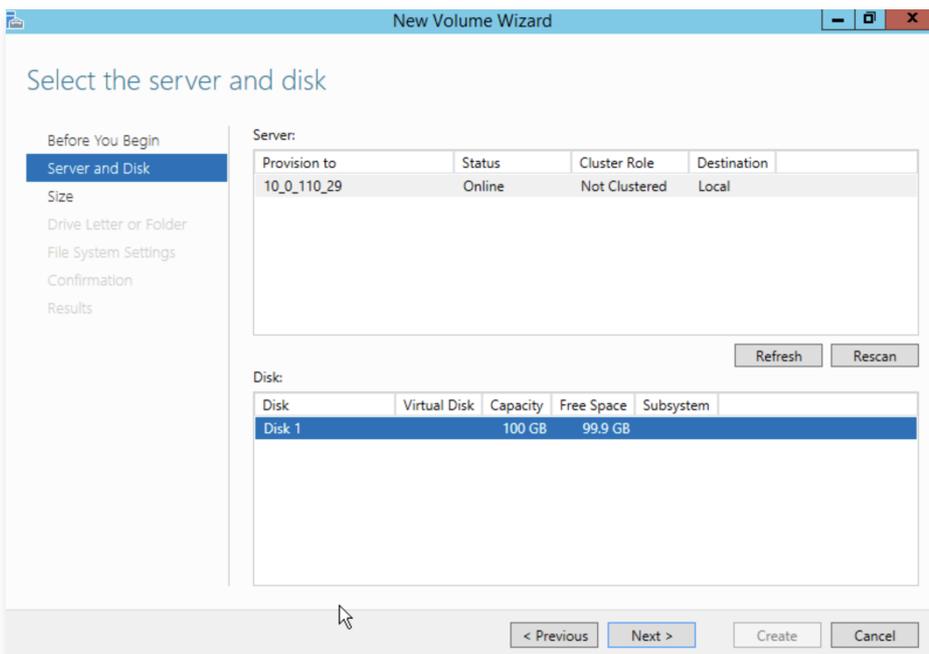
- In the right pane, a list of disks will appear. If the status of the disk is **Offline**, right-click on Disk 1 area and select **Online** from the menu list to bring it online. After going online, Disk changes from **Offline** status to **Not Initialized**.
- Right-click on the Disk 1 area and select **New Volume...** from the drop-down menu, as shown in the figure below:



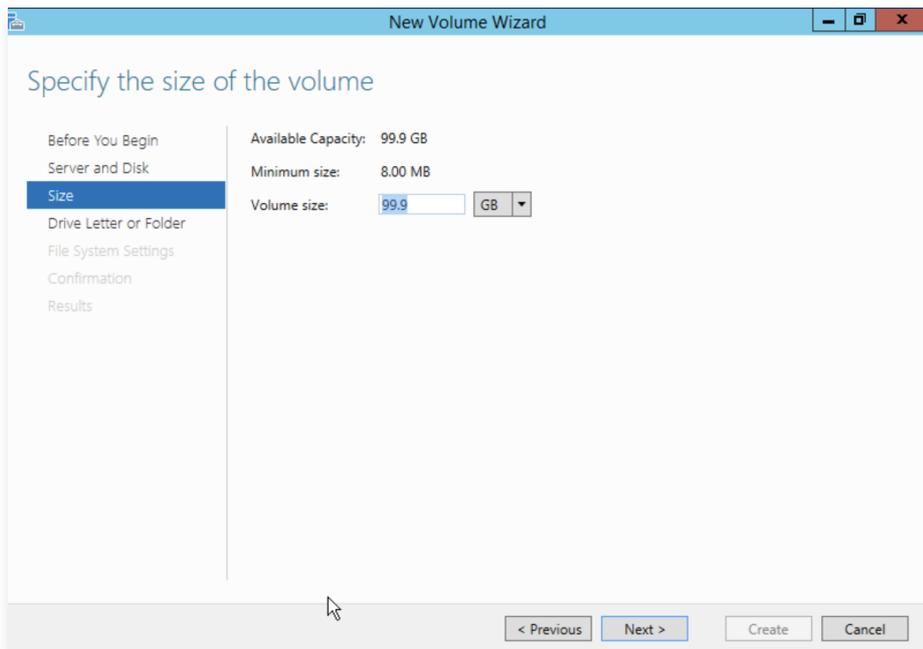
- In the **New Volume...** dialog box, select the disk that needs to be initialized, the confirm dialog appears and click **Confirm**. As shown below:

Note

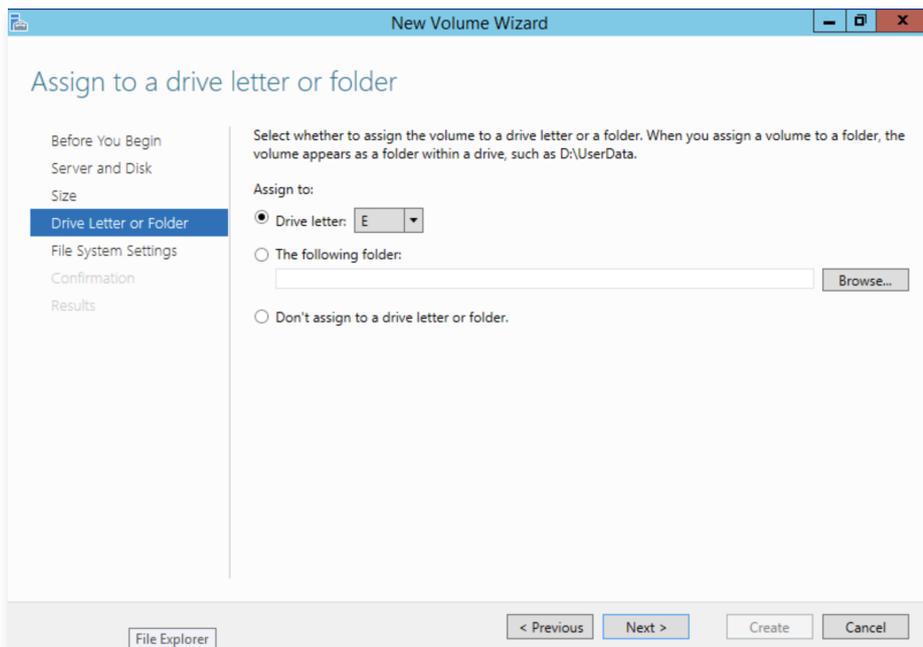
If a disk initializes after the disk is put into use, the original data on the disk will be erased.



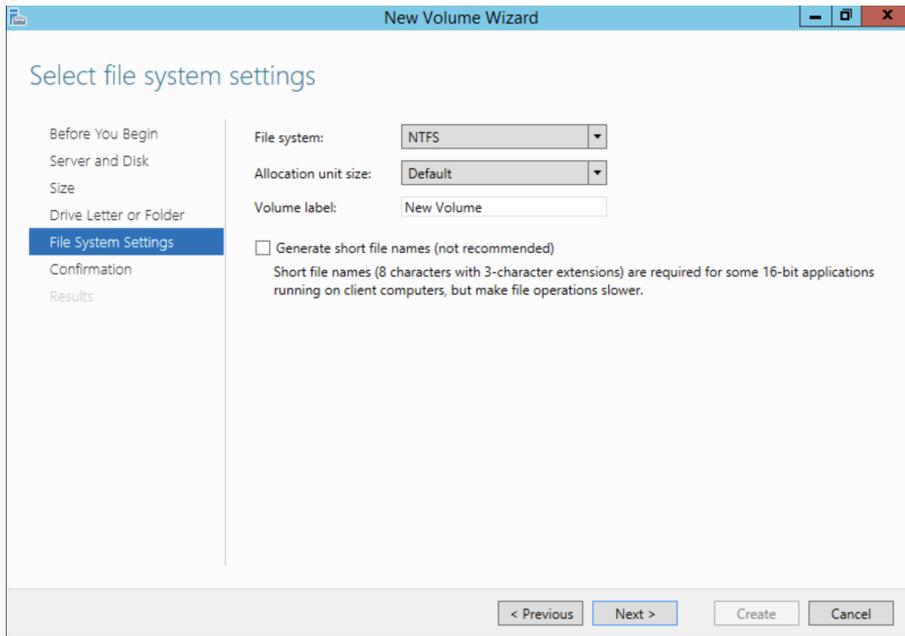
- Specify the volume size according to your needs. By default, the maximum value is selected. Click **Next**.



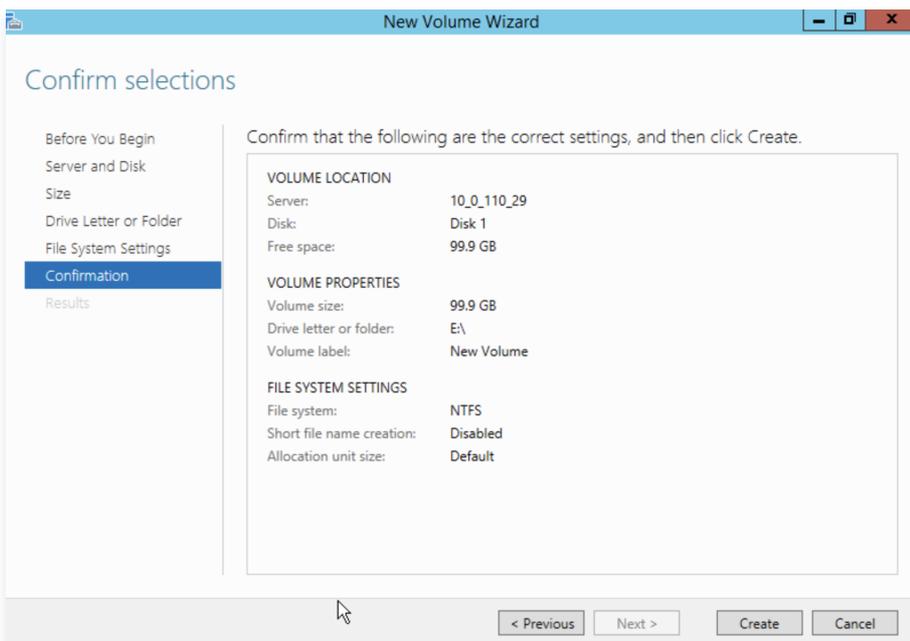
8. Assign a drive letter and click **Next**. As shown in the figure below:



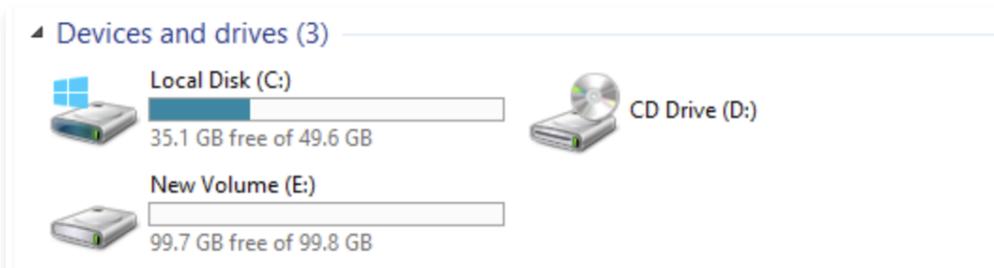
9. Select File System Setting, set the parameters as needed, take default for example and click **Next**.



10. Check all information and click **Create**. Please wait a moment for the system to complete the initialization. When the volume status is **Healthy**, it indicates that the disk initialization was successful.



After successful initialization, you can view the new disk in the **Computer** interface.



Initializing Cloud Disks (Linux)

Please select the initialization method based on your actual usage scenario:

- If the entire disk is presented as a single independent partition (i.e., there are no multiple logical disks, such as vdb1 and vdb2), we strongly recommend that you do not use partitions and directly [build a file system on the raw device](#).
- If the entire disk needs to be represented as multiple logical partitions (i.e., there are multiple logical drives), you must first perform partitioning operations, then [create a file system on the partition](#).

Creating file systems on bare devices

1. [Log into the Linux CVM](#).
2. Run the following command as the root user to view the disk name.

```
fdisk -l
```

If information similar to what is shown below is returned, it indicates that the current cloud server has two disks. `/dev/vda` is the system disk, and `/dev/vdb` is the newly added data disk.

```
[root@UM_16_14_centos ~]# fdisk -l
Disk /dev/vda: 53.7 GB, 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: 0x000d64b4

   Device Boot      Start         End      Blocks   Id  System
/dev/vda1 *          2048     104857599     52427776   83  Linux

Disk /dev/vdb: 64.4 GB, 64424509440 bytes, 125829120 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
```

3. Execute the following command to directly create a file system format for the `/dev/vdb` bare device.

```
mkfs -t <file system format> /dev/vdb
```

Different file systems support different partition sizes. Choose the file system according to your actual needs. For instance, when setting the file system to `EXT4` :

```
mkfs -t ext4 /dev/vdb
```

Note

The formatting takes a while. Pay attention to the system's running status and do not exit.

4. Run the following command to create a new mount point.

```
mkdir <mount point>
```

Taking the newly created mount point `/data` as an example:

```
mkdir /data
```

5. Run the following command to mount the new partition to the newly created mount point.

```
mount /dev/vdb <mount point>
```

For instance, a newly created partition is mounted to the `/data` mount point:

```
mount /dev/vdb /data
```

6. Run the following command to view the mount result.

```
df -TH
```

Note

If you do not need to configure disk automount at startup, skip the following steps.

7. Confirm the mounting method and obtain the corresponding information.

Based on business needs, you can use an elastic cloud disk's soft link, file system's UUID (universally unique identifier), or device name to automatically mount a disk. The descriptions and information acquisition methods are as follows:

Mount method	Pros and cons	Obtaining information
Use the soft link of the elastic cloud disk (recommended)	Advantages: Each elastic cloud disk has a fixed and unique soft link, which will not change with operations such as unmounting, mounting, formatting partitions, etc. Limitations: Only elastic cloud disks support soft links. It is unable to detect the formatting operations of partitions.	Run the following command to view the soft link of the elastic cloud disk. <pre>ls -l /dev/disk/by-id</pre>
Use the UUID of the file system	Automatic mounting configuration may fail due to changes in a file system's UUID. For example, reformatting a file system will change its UUID.	Run the following command to view the UUID of the file system. <pre>blkid /dev/vdb</pre>
Use device name	Auto-attaching configuration may fail due to changes in device name. For example, if an elastic cloud disk on the CVM is unmounted and then remounted, the device name may change when the operating system recognizes the file system again.	Run the following command to view the device name. <pre>fdisk -l</pre>

8. Execute the following command to backup the `/etc/fstab` file. The example below demonstrates backing up to the `/home` directory:

```
cp -r /etc/fstab /home
```

9. Run the following command to open the `/etc/fstab` file using the VI editor.

```
vi /etc/fstab
```

10. Press `i` to enter the editing mode.

11. Move the cursor to the end of the file, press **Enter**, and append the following content.

```
<Device information> <Mount point> <File system format> <File system installation option> <File system dump frequency> <File system check sequence at launch>
```

- **(Recommended)** Using the automatic mounting of the soft link of an elastic cloud disk as an example, combine with the previous examples to add:

```
/dev/disk/by-id/virtio-disk-drkhklpe /data ext4 defaults,nofail 0 0
```

Note

If you have multiple elastic cloud disks, you can differentiate them by comparing the `disk-xxxxx` with the cloud disk ID in the [console](#).

- Take automatic mounting using the UUID of the disk partition as an example. Add the following content:

```
UUID=d489ca1c-5057-4536-81cb-ceb2847f9954 /data ext4 defaults,nofail 0 0
```

- Take automatic mounting using the device name as an example. Add the following content:

```
/dev/vdb /data ext4 defaults,nofail 0 0
```

12. Press **Esc**, input `:wq`, and press **Enter** to save the settings and exit the editor.

13. Run the following command to verify if the `/etc/fstab` file has been written successfully.

```
mount -a
```

If the command runs successfully, the file has been written. The newly created file system will automatically mount when the operating system starts up.

Creating a file system on a partition

Note

This operation will use the `fdisk` partitioning tool in the CentOS 7.5 operating system to set the data disk `/dev/vdb` as the primary partition. The partition format is set to MBR by default, the file system is set to EXT4 format, it is mounted under `/data/newpart`, and it is set to automatically mount at startup. The formatting operations may vary in different operating systems, and this document is for reference only.

1. [Log into the Linux CVM](#).
2. Run the following command as the root user to view the disk name.

```
fdisk -l
```

If information similar to what is shown below is returned, it indicates that the current cloud server has two disks. `/dev/vda` is the system disk, and `/dev/vdb` is the newly added data disk.

```

[root@UM_16_14_centos ~]# fdisk -l

Disk /dev/vda: 53.7 GB, 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: 0x000d64b4

   Device Boot      Start         End      Blocks   Id  System
/dev/vda1  *          2048     104857599     52427776   83  Linux

Disk /dev/vdb: 64.4 GB, 64424509440 bytes, 125829120 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

```

- Run the following command to use the fdisk tool to partition the newly added data disk.

```
fdisk <add new data disk>
```

Taking the newly mounted data disk `/dev/vdb` as an example:

```
fdisk /dev/vdb
```

The echo information is displayed as follows:

```

[root@UM_16_14_centos ~]# 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 0xcb7892a8.

Command (m for help):

```

- Enter `n` and press **Enter** to start creating a new partition.

The returned information is as follows:

```

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

```

This indicates that the disk has two types of partitions:

- `p` denotes the primary partition.
- `e` denotes an extended partition.

- Taking the creation of a primary partition as an example, enter `p` and press **Enter** to start creating a primary partition.

The echo information is shown as follows:

```

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

```

Partition number refers to the primary partition number, which can be selected from 1 to 4.

- As an example, enter the primary partition number `1` and press **Enter**.

The following information will be displayed:

```

Partition number (1-4, default 1): 1
First sector (2048-125829119, default 2048):

```

First sector refers to the initial cylinder area, with a selectable range of 2048 – 125829119, defaulting to 2048.

- Using the default initial cylinder number 2048 as an example, press **Enter**.

The echo information is displayed as follows:

```
First sector (2048-125829119, default 2048):
Last sector, +sectors or +size{K,M,G} (2048-125829119, default 125829119):
```

Last sector refers to the end of the cylinder area, you can choose between 2048 – 125829119, with the default being 125829119.

8. Taking the default end cylinder number 125829119 as an example, press **Enter**.

The returned information is as follows:

```
Last sector, +sectors or +size{K,M,G} (2048-125829119, default 125829119):
Partition 1 of type Linux and of size 60 GiB is set
Command (m for help):
```

This indicates that the partitioning is complete, and a new partition has been created for the 60GB data disk.

9. Enter **p** and press **Enter** to view the detailed information of the newly created partition.

If information similar to the following is returned:

```
Command (m for help): p
Disk /dev/vdb: 64.4 GB, 64424509440 bytes, 125829120 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: 0xcb7892a8

   Device Boot      Start         End      Blocks   Id  System
/dev/vdb1           2048     125829119     62913536   83   Linux
Command (m for help):
```

it indicates the detailed information of the newly created partition `/dev/vdb1`.

Note

If the partition operation above is incorrect, enter **q** to exit the fdisk partition tool, and the previous partition results will not be retained.

10. Enter **w** and press **Enter** to write the partition results to the partition table.

The following returned information indicates that the partition has been created.

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

Calling ioctl() to re-read partition table.
Syncing disks.
[root@UM_16_14_centos ~]#
```

11. Execute the following command to synchronize the new partition table changes to the operating system.

```
partprobe
```

12. Execute the following command to set the file system of the newly created partition to the format required by the system.

```
mkfs -t <file system format> /dev/vdb1
```

Different file systems support different partition sizes. Choose a file system that suits your needs. For instance, when setting up the EXT4 file system:

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

The echo message is displayed as follows:

```

[root@UM_16_14_centos ~]# 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
3932160 inodes, 15728384 blocks
786419 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2164260864
480 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000, 7962624, 11239424

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

```

Formatting requires some time. Please monitor the system's running status and do not exit.

13. Run the following command to create a new mount point.

```
mkdir <mount point>
```

Taking the newly created mount point `/data/newpart` as an example:

```
mkdir /data/newpart
```

14. Run the following command to mount the new partition to the newly created mount point.

```
mount /dev/vdb1 <mount point>
```

For instance, mounting a newly created partition to the mount point `/data/newpart` :

```
mount /dev/vdb1 /data/newpart
```

15. Run the following command to view the mount result.

```
df -TH
```

The following example of returned information:

```

[root@UM_16_14_centos ~]# mkdir /data/newpart
[root@UM_16_14_centos ~]# mount /dev/vdb1 /data/newpart
[root@UM_16_14_centos ~]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/vda1      ext4      53G   1.6G   49G   4% /
devtmpfs       devtmpfs  510M     0   510M   0% /dev
tmpfs          tmpfs     520M   25k   520M   1% /dev/shm
tmpfs          tmpfs     520M  492k   520M   1% /run
tmpfs          tmpfs     520M     0   520M   0% /sys/fs/cgroup
tmpfs          tmpfs     104M     0   104M   0% /run/user/0
/dev/vdb1      ext4      64G   55M   60G   1% /data/newpart

```

indicates that the newly created partition `/dev/vdb1` has been mounted to `/data/newpart` .

Note

If you do not need to configure disk automount at startup, skip the following steps.

16. Confirm the mounting method and obtain the corresponding information.

Based on business needs, you can use an elastic cloud disk's soft link, file system's UUID (universally unique identifier), or device name to automatically mount a disk. The descriptions and information acquisition methods are as follows:

Mount method	Pros and cons	Obtaining information
Utilizing the soft link of an elastic cloud disk (recommended)	Advantages: Each elastic cloud disk has a fixed and unique soft link, which will not change with operations such as unmounting, mounting, formatting partitions, etc. Limitations: Only elastic cloud disks support soft links. It is unable to detect the formatting operations of partitions.	Run the following command to view the soft link of the elastic cloud disk. <pre>ls -l /dev/disk/by-id</pre>
Use the UUID of the file system	Automatic mounting configuration may fail due to changes in a file system's UUID. For example, reformatting a file system will change its UUID.	Run the following command to view the UUID of the file system. <pre>blkid /dev/vdb1</pre>
Use device name	Auto-attaching configuration may fail due to changes in device name. For instance, if an elastic cloud disk on the cloud server is unmounted and then remounted, the device name may change when the operating system recognizes the device again.	Run the following command to view the device name. <pre>fdisk -l</pre>

17. Execute the following command to backup the `/etc/fstab` file . The example below demonstrates how to backup to the `/home` directory:

```
cp -r /etc/fstab /home
```

18. Run the following command to open the `/etc/fstab` file using the VI editor.

```
vi /etc/fstab
```

19. Press `i` to enter the editing mode.

20. Move the cursor to the end of the file, press **Enter**, and append the following content.

```
<Device information> <Mount point> <File system format> <File system installation option> <File system dump frequency> <File system check sequence at launch>
```

- **(Recommended)** Using the automatic mounting of the soft link of an elastic cloud disk as an example, combine with the previous examples to add:

```
/dev/disk/by-id/virtio-disk-drkhlpe /data/newpart ext4 defaults,nofail 0 2
```

Note

If you have multiple elastic cloud disks, you can differentiate them by comparing the `disk-xxxxx` with the cloud disk ID in the [console](#).

- Take automatic mounting using the UUID of the disk partition as an example. Add the following content:

```
UUID=d489ca1c-5057-4536-81cb-ceb2847f9954 /data/newpart ext4 defaults,nofail 0 2
```

- Take automatic mounting using the device name as an example. Add the following content:

```
/dev/vdb1 /data/newpart ext4 defaults,nofail 0 2
```

21. Press **Esc**, input `:wq`, and press **Enter** to save the settings and exit the editor.

22. Run the following command to verify if the `/etc/fstab` file has been written successfully.

```
mount -a
```

If the command runs successfully, the file has been written. The newly created file system will automatically mount when the operating system starts up.

Related Actions

[Initializing cloud disks \(≥ 2 TB\)](#)

Initializing Cloud Disks (≥ 2 TB)

Last updated: 2023-09-25 17:23:13

Scenario

This document provides guidance on initializing cloud disks with a capacity of 2TB or more. For more information on cloud disk initialization scenarios, refer to [Introduction to Initialization Scenarios](#).

The maximum disk capacity supported by MBR is 2TB, so when partitioning a disk with a capacity greater than 2TB, please use the GPT partition format. For the Linux operating system, when the GPT partition format is chosen, the fdisk partition tool cannot be used and the parted tool must be utilized instead.

Preparations

The [cloud disk](#) has been attached to the cloud server.

Supports and Limits

- Before performing any operations on the cloud disk, it is advisable to familiarize yourself with the [Precautions for Using Cloud Disks](#) to prevent damage to important data.
- Formatting a data disk will erase all data. Ensure that the disk is either empty or that important data has been backed up.
- To avoid service exceptions, ensure before formatting that the CVM has stopped external services.

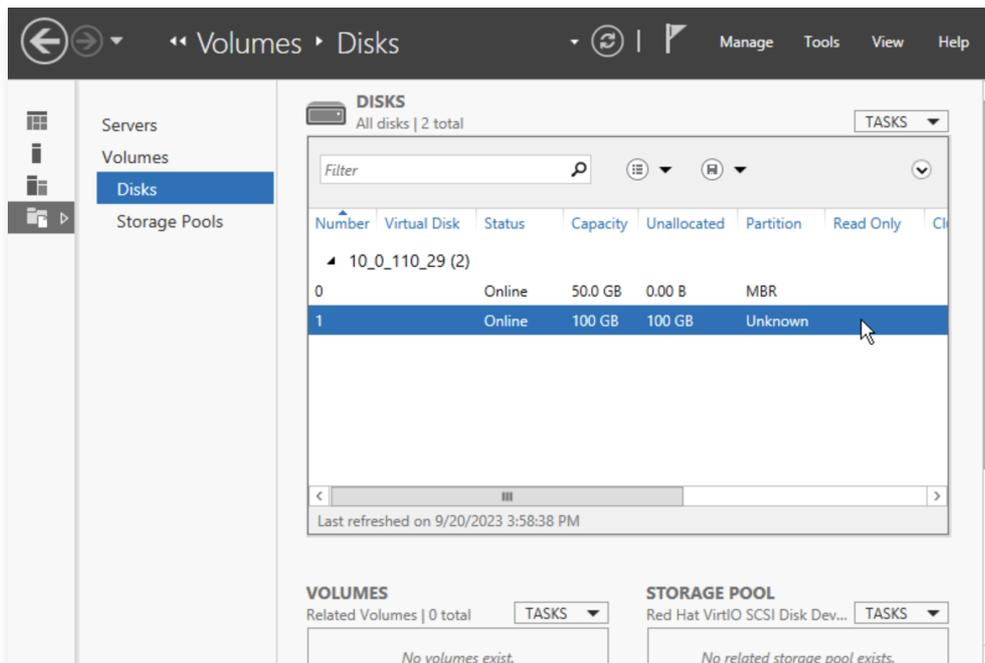
Instructions

Initializing Cloud Disk (Windows)

Note

This document uses a CVM with Windows Server 2012 installed as an example. Note that the steps may vary according to the operating system version.

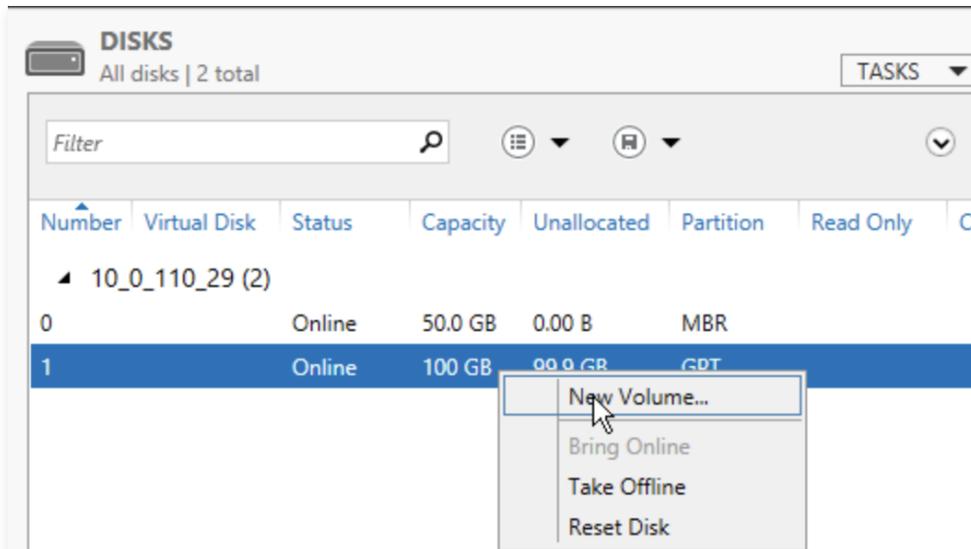
1. [Log in to the Windows CVM](#).
2. On the cloud server desktop, click .
3. Navigate to the **Server Manager** page and click on **Files and Storage Services** in the left navigation tree.
4. In the left navigation tree, select **Volumes > Disks**.



Note

If the newly added disk is in an offline state, you need to first execute [Step 5](#) to bring it online before executing [Step 6](#) for initialization. Otherwise, proceed directly to [Step 6](#) for initialization.

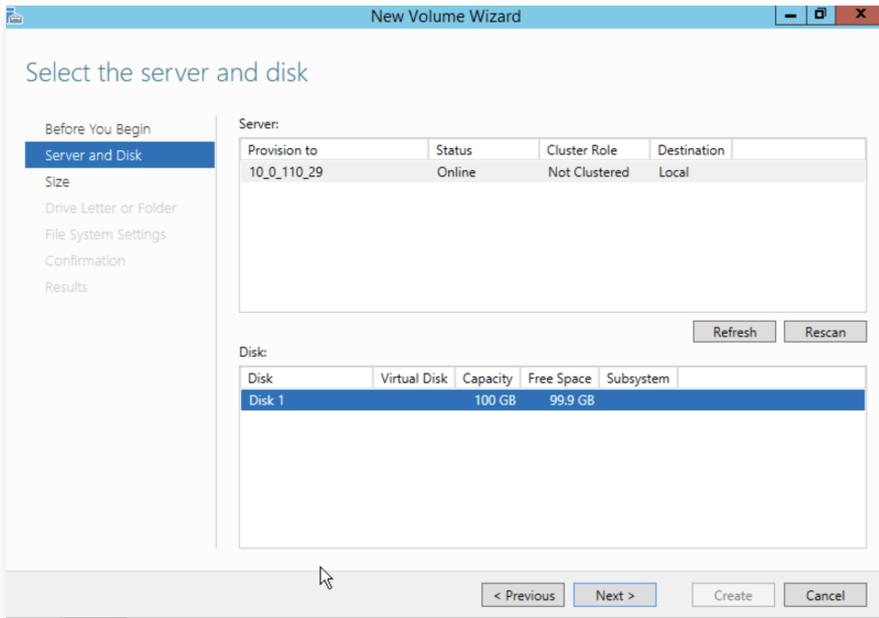
- In the right pane, a list of disks will appear. If the status of the disk is **Offline**, right-click on Disk 1 area and select **Online** from the menu list to bring it online. After going online, Disk changes from **Offline** status to **Not Initialized**.
- Right-click on the Disk area and select **New Volume...** from the drop-down menu, as shown in the figure below:



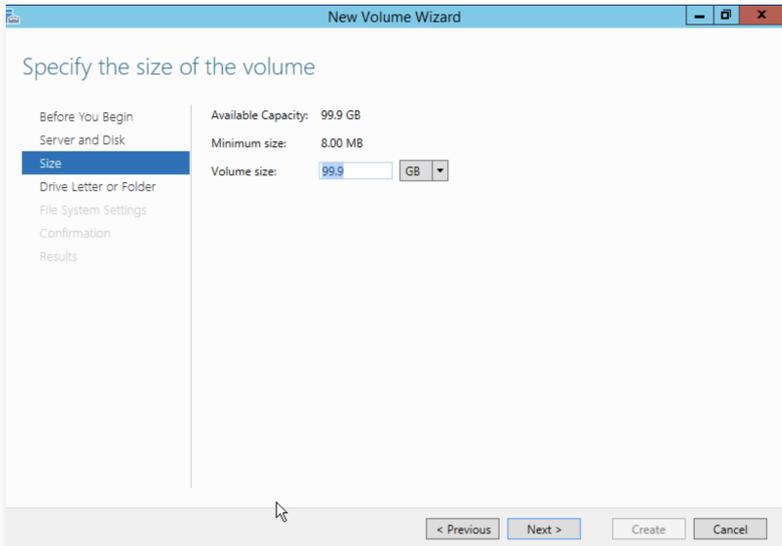
- In the **New Volume...** dialog box, select the disk that needs to be initialized, the confirm dialog appears and click **Confirm**. As shown below:

Note

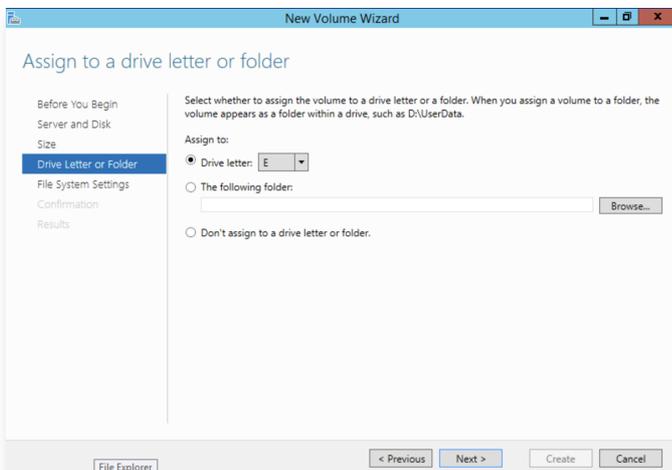
If a disk initializes after the disk is put into use, the original data on the disk will be erased.



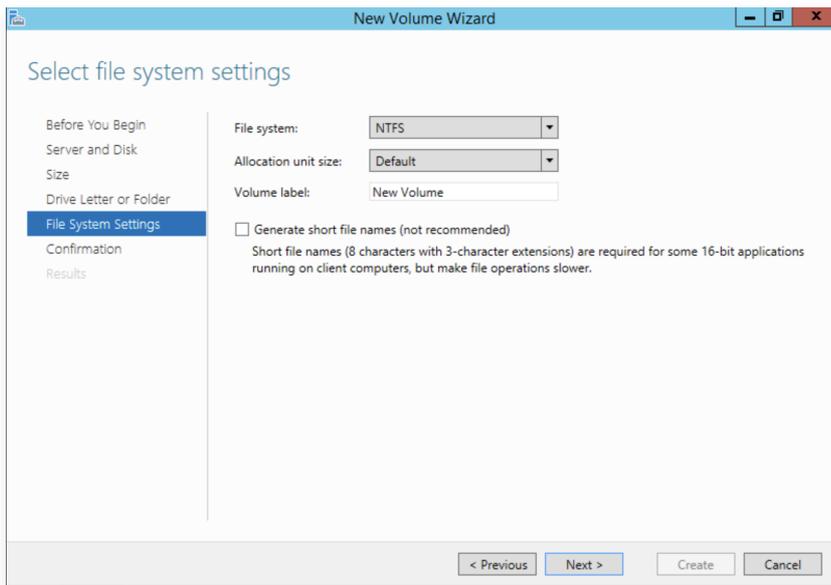
8. Specify the volume size according to your needs. By default, the maximum value is selected. Click **Next**.



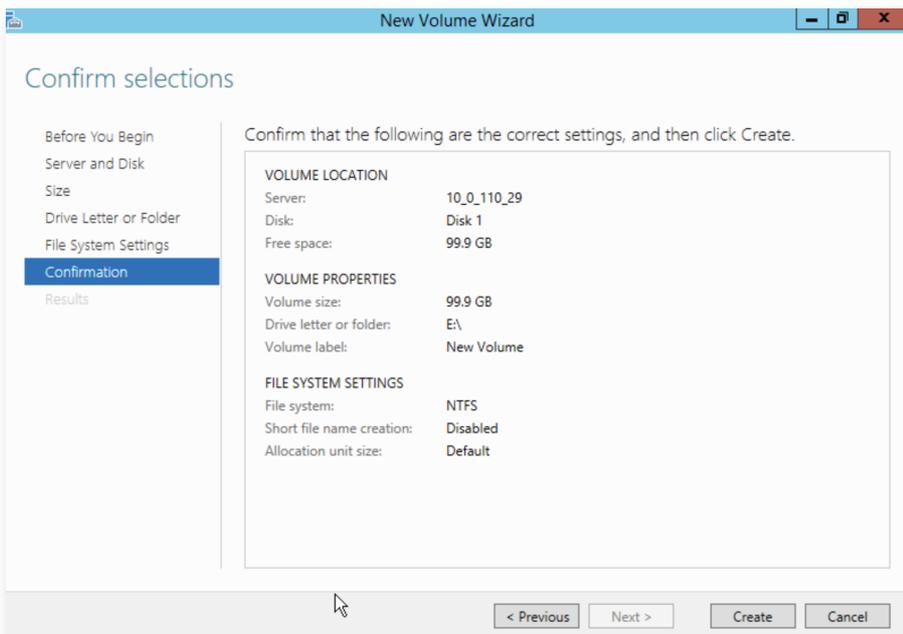
9. Assign a drive letter and click **Next**. As shown in the figure below:



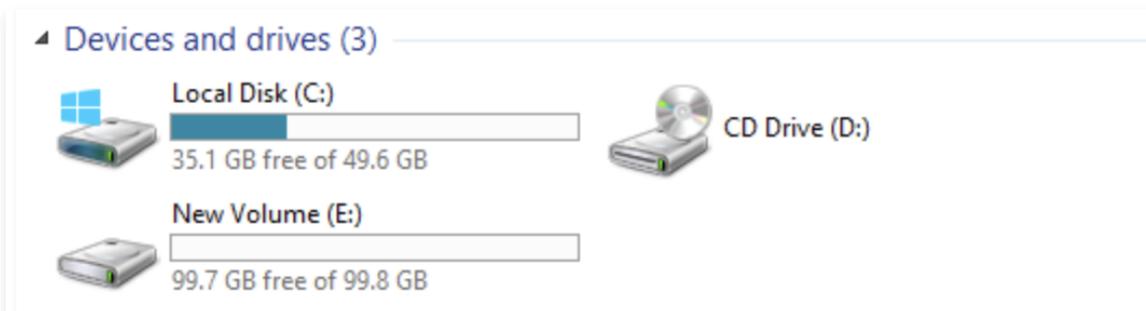
10. Select File System Setting, set the parameters as needed, take default for example and click Next.



11. Check all information and click **Create**. Please wait a moment for the system to complete the initialization. When the volume status is **Healthy**, it indicates that the disk initialization was successful.



After successful initialization, you can view the new disk in the **Computer** interface.



Initializing Cloud Disks (Linux)

Select the initialization method according to your actual use cases:

- If the entire disk is presented as a single independent partition (i.e., there are no multiple logical disks such as vdb1 and vdb2), we strongly recommend that you do not use partitions and directly [build a file system on the raw device](#).
- If the entire disk needs to be represented as multiple logical partitions (i.e., there are multiple logical drives), you must first perform partitioning operations, then [create a file system on the partition](#).

Creating file systems on bare devices

1. [Log into the Linux CVM](#).
2. Run the following command as the root user to view the disk name.

```
fdisk -l
```

If information similar to what is shown below is returned, the current CVM has two disks, where "/dev/vda" is the system disk and "/dev/vdb" is the newly added data disk.

```
root@UM_16_14_centos ~]# fdisk -l
Disk /dev/vda: 53.7 GB, 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: 0x000d64b4

   Device Boot      Start         End      Blocks   Id  System
/dev/vda1  *           2048     104857599     52427776   83  Linux

Disk /dev/vdb: 64.4 GB, 64424509440 bytes, 125829120 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
```

3. Run the following command to create a file system on the `/dev/vdb` bare device.

```
mkfs -t <file system format> /dev/vdb
```

Different file systems support different partition sizes. Please choose the file system according to your actual needs. For instance, when setting the file system to `EXT4` :

```
mkfs -t ext4 /dev/vdb
```

Note

The formatting takes a while. Pay attention to the system's running status and do not exit.

4. Run the following command to create a new mount point.

```
mkdir <mount point>
```

Taking the newly created mount point `/data` as an example:

```
mkdir /data
```

5. Run the following command to mount the new partition to the newly created mount point.

```
mount /dev/vdb <mount point>
```

Taking the newly created mount point `/data` as an example:

```
mount /dev/vdb /data
```

6. Run the following command to view the mount result.

```
df -TH
```

Note

If you do not need to configure disk automount at startup, skip the following steps.

7. Confirm the mounting method and obtain the corresponding information.

Based on business needs, you can use an elastic cloud disk's soft link, file system's UUID (universally unique identifier), or device name to automatically mount a disk. The descriptions and information acquisition methods are as follows:

Mount method	Pros and cons	Obtaining information
Utilizing the soft link of an elastic cloud disk (recommended)	<p>Advantages: Each elastic cloud disk has a fixed and unique soft link that does not change with operations such as unmounting, mounting, formatting partitions, etc.</p> <p>Limitations: Only elastic cloud disks support soft links. It is unable to detect the formatting operations of partitions.</p>	<p>Run the following command to view the soft link of the elastic cloud disk.</p> <pre>ls -l /dev/disk/by-id</pre>
Use the UUID of the file system	<p>Automatic mounting configuration may fail due to changes in a file system's UUID. For example, reformatting a file system will change its UUID.</p>	<p>Run the following command to view the UUID of the file system.</p> <pre>blkid /dev/vdb</pre>
Use device name	<p>Auto-attaching configuration may fail due to changes in device name. For example, if an elastic cloud disk on the CVM is unmounted and then remounted, the device name may change when the operating system recognizes the file system again.</p>	<p>Run the following command to view the device name.</p> <pre>fdisk -l</pre>

8. Run the following command to back up the `/etc/fstab` file. The example below demonstrates how to back it up to the `/home` directory:

```
cp -r /etc/fstab /home
```

9. Run the following command to open the `/etc/fstab` file using the VI editor.

```
vi /etc/fstab
```

10. Press `i` to enter the editing mode.

11. Move the cursor to the end of the file, press **Enter**, and append the following content.

```
<Device information> <Mount point> <File system format> <File system installation option> <File system
dump frequency> <File system check sequence at launch>
```

- **(Recommended)** Using the automatic mounting of the soft link of an elastic cloud disk as an example, combine with the previous examples to add:

```
/dev/disk/by-id/virtio-disk-drkhk|pe /data ext4 defaults,nofail 0 0
```

- Take automatic mounting using the UUID of the disk partition as an example. Add the following content:

```
UUID=d489ca1c-5057-4536-81cb-ceb2847f9954 /data ext4 defaults,nofail 0 0
```

- Take automatic mounting using the device name as an example. Add the following content:

```
/dev/vdb /data ext4 defaults,nofail 0 0
```

12. Press **ESC**, enter **:wq**, and press **Enter** to save the configuration and exit the editor.

13. Run the following command to verify if the `/etc/fstab` file has been written successfully.

```
mount -a
```

If the command runs successfully, the file has been written. The newly created file system will automatically mount when the operating system starts up.

Creating a file system on a partition

Note

This guide will demonstrate how to use the parted partition tool in CentOS 7.5 operating system to set the data disk `/dev/vdc` as the primary partition. The partition format is set to GPT by default, the file system is set to EXT4 format, it is mounted under `/data/newpart2`, and is set to automatically mount at startup. The formatting operations may vary across different operating systems, so this guide should be used for reference only.

1. [Log into the Linux CVM](#)
2. Run the following command as the root user to view the disk name.

```
lsblk
```

If information similar to what is shown below is returned, the current CVM has two disks, where `/dev/vda` is the system disk and `/dev/vdc` is the newly added data disk.

```
[root@VM_16_14_centos ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sr0 11:0 1 37M 0 rom
vda 253:0 0 50G 0 disk
└─vda1 253:1 0 50G 0 part /
vdc 253:32 0 3T 0 disk
```

3. Run the following command to use the parted tool to partition the newly added data disk.

```
parted <new data disk>
```

The newly mounted data disk `/dev/vdc` is used as an example:

```
parted /dev/vdc
```

The returned information is similar to what is shown below:

```
[root@UM_16_14_centos ~]# parted /dev/vdc
GNU Parted 3.1
Using /dev/vdc
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) _
```

4. Enter `p` and press **Enter** to view the current disk partition format. The echo information should resemble the

```
(parted) p
Error: /dev/vdc: unrecognised disk label
Model: Virtio Block Device (virtblk)
Disk /dev/vdc: 3221GB
Sector size (logical/physical): 512B/512B
Partition Table: unknown
Disk Flags:
(parted) _
```

following image:

Partition Table: unknown indicates that the disk partition format is unknown.

5. Run the following command to configure the disk partition format.

```
mklabel <disk partition format>
```

If the disk capacity is larger than or equal to 2TB, only GPT partition format can be used:

```
mklabel gpt
```

6. Enter `p` and press **Enter** to check if the disk partition format has been successfully set.

The echo information should resemble the following figure:

```
(parted) mklabel gpt
(parted) p
Model: Virtio Block Device (virtblk)
Disk /dev/vdc: 3221GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

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

Partition Table: gpt indicates that the disk partition format is GPT.

7. Enter `unit s` and press **Enter** to set the disk's measurement unit to cylinders.
8. For instance, to create a partition for the entire disk, enter `mkpart opt 2048s 100%` and press **Enter**. Here, 2048s represents the starting capacity of the disk, and 100% represents the ending capacity. This is just a reference, and you can plan the number and capacity of disk partitions according to your business needs.
9. Enter `p` and press **Enter** to view the details of the newly created partition. The returned information should look similar to the following:

```
(parted) unit s
(parted) mkpart opt 2048s 100%
(parted) p
Model: Virtio Block Device (virtblk)
Disk /dev/vdc: 6291456000s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start  End              Size          File system  Name  Flags
  1      2048s  6291453951s    6291451904s                opt

(parted) _
```

This indicates the details of the newly created partition `/dev/vdc1`.

10. Enter `q` and press **Enter** to exit the parted partition tool.
11. Run the following command to view the disk name.

```
lsblk
```

If information similar to what is shown below is returned, you can now see the new partition `/dev/vdc1`.

```
[root@UM_16_14_centos ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sr0   11:0    1  37M  0 rom
vda   253:0    0  50G  0 disk
├─vda1 253:1    0  50G  0 part /
└─vdc   253:32   0   3T  0 disk
   └─vdc1 253:33   0   3T  0 part
```

12. Execute the following command to set the file system of the newly created partition to the format required by the system.

```
mkfs -t <file system format> /dev/vdc1
```

Different file systems support different partition sizes. Please choose the file system according to your actual needs. The following example uses the EXT4 file system:

```
mkfs -t ext4 /dev/vdc1
```

The returned information is similar to what is shown below:

```
[root@UM_16_14_centos ~]# mkfs -t ext4 /dev/vdc1
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
196608000 inodes, 786431488 blocks
39321574 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2933915648
24000 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 while. Pay attention to the system's running status and do not exit.

13. Run the following command to create a new mount point.

```
mkdir <mount point>
```

Taking the newly created mount point `/data/newpart2` as an example:

```
mkdir /data/newpart2
```

14. Run the following command to mount the new partition to the newly created mount point.

```
mount /dev/vdc1 <mount point>
```

Taking the newly created mount point `/data/newpart2` as an example:

```
mount /dev/vdc1 /data/newpart2
```

15. Run the following command to view the mount result.

```
df -TH
```

If information similar to what is shown below is returned:

```

[root@UM_16_14_centos ~]# mkdir /data/newpart2
[root@UM_16_14_centos ~]# mount /dev/vdc1 /data/newpart2
[root@UM_16_14_centos ~]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/vda1       ext4      53G   1.6G  49G   4% /
devtmpfs        devtmpfs  510M   0     510M  0% /dev
tmpfs           tmpfs     520M   25k   520M  1% /dev/shm
tmpfs           tmpfs     520M  459k   520M  1% /run
tmpfs           tmpfs     520M   0     520M  0% /sys/fs/cgroup
tmpfs           tmpfs     104M   0     104M  0% /run/user/0
/dev/vdc1       ext4      3.2T   93M   3.1T  1% /data/newpart2

```

It indicates that the newly created partition `/dev/vdc1` has been mounted to `/data/newpart2`.

Note

If you do not need to configure disk automount at startup, skip the following steps.

16. Confirm the mounting method and obtain the corresponding information.

Based on business needs, you can use an elastic cloud disk's soft link, file system's UUID (universally unique identifier), or device name to automatically mount a disk. The descriptions and information acquisition methods are as follows:

Mount method	Pros and cons	Obtaining information
Utilizing the soft link of an elastic cloud disk (recommended)	<p>Advantages: Each elastic cloud disk has a fixed and unique soft link, which will not change with operations such as unmounting, mounting, formatting partitions, etc.</p> <p>Limitations: Only elastic cloud disks support soft links. It is unable to detect the formatting operations of partitions.</p>	<p>Run the following command to view the soft link of the elastic cloud disk.</p> <pre>ls -l /dev/disk/by-id</pre>
Use the UUID of the file system	Automatic mounting configuration may fail due to changes in a file system's UUID.	Run the following command to view the

	For example, reformatting a file system will change its UUID.	UUID of the file system. <pre>blkid /dev/vdc1</pre>
Use device name	Auto-attaching configuration may fail due to changes in device name. For instance, if an elastic cloud disk on the cloud server is unmounted and then remounted, the device name may change when the operating system recognizes the device again.	Run the following command to view the device name. <pre>fdisk -l</pre>

17. Run the following command to backup the `/etc/fstab` file. The example below demonstrates how to backup to the `/home` directory:

```
cp -r /etc/fstab /home
```

18. Run the following command to open the `/etc/fstab` file using the VI editor.

```
vi /etc/fstab
```

19. Press `i` to enter the editing mode.

20. Move the cursor to the end of the file, press **Enter**, and append the following content.

```
<Device information> <Mount point> <File system format> <File system installation option> <File system  
dump frequency> <File system check sequence at launch>
```

- **(Recommended)** Take automatic mounting using the soft link of an elastic cloud disk as an example. Enter the device information `/dev/disk/by-id/virtio-disk-xxxxxxx`, where `disk-xxxxxxx` is the ID of the cloud disk you need to set for automatic mounting, which can be viewed on the cloud disk [console](#).

```
/dev/disk/by-id/virtio-disk-bm42ztpm /data/newpart2 ext4 defaults,nofail 0 2
```

- Take automatic mounting using the UUID of the disk partition as an example. Add the following content:

```
UUID=fc3f42cc-2093-49c7-b4fd-c616ba6165f4 /data/newpart2 ext4 defaults,nofail 0 2
```

- Take automatic mounting using the device name as an example. Add the following content:

```
/dev/vdc1 /data/newpart2 ext4 defaults,nofail 0 2
```

21. Press **ESC**, enter `:wq`, and press **Enter** to save the configuration and exit the editor.

22. Run the following command to verify if the `/etc/fstab` file has been written successfully.

```
mount -a
```

If the command runs successfully, the file has been written. The newly created file system will automatically mount when the operating system starts up.

Related Actions

[Initialize Cloud Disks \(less than 2 TB\)](#)

Expanding cloud disks

Overview

Last updated: 2023-09-22 14:44:54

Expanding Cloud System Disks

When expanding a cloud system disk, you have three methods at your disposal.

- [Expand via the Cloud Server Console](#)
- [Expand via the Cloud Disk Console](#)
- [Expand via API](#)

Expanding Cloud Data Disks

To expand a data disk, the following three methods are provided.

- [Expand via the Cloud Server Console](#)
- [Expand via the Cloud Disk Console](#)
- [Expand via API](#)

Depending on the **detachable** attribute of the CBS data disk, you can choose different operation entrances to expand the data disk.

- If the current disk is a **detachable** CBS data disk, you can expand it through the Cloud Disk Console or by using [Expanding Cloud Disks](#).
- If the current disk is a **non-detachable** CBS data disk, you can expand it through the Cloud Server Instance Console or by using [Expanding Cloud Disks](#).

Note

If the maximum capacity of the cloud disk still cannot meet your business needs, you can opt to either [build a RAID group with multiple elastic cloud disks](#) or [construct LVM logical volumes with multiple elastic cloud disks](#).

Once the data disk capacity is expanded, you must perform the following operations for the instance to recognize and use the data disk:

Before Expansion	After Expansion	See Also
File system not created	Disk capacity is less than 2 TB.	Initialize Cloud Disks (less than 2 TB)
	Disk capacity \geq 2 TB	Initializing cloud disks (\geq 2 TB)
File system created	Disk capacity is less than 2 TB.	<ul style="list-style-type: none"> • If you are expanding a cloud disk on a Windows cloud server: Extending Partitions and File Systems (Windows) • If you are expanding a cloud disk on a Linux cloud server: Extending Partitions and File Systems (Linux)
	Disk capacity \geq 2 TB	<ul style="list-style-type: none"> • For GPT partition format: Extend Partition and File System (Windows) or Extend Partition and File System (Linux). • Adopting MBR partition format: Not supported. The maximum disk capacity supported by MBR partition format is 2TB. If your disk partition is in MBR format and needs to be expanded beyond 2TB, it is recommended to create and mount a new data disk, then copy the data to the new disk after using the GPT partition format.

Billing Description

Expanded capacity fees will be charged after cloud disk expansion. For cloud disks in different billing modes:

- **Monthly Subscribed Cloud Disks:** The expansion capacity is priced based on the remaining time of the lifecycle, making up the difference between the new and old configurations. The specific situation depends on the actual circumstances, which can be viewed on the payment page.
- **Pay-as-you-go cloud disks:** The changes take effect immediately and billing commences based on the newly configured price.
- Billing rules for monthly subscribed cloud disks:
 - The price difference should be made up based on the number of days. Configuration upgrade fees = price difference for upgrade per month *number of months for upgrade* applicable discount.
 - Price difference for the upgrade per month: Unit price difference between the new and old configurations.
 - Number of months for upgrade: Upgrade fees are calculated by month based on the number of days.
 - Number of days for upgrade = resource expiration time – current time
 - Number of months for upgrade = number of days for upgrade / (365/12)
 - Applicable discount: The applicable discount is matched based on the number of months for upgrade as effective at the official website.

Note

- This operation doesn't affect the resource expiration time.
- This operation allows you to use vouchers and free credits for fees deduction.

Billing Example

Note

The following prices are for illustrative purposes only and do not reflect the actual prices on the official website. For actual cloud disk prices, see [Cloud Disk Storage Space Price Overview](#).

A user in the Guangzhou region purchased a 200GB monthly subscription Balanced SSD cloud disk for a duration of one month on May 1, 2022, to be used as a data disk. On May 5, 2022, the disk was expanded to 300GB. The cost calculation is as follows:

Price difference for upgrade per month = 150 – 100 = 50 CNY/month

Number of days for upgrade = 27 days

Applicable discount = No discount available

Final total upgrade cost = 50 x (27 / (365/12)) x 1 = 44.38 CNY

Expanding cloud disks

Last updated: 2023-09-21 17:55:22

Scenario

Cloud Block Storage is a scalable storage device in the cloud, allowing users to expand its size at any time after creation to increase storage space without losing existing data on the disk. After the cloud disk expansion is complete, it is necessary to extend the partition and file system. You can allocate the expanded capacity to an existing partition or format the expanded capacity into a separate new partition.

Note

The maximum disk capacity supported by MBR partition format is 2TB. If your disk partition is in MBR format and needs to be expanded to more than 2TB, it is recommended to create and mount a new data disk, use the GPT partition method, and copy the data to the new disk.

Expanding Data Disks

To expand a data disk, the following three methods are provided.

Note

If your cloud server has multiple cloud disks with the same capacity and type mounted, you can refer to [Distinguishing Data Disks](#) for differentiation. After selecting the data disk to be expanded, proceed with the expansion using the following methods.

Expand via Cloud Server Console (Recommended)

1. Log in to the [CVM console](#).
2. In the row of the target cloud server, select **More > Resource Adjustment > Expand Cloud Disk**.
3. In the **Expand Cloud Disk** pop-up window, select the data disk to be expanded and click **Next**.
4. In the **Adjust Capacity** step, set the target capacity (must be greater than or equal to the current capacity) and click **Next**.
5. In the **Expand Partition and File System** step, review the notes and click **Start Adjusting**.
6. Depending on the target cloud service's operating system type, you need to [extend partitions and file systems \(Windows\)](#) or [extend partitions and file systems \(Linux\)](#) to allocate the expanded capacity to an existing partition, or format the expanded capacity into a separate new partition.

Expand via Cloud Block Storage Console

1. Log in to the [Cloud Block Storage Console](#).
2. Select **More > Expand** for the target cloud disk.
3. Select a new capacity. It must be greater than or equal to the current capacity.
4. Complete the payment.
5. Depending on the operating system type of the target cloud service, you need to execute [Extending Partitions and File Systems \(Windows\)](#) or [Extending Partitions and File Systems \(Linux\)](#) to allocate the expanded capacity to an existing partition, or format the expanded capacity into a new independent partition.

Expand via API

You can use the `ResizeDisk` interface to expand the specified elastic cloud disk. For detailed directions, see [Expanding Cloud Disks \(ResizeDisk\)](#).

Expanding System Disks

1. Log in to the [CVM console](#). In the row of the target CVM, select **More > Resource Adjustment > Expand Cloud Disk**.
2. In the **Expand Cloud Disk** pop-up window, select the system disk to be expanded and click **Next**.
3. In the **Adjust Capacity** step, set the target capacity (must be greater than or equal to the current capacity) and click **Next**.
4. Expand the cloud disk as instructed below.

Expand via Cloud Server Console (Recommended)

Note

Cloud servers support online expansion of cloud disks used as system disks, that is, expansion without downtime. To use this feature, please [submit an application](#). Once approved, you can start using it.

1. In the **Expand Partition and File System** step, review the notes and click **Start Adjusting**.
2. After completing the expansion operation in the console, please log in to the instance to confirm whether the file system has been automatically expanded. If not, please refer to [Online Expansion of System Disk and File System](#) for partition and file system expansion operations.

Expand via Cloud Block Storage Console

1. Log in to the [Cloud Block Storage Console](#).
2. Select **More > Expand** for the target cloud disk.
3. Select a new capacity. It must be greater than or equal to the current capacity.
4. Complete the payment.
5. Depending on the operating system type of the target cloud service, you need to execute [Extending Partitions and File Systems \(Windows\)](#) or [Extending Partitions and File Systems \(Linux\)](#) to allocate the expanded capacity to an existing partition, or format the expanded capacity into a new independent partition.

Expand via API

You can use the `ResizeInstanceDisks` API to expand a specified non-elastic cloud disk. For detailed instructions, please refer to [Expanding Instance Disks \(ResizeInstanceDisks\)](#).

Related Actions

Distinguishing data disks

Check cloud disks according to the operating system of the CVM.

Linux

1. [Log in to the Linux Instance](#).
2. Run the following command to view the relationship between the elastic cloud disks and the device name.

```
ls -l /dev/disk/by-id
```

The result is shown in the following figure:

```
[root@VM_63_126_centos ~]# ls -l /dev/disk/by-id/
total 0
lrwxrwxrwx 1 root root 9 Mar  1 17:31 virtio-disk-35t32l8g -> ../../vdf
lrwxrwxrwx 1 root root 9 Mar  1 17:31 virtio-disk-jel3nl0g -> ../../vdc
lrwxrwxrwx 1 root root 9 Mar  1 17:31 virtio-disk-jwz43lpg -> ../../vde
lrwxrwxrwx 1 root root 9 Mar  1 17:31 virtio-disk-punhzcju -> ../../vdd
```

Where, `disk-xxxx` is the cloud disk ID, which can be viewed in the [Cloud Disk Console](#).

Windows

1. [Log in to the Windows instance](#).
2. Right-click  and select **Run**.
3. In the **Run** window, enter `cmd` and press **Enter**.
4. Run the following command to view the relationship between the cloud disks and the device name.

```
wmic diskdrive get caption,deviceid,serialnumber
```

Or execute the following command:

```
wmic path win32_physicalmedia get SerialNumber,Tag
```

The result is shown in the following figure:

```
C:\Users\Administrator>wmic diskdrive get caption,deviceid,serialnumber
Caption                DeviceID                SerialNumber
Red Hat VirtIO SCSI Disk Device  \\.\PHYSICALDRIVE0
Red Hat VirtIO SCSI Disk Device  \\.\PHYSICALDRIVE1  disk-hnvcmqrm
```

Where, `disk-xxxx` is the cloud disk ID, which can be viewed in the [Cloud Disk Console](#).

Checking the cloudinit configuration

Check cloud disks according to the operating system of the CVM.

Viewing the Linux Instance's cloudinit Configuration

After completing the expansion, please [log in to the Linux instance](#) to confirm whether `/etc/cloud/cloud.cfg` contains the `growpart` and `resizefs` configuration items.

- If yes, ignore other operations. See below:

```
cloud_init_modules:  
- migrator  
- bootcmd  
- write-files  
- growpart  
- resizefs  
- set_hostname  
- update_hostname  
- ['update_etc_hosts', 'once-per-instance']  
- rsyslog  
- users-groups  
- ssh
```

- **growpart**: Extend partition size to match disk size.
- **resizefs**: Expand and adjust the / partition file system to the partition size.
- Otherwise, you need to manually extend the file system and partition according to the operating system type of the target cloud service. You need to execute [Extending Partitions and File Systems \(Linux\)](#), allocating the expanded capacity to an existing partition or formatting the expanded capacity into a new independent partition.

Viewing the cloudinit configuration for Windows instances

After completing the expansion operation, please [log in to the Windows instance](#) to confirm whether the plugin in `C:\Program Files\Cloudbase Solutions\Cloudbase-Init\conf\cloudbase-init.conf` includes the `ExtendVolumesPlugin` configuration item.

- Yes, if the `cloudbase-init` configuration file `cloudbase-init.conf` includes the `ExtendVolumesPlugin` configuration item, you need to restart the machine for `cloudbase-init` to automatically extend the volume and add the blank space behind the C partition to the C partition. This requires no other partitions to interfere between the C partition and the blank space. If there are no other partitions between the C partition and the blank space and you do not want to restart, or if a third-party security software blocks `cloudbase-init` from completing the extend volume, you need to manually execute the following powershell command.

```
$DiskOps="@  
select disk 0  
select volume c  
extend  
exit  
@"  
$DiskOps | diskpart.exe | Out-Null
```

- Otherwise, you need to manually extend the file system and partition according to the operating system type of the target cloud service. You need to perform [Extending Partitions and File Systems \(Windows\)](#), allocating the expanded capacity to an existing partition or formatting the expanded capacity into a new independent partition.

Extending System Disk Partitions and File Systems Online

Last updated: 2023-09-21 14:20:22

Scenario

After expanding the storage capacity of a cloud disk used as a system disk through the console, you still need to expand the cloud disk partition or file system. This document will guide you on how to expand the partition and file system without interrupting the service.

Preparations

- Before proceeding with the steps outlined in this document, please create a snapshot for the corresponding instance system disk to back up the data. For more information, see [Creating Snapshots](#). In case of data loss due to an error, you can roll back the snapshot to recover the data.
- You have already expanded the cloud disk capacity through the console. For more information, see [Expanding System Disk](#).
- The Linux CVM kernel version should not be lower than 3.6.0. You can use the `uname -a` command to check the kernel version. If the kernel version is lower than 3.6.0, you can refer to [Offline Expansion of MBR Partition and File System](#) for operations.

Operating Environment

Linux Instances

Resources	Note
Operating System	CentOS 8.0 64-bit
Cloud disk (system disk)	/dev/vda: Uses the MBR partition and the EXT4 file system, and expands from 50 GB to 60 GB online in the console.

Windows Instances

Resources	Note
Operating System	Windows Server 2012 R2 Datacenter 64-bit Chinese
Cloud disk (system disk)	Drive C : Utilizing MBR partition and NTFS file system, its capacity has been expanded from 50GB to 100GB online through the console.

Instructions

Please execute the following steps according to the operating system of your instance:

Linux Instances

1. Log in to the cloud server. For more information, see [Logging in to Linux Instance Using Standard Login Method \(Recommended\)](#).
2. Run the following command to query partitions of the cloud disk.

```
fdisk -l
```

The result returned is as shown below, indicating that the `dev/vda` data disk has a capacity of 60GB, which includes the MBR partition `/dev/vda1` with a capacity of 50GB.

```
[root@VM-5-17-centos ~]# fdisk -l
Disk /dev/vda: 60 GiB, 64424509440 bytes, 125829120 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
Disklabel type: dos
Disk identifier: 0x89ee0607

Device      Boot Start      End  Sectors  Size Id Type
/dev/vda1  *        2048 104857566 104855519   50G 83 Linux
```

3. Run the following command to confirm the file system that has a partition.

```
df -TH
```

The following result shows that the file system type of `/dev/vda1` is ext4.

```
[root@VM-5-17-centos ~]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
devtmpfs        devtmpfs  937M   0  937M   0% /dev
tmpfs           tmpfs     953M  25k  953M   1% /dev/shm
tmpfs           tmpfs     953M  447k  953M   1% /run
tmpfs           tmpfs     953M   0  953M   0% /sys/fs/cgroup
/dev/vda1       ext4      53G   3.2G  48G   7% /
tmpfs           tmpfs     191M   0  191M   0% /run/user/0
```

4. Run the following command to install the growpart tool according to the operating system of the CVM.

- CentOS

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

- Ubuntu or Debian

```
apt-get install -y cloud-guest-utils
```

5. Run the following command to use the growpart tool to expand the partition `/dev/vda1`. In the command, `/dev/vda` and `1` must be separated by a space.

```
growpart /dev/vda 1
```

The following information will appear:

```
[root@VM-5-17-centos ~]# growpart /dev/vda 1
CHANGED: partition=1 start=2048 old: size=104855519 end=104857567 new: size=125827039 end=125829087
```

6. Run the following command to extend the EXT4 file system.

```
resize2fs /dev/vda1
```

The following information will appear:

```
[root@VM-5-17-centos ~]# resize2fs /dev/vda1
resize2fs 1.45.6 (20-Mar-2020)
Filesystem at /dev/vda1 is mounted on /; on-line resizing required
old_desc_blocks = 7, new_desc_blocks = 8
The filesystem on /dev/vda1 is now 15728379 (4k) blocks long.
```

7. Run the following command to view the result.

```
df -TH
```

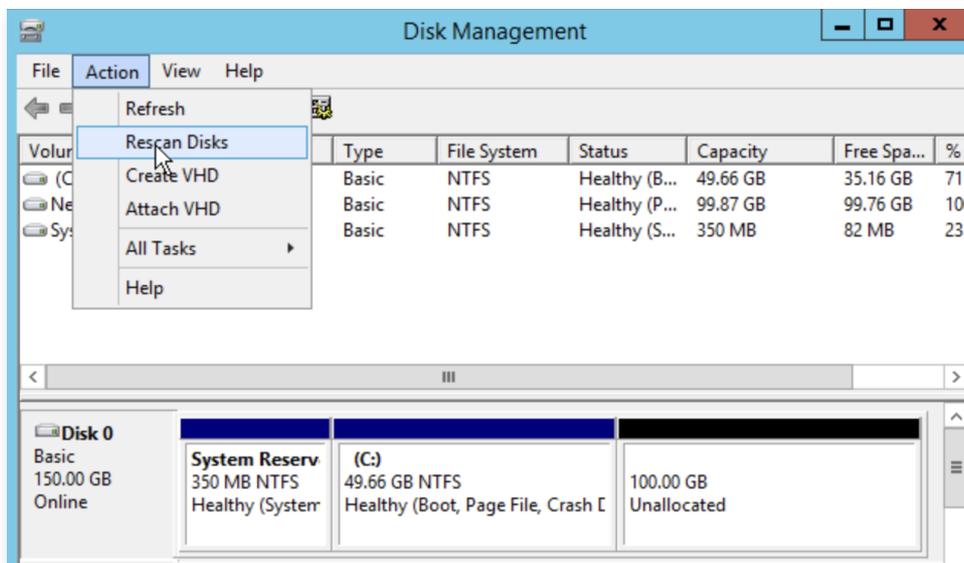
If information similar to what is shown below is returned, the file system has been expanded successfully.

```
[root@VM-5-17-centos ~]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
devtmpfs        devtmpfs  937M   0  937M   0% /dev
tmpfs           tmpfs     953M  25k  953M   1% /dev/shm
tmpfs           tmpfs     953M  447k  953M   1% /run
tmpfs           tmpfs     953M   0  953M   0% /sys/fs/cgroup
/dev/vda1       ext4      64G   3.2G   58G   6% /
tmpfs           tmpfs     191M   0  191M   0% /run/user/0
```

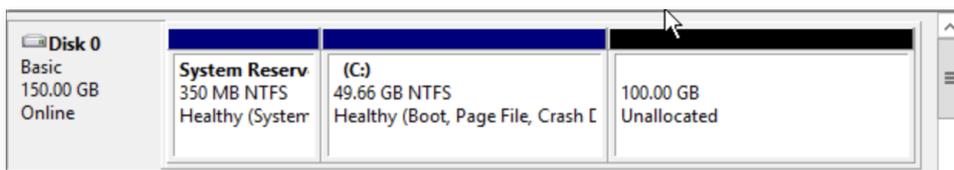
Check data integrity and CVM running status after expansion.

Windows Instances

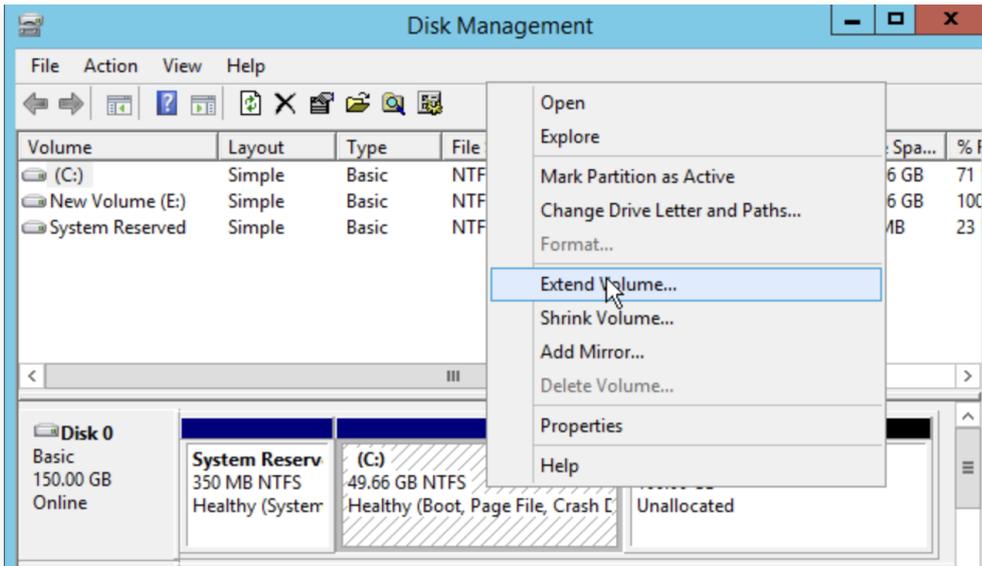
1. Log in to the cloud server. For more information, see [Logging in to Windows Instance Using Standard Method](#).
2. Right-click on the  at the lower left corner of the desktop, search and select **Disk Management**.
3. In the window, choose the expanded disk and click right button select **Action > Rescan Disks** from the top of the interface. See below:



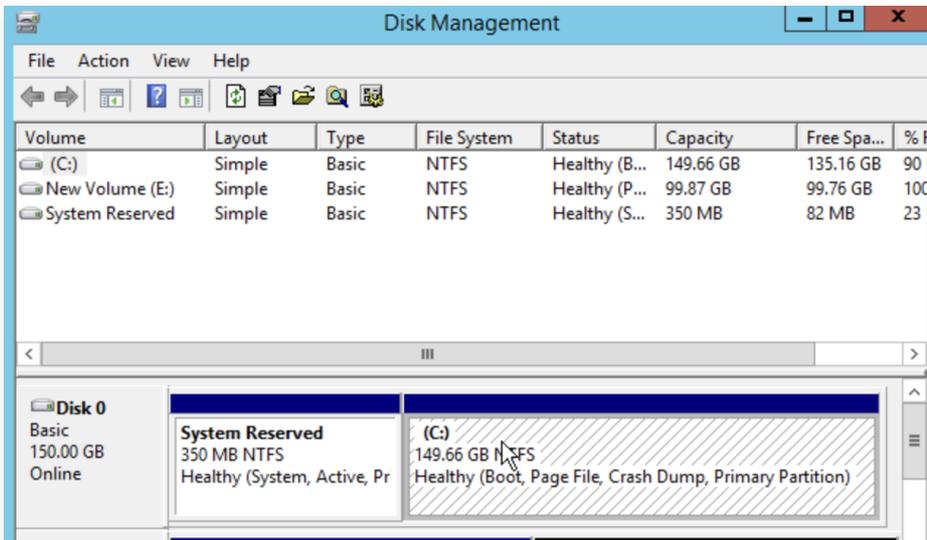
After the scan is completed, you can view the added space. See below:



4. Right-click on the C drive area and select **Extend Volume** from the pop-up menu, as shown in the figure below:



5. Follow the **Expand Volume Wizard** to expand the volume. After the operation, the newly added space will be merged into the original volume. See below:



Check data integrity and CVM running status after expansion.

Extending Data Disk Partitions and File Systems (Windows)

Last updated: 2023-09-21 14:27:23

Scenario

After expanding the cloud disk through the [Cloud Disk Expansion](#) console, you need to allocate the expanded capacity to an existing partition, or format the expanded capacity into a separate new partition.

- If you have expanded the cloud disk while it is attached to a running cloud server, you need to perform the [Rescan Disk](#) operation to recognize the disk space after expansion.
- If the cloud disk to be expanding is not attached to a CVM or the attached CVM is shut down, the disk capacity after expansion will be automatically recognized.

Note

- Improper file system expansion operations may affect existing data. Therefore, we strongly recommend that you manually [create a snapshot](#) to back up data before proceeding.
- Extending the file system requires a [restart of the instance](#) or a rescan of the disk, which will cause a temporary interruption of services. We advise you to choose an appropriate time to proceed with caution.
- After completing the expansion operation, we strongly recommend you follow the [Rescan Disk](#) procedure to recognize the expanded capacity. Performing other operations such as "Refresh" cannot guarantee that the system will recognize the expanded capacity.

Preparations

- You have already [expanded the cloud disk](#) via the console.
- This cloud disk has been attached to a Windows cloud server and a file system has been created through the console. For more details, please refer to [Attaching Cloud Disks](#).
- You have logged into the Windows CVM where the partition and file system to be expanded are located. For more information, see [Logging in to a Windows Instance Using RDP \(Recommended\)](#).

Note

This document describes how to expand a disk attached to a CVM with Windows Server 2012 R2 installed. Note that the steps may vary according to the operating system version.

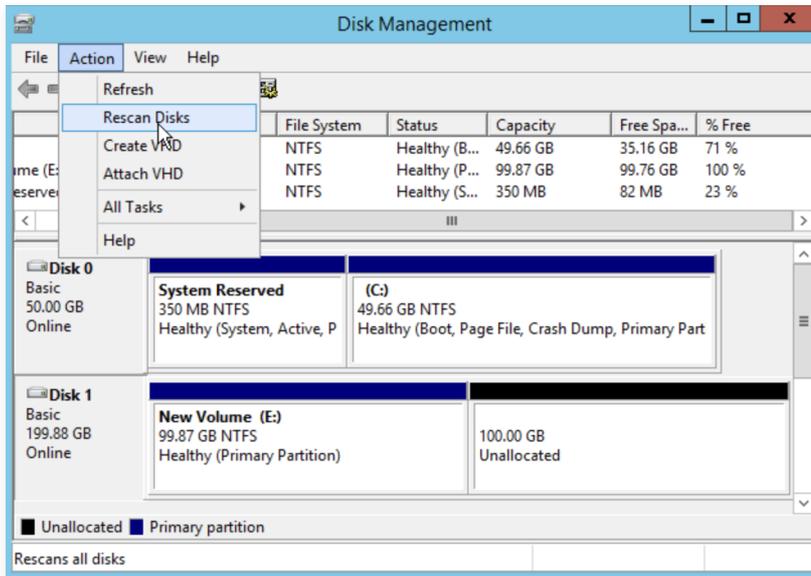
Instructions

Note

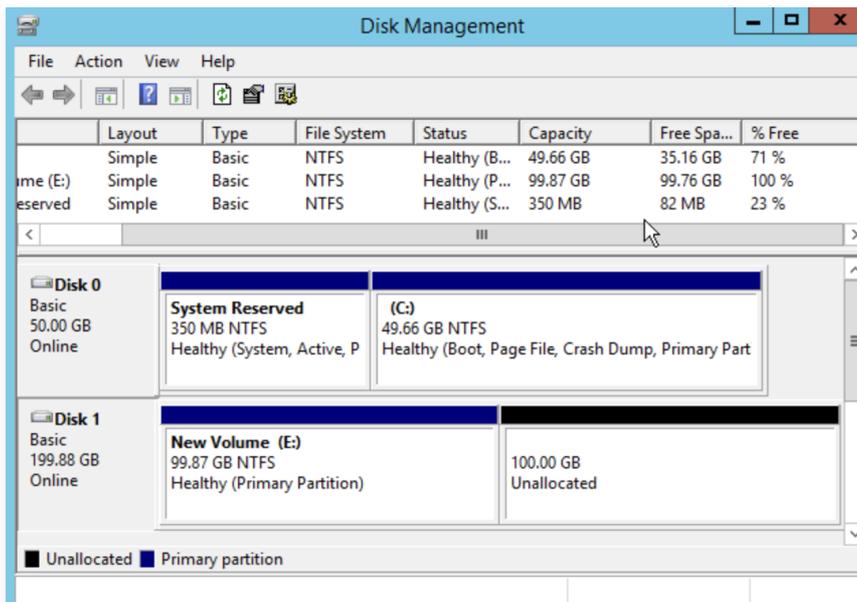
- If the cloud server to which the disk is attached is running normally when expanding the cloud disk through the [Cloud Disk Expansion](#) console, you need to [Rescan Disk](#) to recognize the expanded cloud disk space before [expanding the file system of the original partition or creating a new partition](#).
- If the cloud disk is in a pending mount state or the cloud server to which the disk is attached is shut down when expanding the cloud disk through the [Cloud Disk Expansion](#) console, you can directly [expand the file system of the original partition or create a new partition](#).
- If the Virtio driver version of the cloud server's storage controller is lower than 58003, please [restart the instance](#) before proceeding with the following operations. Refer to [Checking Virtio Driver Version](#) to determine the Virtio driver version in use.

Rescanning the disk

1. Right-click  and select **Computer Management**.
2. In the left navigation bar of the "Computer Management" window, select **Storage > Disk Management**.
3. Right-click on **Disk Management** and select **Rescan Disks**. As shown in the figure below:



4. After the scan is completed, check whether the data disk has expanded to the specified size. (In this example, the scan shows that the cloud disk is expanded from 100 GB to 200 GB). See below:

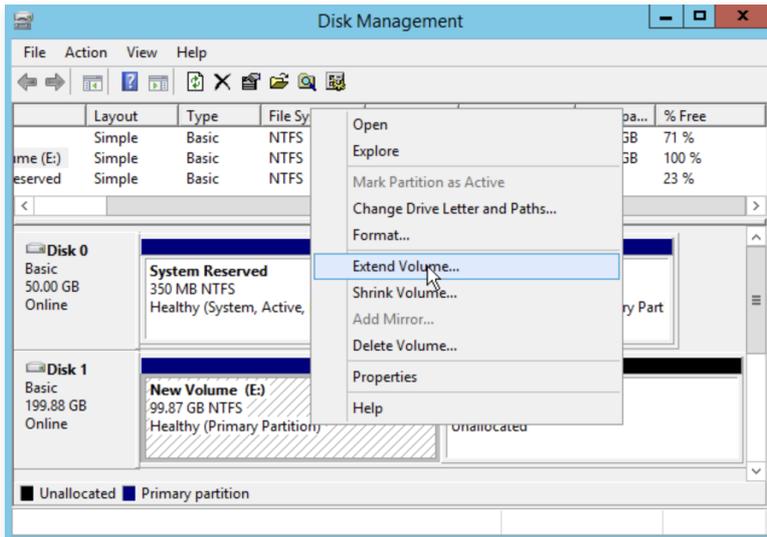


Expanding the File System of an Existing Partition or Creating a New Partition

You can extend the file system of an existing partition or create a partition as instructed below:

Expanding the File System of the Original Partition

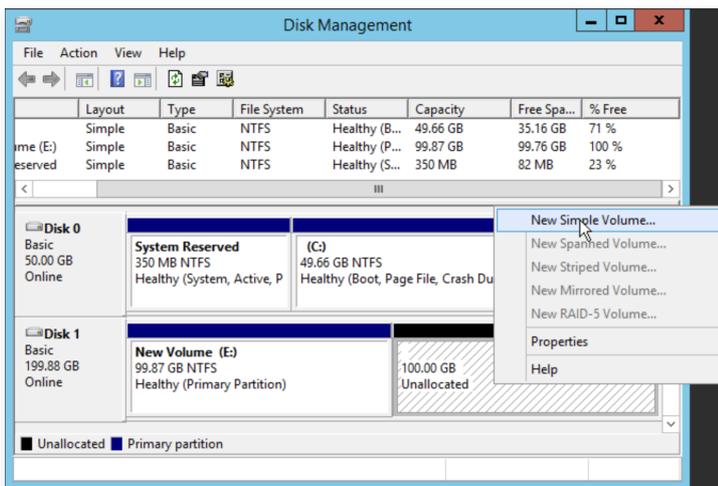
1. Right-click on any blank area of the disk space and select **Extend Volume**, as shown in the figure below:



2. Follow the **Extend Volume Wizard** to extend the volume.
The new data disk capacity will be added to the original volume.

Creating a New Partition

1. Right-click on the unallocated space of the disk and select **New Simple Volume**, as shown in the figure below:

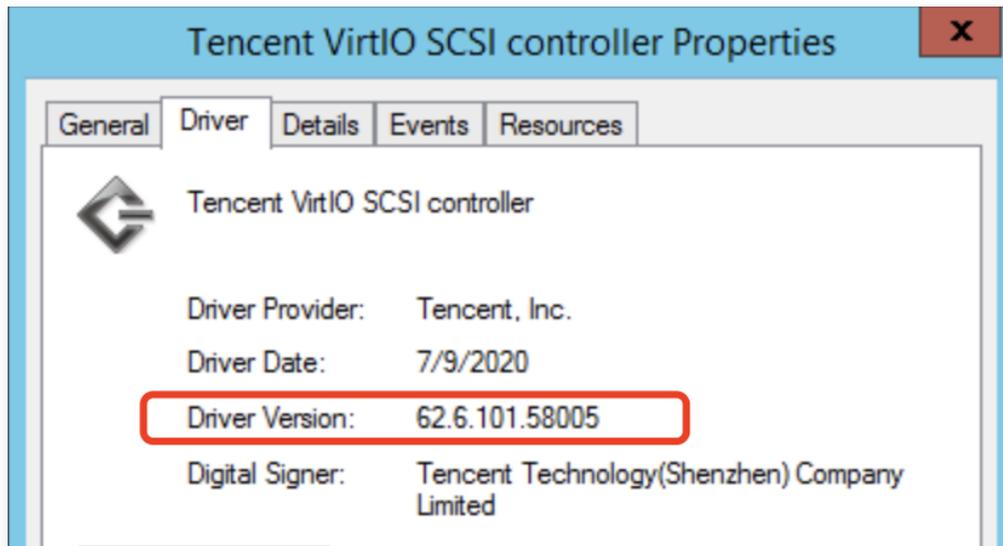


2. Follow the **New Simple Volume Wizard** to create a simple volume with default settings.
The new data disk capacity will be formatted into a new partition.

Related Actions

View Virtio Driver Version

1. Right-click  and select **Device Manager**.
2. In the **Device Manager** window, expand the **Storage Controllers** item, and double-click on **Tencent VirtIO SCSI Controller**.
3. In the **Tencent VirtIO SCSI Controller Properties** window, select **Driver** to view the current version. As shown below, the current version is 58005.



Documentation

- [Scaling up Cloud Disks](#)
- [Extending Partitions and File Systems \(Linux\)](#)

Extending Data Disk Partitions and File Systems (Linux)

Extending Data Disk Partitions and File Systems Online

Last updated: 2023-09-19 20:20:33

Scenario

After expanding the storage capacity of a cloud disk used as a data disk through the console, it is necessary to further expand the partition or file system of the cloud disk. This document provides steps for expanding the partition and file system without downtime. You may refer to this document for operations.

Preparations

- Before proceeding with the steps outlined in this document, please create a snapshot for the corresponding cloud disk to back up the data. For more information, see [Creating Snapshots](#).
In case of data loss due to incorrect operations, you can roll back the snapshot to recover the data.
- The cloud disk capacity has already been expanded through the console, and the disk has been attached to the cloud server. For more information, see [Expanding Cloud Disks](#).
- The Linux CVM kernel version is not lower than 3.6.0. You can use the `uname -a` command to check the kernel version. If the kernel version is lower than 3.6.0, you can refer to [Expanding Partitions and File Systems \(Linux\)](#) for operations.

Operating Environment

Resources	Note
Operating System	CentOS 8.0 64-bit
Cloud disk (data disk)	<ul style="list-style-type: none">• <code>/dev/vdb</code> : Utilizing MBR partition and ext4 file system, the capacity has been expanded from 50GB to 60GB through the console.• <code>/dev/vdc</code> : Utilizing GPT partition and xfs file system, the capacity has been expanded from 50GB to 60GB through the console.

Instructions

Viewing partitions of the cloud disk

1. Log in to the cloud server. For more information, see [Logging in to Linux Instance Using Standard Login Method \(Recommended\)](#).
2. Run the following command to query partitions of the cloud disk.

```
fdisk -l
```

The following information will appear:

```
[root@VM-8-61-centos ~]# fdisk -l
Disk /dev/vda: 20 GiB, 21474836480 bytes, 41943040 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 262144 bytes / 262144 bytes
Disklabel type: dos
Disk identifier: 0x89ee0607

Device      Boot Start          End  Sectors  Size Id Type
/dev/vda1   *      2048 41943006 41940959   20G 83 Linux

Disk /dev/vdb: 60 GiB, 64424509440 bytes, 125829120 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 262144 bytes / 262144 bytes
Disklabel type: dos
Disk identifier: 0x6d871946

Device      Boot Start          End  Sectors  Size Id Type
/dev/vdb1   *      2048 104857599 104855552   50G 83 Linux

Disk /dev/vdc: 60 GiB, 64424509440 bytes, 125829120 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 262144 bytes / 262144 bytes
Disklabel type: gpt
Disk identifier: 8D34962D-3D6F-47E4-A3FF-7F2DC0441FE3

Device      Start          End  Sectors  Size Type
/dev/vdc1   2048 104855551 104853504   50G Linux filesystem
```

As shown in the figure,

- The `/dev/vdb` data disk has a capacity of 60GB, which includes the MBR partition `/dev/vdb1` with a capacity of 50GB.
- The `/dev/vdc` data disk has a capacity of 60GB, which includes a GPT partition `/dev/vdc1` with a capacity of 50GB.

3. Run the following command to confirm the file system type of the existing partition.

```
df -TH
```

The following information will appear:

```
[root@VM-8-61-centos ~]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
devtmpfs        devtmpfs  946M   0  946M   0% /dev
tmpfs           tmpfs     960M  25k  960M   1% /dev/shm
tmpfs           tmpfs     960M 443k  959M   1% /run
tmpfs           tmpfs     960M   0  960M   0% /sys/fs/cgroup
/dev/vda1       ext4      22G   2.1G   19G  11% /
/dev/vdb1       ext4      53G   55M   50G   1% /mnt/disk1
/dev/vdc1       xfs       54G  409M   54G   1% /mnt/disk2
tmpfs           tmpfs     192M   0  192M   0% /run/user/0
```

As shown in the figure,

- The file system type of `/dev/vdb1` is ext4, and it has been mounted to `/mnt/disk1`.
- The file system type of `/dev/vdc1` is xfs, and it has been mounted to `/mnt/disk2`.

Extending a partition

1. Use the command as needed to install the `gdisk` tool.

- For a MBR partition, skip this step.
- For a GPT partition, run the following command according to the operating system of the CVM.

```
CentOS
```

```
yum install gdisk -y
```

Ubuntu or Debian

```
apt-get install gdisk -y
```

2. Run the following command to install the growpart tool according to the operating system of the CVM.

CentOS

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

Ubuntu or Debian

```
apt-get install -y cloud-guest-utils
```

3. Run the following command to use the growpart tool to expand the partition.

This document uses the expansion of the `/dev/vdb1` partition as an example. There should be a space between `/dev/vdb` and `1` in the command. You can modify the command as needed.

```
growpart /dev/vdb 1
```

If information similar to what is shown below is returned, the partition has been extended.

```
[root@VM-8-61-centos ~]# growpart /dev/vdb 1
CHANGED: partition=1 start=2048 old: size=104855552 end=104857600 new: size=1257039,end=125829087
```

Extending the file system

1. Based on the file system type obtained from [Step 3](#), execute the corresponding command to expand the file system:

Expanding the EXT File System

Run the following command to extend the EXT file system.

```
resize2fs /dev/vdb1
```

The following information will appear:

```
[root@VM-8-61-centos ~]# resize2fs /dev/vdb1
resize2fs 1.44.3 (10-July-2018)
Filesystem at /dev/vdb1 is mounted on /mnt/disk1; on-line resizing required
old_desc_blocks = 7, new_desc_blocks = 8
The filesystem on /dev/vdb1 is now 15728379 (4k) blocks long.
```

Extending the XFS File System

Run the following command to extend the XFS file system.

```
xfs_growfs <mount point>
```

In this example, `/dev/vdc1` is mounted to `/mnt/disk2`. Execute the following command.

```
xfs_growfs /mnt/disk2
```

The following information will appear:

```
[root@VM-8-61-centos ~]# xfs_growfs /mnt/disk2
meta-data=/dev/vdc1          isize=512    agcount=4, agsize=3276672 blks
=                               sectsz=512   attr=2, projid32bit=1
=                               crc=1       finobt=1, sparse=1, rmapbt=0
=                               reflink=1
data      =                   bsize=4096  blocks=13106688, imaxpct=25
=                               sunit=0    swidth=0 blks
naming    =version 2          bsize=4096  ascii-ci=0, ftype=1
log       =internal log     bsize=4096  blocks=6399, version=2
=                               sectsz=512  sunit=0 blks, lazy-count=1
realtime  =none             extsz=4096  blocks=0, rtextents=0
```

2. Run the following command to view the result.

```
df -TH
```

If information similar to what is shown below is returned, the file system has been extended.

```
[root@VM-8-61-centos ~]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
devtmpfs        devtmpfs  946M   0  946M   0% /dev
tmpfs           tmpfs     960M  25k  960M   1% /dev/shm
tmpfs           tmpfs     960M  443k  959M   1% /run
tmpfs           tmpfs     960M   0  960M   0% /sys/fs/cgroup
/dev/vda1       ext4      22G   2.1G  19G  11% /
/dev/vdb1       ext4      64G   55M   60G   1% /mnt/disk1
/dev/vdc1       xfs       65G  484M   64G   1% /mnt/disk2
tmpfs           tmpfs     192M   0  192M   0% /run/user/0
```

3. After successful expansion, please check the data integrity and observe whether the business in the cloud server is running normally.

If there are any abnormalities, you can use the snapshot rollback method for data recovery. For more information, see [Rolling Back Data from Snapshots](#).

Extending Data Disk Partitions and File Systems Offline

Determining the Expansion Method

Last updated: 2023-09-22 14:47:06

Scenario

Cloud Block Storage (CBS) is a scalable storage device in the cloud. You can expand its size at any time after creation to increase storage space without losing the existing data on the CBS.

After completing the [Cloud Disk Expansion](#) via the console, you also need to allocate the expanded capacity to the existing partitions within the cloud server instance. You should choose the best CBS expansion method based on your actual needs. This document explains how to determine the CBS expansion method on a Linux cloud server.

Note

The process of expanding the file system could potentially affect existing data. Therefore, it is advisable to manually [create a snapshot](#) to back up your data prior to initiating the operation.

Preparations

- You have already [expanded the cloud disk](#) via the console.
- The cloud disk has been attached to a Linux cloud server and a file system has been created. For more details, please refer to [Attaching Cloud Disks](#).
- You have logged into the Linux cloud server where the partition and file system to be expanded are located. For more details, please refer to [Logging in to Linux Instance Using Standard Login Method \(Recommended\)](#).

Instructions

1. Execute the following command as a root user to query the partition format used by the cloud disk.

```
fdisk -l
```

- If the result shows no partition (only /dev/vdb is displayed as in the image below), follow the steps in the [Expand File System](#) section below.

```
[root@VM_0_102_centos ~]# fdisk -l
Disk /dev/vda: 53.7 GB, 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: 0x000d64b4

   Device Boot      Start         End      Blocks   Id  System
/dev/vda1  *           2048     104857599     52427776   83  Linux

Disk /dev/vdb: 10.7 GB, 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

[root@VM_0_102_centos ~]#
```

- If the results are as shown in the two images below (which may vary depending on the operating system), follow the steps for [GPT partition](#) as outlined in the table below.

```

Disk /dev/vdb: 32.2 GB, 32212254720 bytes, 7864320 sectors
Units = sectors of 1 * 4096 = 4096 bytes
Sector size (logical/physical): 4096 bytes / 4096 bytes
I/O size (minimum/optimal): 4096 bytes / 4096 bytes
Disk label type: dos
Disk identifier: 0x00000000

   Device Boot      Start         End      Blocks   Id  System
/dev/vdb1          1      2621439     10485756   ee    GPT
    
```

```

WARNING: GPT (GUID Partition Table) detected on '/dev/vdb'! The util fdisk doesn't support GPT. Use GNU Parted.
    
```

- If the following result is shown (which may vary by operation system), proceed according to the **MBR Partition** steps in the table below.

```

[root@VM_20_49_tlinux ~]# fdisk -l

Disk /dev/xvda: 8589 MB, 8589934592 bytes
255 heads, 63 sectors/track, 1044 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0xf5a25329

   Device Boot      Start         End      Blocks   Id  System
/dev/xvda1    *           1          1044     8385898+   83  Linux

Disk /dev/xvdc: 107.4 GB, 107374182400 bytes
255 heads, 63 sectors/track, 13054 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x06a4a875

   Device Boot      Start         End      Blocks   Id  System
/dev/xvdc1          1          6526     52420063+   83  Linux
    
```

2. Based on the cloud disk partition format found in [Step 1](#), select the appropriate expansion method in accordance with the actual situation of the cloud disk.

Note

- MBR partition supports disk with a maximum capacity of 2 TB.
- When you partition a disk with a capacity greater than 2 TB, we recommend you create and mount a new data disk and use the GPT partition format to copy data.

Partition format	Scaling	Note
-	Expanding the File System	This is applicable for scenarios where no partition has been created and a file system has been directly created on the raw device.
GPT	Extending GPT Partitions and File Systems (Larger than 2 TB)	Expand the existing GPT partition.
	Extending GPT Partitions and File Systems (Larger than 2 TB)	Maintain the original partition unchanged, and create a new GPT partition for the expanded part.
MBR	Extending MBR Partitions and File Systems (Smaller than 2 TB)	Expand the existing MBR partition.

[Extending MBR Partitions and File Systems \(Smaller than 2 TB\)](#)

Maintain the original partition unchanged and create a new MBR partition with the expanded part.

Extending MBR Partitions and File Systems Offline (Smaller than 2 TB)

Last updated: 2023-09-26 16:38:13

Scenario

When your cloud disk has an existing MBR partition with a created file system and has been expanded to less than 2 TB, you can extend the partition and file system in one of the following ways, depending on your actual needs:

- [Assigning the expanded capacity to an existing MBR partition](#)
- [Formatting the expanded capacity into an independent new MBR partition](#)

Preparations

You can use automatic expansion tools including `fdisk`, `e2fsck` and `resize2fs` to add the expanded cloud disk capacity to the existing file system on a Linux CVM. To ensure a successful expansion, the following requirements must be met:

- Ensure that the partition format for expansion has been confirmed. For more information, refer to [Confirming the Expansion Method](#).
- The file system is EXT2, EXT3, EXT4, or XFS.
- The current file system does not have any error.
- The disk size after expansion does not exceed 2 TB.
- Only use Python version 2 because of compatibility with expansion tools in this document.

Instructions

Assigning the expanded capacity to an existing MBR partition

Run the following command as the root user to query partitions of the cloud disk.

```
lsblk
```

- If the information as shown in the following figure is returned, it indicates that there is only one partition. You can use the tool for automatic expansion. For details, please refer to [Using the Tool for Expansion](#).

```
[root@VM-8-57-centos ~]# lsblk
NAME MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sr0    11:0    1 118.6M  0  rom
vda    253:0    0   10G  0  disk
└─vda1 253:1    0   10G  0  part /
vdb    253:16   0   10G  0  disk
└─vdb1 253:17   0   10G  0  part
```

- If the returned information is similar to the image below, it indicates that there are two partitions, `vdb1` and `vdb2`. If you have two or more partitions, please refer to [manual expansion](#) to select the partition for expansion.

```
[root@VM-8-128-centos ~]# lsblk
NAME MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sr0    11:0    1 118.6M  0  rom
vda    253:0    0   10G  0  disk
└─vda1 253:1    0   10G  0  part /
vdb    253:16   0   30G  0  disk
└─vdb1 253:17   0   10G  0  part
└─vdb2 253:18   0   10G  0  part
```

Utilizing Tools for Expansion

Note

The tool-based expansion method supports scenarios with only one partition. If there are two or more partitions, please use the [manual expansion](#) method.

1. Run the following command as the root user to unmount the partition:

```
umount <mount point>
```

Taking the `/data` mount point as an example, execute the following:

```
umount /data
```

2. Run the following command to download a tool:

```
wget -O /tmp/devresize.py https://raw.githubusercontent.com/tencentyun/tencentcloud-cbs-tools/master/devresize/devresize.py
```

3. Run the following command to use the expansion tool for expansion:

```
python /tmp/devresize.py <disk path>
```

This document uses the disk path `/dev/vdb` and the file system on `vdb1` as an example. Proceed as follows:

```
python /tmp/devresize.py /dev/vdb
```

4. If the output is "The filesystem on `/dev/vdb1` is now XXXXX blocks long." as shown below, the expansion was successful. Please proceed to [Step 6](#).

```
[root@VM-2-4-centos Python-2.7.16]# python /tmp/devresize.py /dev/vdb
[INFO] - checking filesystem healthy
/dev/vdb1: 11/655360 files (0.0% non-contiguous), 66753/2621184 blocks
This operation will extend /dev/vdb1 to the last sector of device.
To ensure the security of your valuable data,
please create a snapshot of this volume before resize its file system, continue? [Y/n]
Y
It will resize (/dev/vdb1).
This operation may take from several minutes to several hours, continue? [Y/n]
Y
[INFO] - Backup MBR to /tmp/MBR_vdb1_2021-03-03_16:13:46_bak
[INFO] - resize filesystem
resize2fs 1.44.3 (10-July-2018)
Resizing the filesystem on /dev/vdb1 to 5242624 (4k) blocks.
The filesystem on /dev/vdb1 is now 5242624 (4k) blocks long.
[INFO] - Finished
```

5. If the output is `[ERROR] - e2fsck failed!!`, please follow the steps below:

a. Run the following command to repair the partition where the file system is located.

```
fsck -a <partition path>
```

This guide assumes that the disk path is `/dev/vdb` and the file system is on `vdb1`. Proceed as follows:

```
fsck -a /dev/vdb1
```

b. After the partition is fixed, run the following command again to use the tool for expansion.

```
python /tmp/devresize.py /dev/vdb
```

6. Run the following command to manually mount the expanded partition. In this document, `/data` is used as an example mount point.

```
mount <partition path> <mount point>
```

If there is an existing partition before expansion and the partition path is `/dev/vdb1` for instance, execute the following:

```
mount /dev/vdb1 /data
```

7. Run the following command to view the partition capacity after expansion:

```
df -h
```

If a message similar to the one shown in the following figure is returned, it indicates that the mounting was successful, and the data disk is now visible:

```
[root@VM-2-4-centos ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
devtmpfs        902M   0  902M   0% /dev
tmpfs           915M  24K  915M   1% /dev/shm
tmpfs           915M  416K  915M   1% /run
tmpfs           915M   0  915M   0% /sys/fs/cgroup
/dev/vda1       9.8G  2.2G  7.2G  23% /
tmpfs           183M   0  183M   0% /run/user/0
/dev/vdb1       20G   44M  19G   1% /data
```

8. Run the following command to view the data information of the original partition after expansion and check whether the newly added storage space has been expanded into the file system.

```
ll /data
```

Manual

1. Run the following command as the root user to unmount the partition:

```
umount <mount point>
```

Taking the `/data` mount point as an example, execute the following:

```
umount /data
```

2. Execute the following command to extend the `vdb2` partition. This document uses the expansion of the `vdb2` partition as an example, you can modify the command according to your actual situation.

```
growpart /dev/vdb 2
```

3. Run the following command to extend the file system of the partition.

```
resize2fs /dev/vdb2
```

If information similar to what is shown below is returned, the expansion has been successful.

```
[root@VM-8-128-centos ~]# resize2fs /dev/vdb2
resize2fs 1.42.9 (28-Dec-2013)
Resizing the filesystem on /dev/vdb2 to 5242875 (4k) blocks.
The filesystem on /dev/vdb2 is now 5242875 blocks long.
```

4. Run the following command to manually mount the expanded partition. In this document, we use `/data` as an example mount point.

```
mount <partition path> <mount point>
```

If there is an existing partition before expansion, and the partition path is `/dev/vdb2` for example, then execute:

```
mount /dev/vdb2 /data
```

5. Run the following command to view the partition capacity after expansion:

```
df -h
```

If a message similar to the one shown in the following figure is returned, it indicates that the mounting was successful, and the data disk is now visible:

```
[root@VM-8-57-centos ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
devtmpfs        909M   0  909M   0% /dev
tmpfs           919M  24K  919M   1% /dev/shm
tmpfs           919M 472K  919M   1% /run
tmpfs           919M   0  919M   0% /sys/fs/cgroup
/dev/vda1       9.8G  2.0G  7.4G  21% /
tmpfs           184M   0  184M   0% /run/user/0
/dev/vdb2       20G   44M  19G   1% /data
```

6. Run the following command to view the data information of the original partition after expansion and check whether the newly added storage space has been expanded into the file system.

```
ll /data
```

Formatting the expanded capacity into an independent new MBR partition

1. Run the following command as a root user to view the partition information of the mounted data disk:

```
df -h
```

The data disk partition of 20GB has been mounted, as shown in the following figure:

```
[root@VM-2-4-centos ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
devtmpfs        902M   0  902M   0% /dev
tmpfs           915M  24K  915M   1% /dev/shm
tmpfs           915M 416K  915M   1% /run
tmpfs           915M   0  915M   0% /sys/fs/cgroup
/dev/vda1       9.8G  2.2G  7.2G  23% /
tmpfs           183M   0  183M   0% /run/user/0
/dev/vdb1       20G   44M  19G   1% /data
```

2. Run the following command to view the information of the data disk without partitions after expansion:

```
fdisk -l
```

The data disk has been expanded to 30 GB, as shown below:

```
[root@VM-2-4-centos ~]# fdisk -l
Disk /dev/vda: 10 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): 262144 bytes / 262144 bytes
Disklabel type: dos
Disk identifier: 0x89ee0607

Device      Boot Start      End  Sectors  Size Id Type
/dev/vda1   *        2048 20971486 20969439  10G 83 Linux

Disk /dev/vdb: 30 GiB, 32212254720 bytes, 62914560 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 262144 bytes / 262144 bytes
Disklabel type: dos
Disk identifier: 0x7bb67b98

Device      Boot Start      End  Sectors  Size Id Type
/dev/vdb1   *        2048 41943039 41940992  20G 83 Linux
```

3. Run the following command to unmount all mounted partitions:

```
umount <mount point>
```

Taking the `/data` mount point as an example, execute the following:

```
umount /data
```

Note

Please ensure that all partitions on the cloud disk are unmounted before proceeding with [Step 4](#).

4. Run the following command to create a new partition.

```
fdisk <disk path>
```

This guide uses the disk path `/dev/vdb` as an example. Proceed as follows:

```
fdisk /dev/vdb
```

Perform the following steps in sequence when prompted.

- i. Enter **p**: to view the existing partition information. In this document, the existing partition is `/dev/vdb1`.
- ii. Enter **n**: to create a new partition.
- iii. Enter **p**: to create a new primary partition.
- iv. Enter **2**: to create the second primary partition.
- v. Press **Enter** twice: The partition size will use the default configuration.
- vi. Enter **w**: Save the partition table and start partitioning.

As shown below:

```
[root@VM-2-4-centos ~]# fdisk /dev/vdb
Welcome to fdisk (util-linux 2.32.1).
Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.

Command (m for help): p
Disk /dev/vdb: 30 GiB, 32212254720 bytes, 62914560 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 262144 bytes / 262144 bytes
Disklabel type: dos
Disk identifier: 0x7bb67b98

Device      Boot Start          End  Sectors  Size Id Type
/dev/vdb1           2048 41943039 41940992   20G 83 Linux

Command (m for help): n
Partition type
  p   primary (1 primary, 0 extended, 3 free)
  e   extended (container for logical partitions)
Select (default p): p
Partition number (2-4, default 2): 2
First sector (41943040-62914559, default 41943040):
Last sector, +sectors or +size{K,M,G,T,P} (41943040-62914559, default 62914559):

Created a new partition 2 of type 'Linux' and of size 10 GiB.

Command (m for help): w
The partition table has been altered.
Calling ioctl() to re-read partition table.
Syncing disks.
```

Note

This document takes creating one partition as an example. You can also create multiple partitions to meet your needs.

5. Run the following command to view the new partition:

```
fdisk -l
```

The new partition `vdb2` has been successfully created, as shown in the figure below:

```
[root@VM-2-4-centos ~]# fdisk -l
Disk /dev/vda: 10 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): 262144 bytes / 262144 bytes
Disklabel type: dos
Disk identifier: 0x89ee0607

Device      Boot Start          End  Sectors  Size Id Type
/dev/vda1   *      2048 20971486 20969439   10G 83 Linux

Disk /dev/vdb: 30 GiB, 32212254720 bytes, 62914560 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 262144 bytes / 262144 bytes
Disklabel type: dos
Disk identifier: 0x7bb67b98

Device      Boot Start          End  Sectors  Size Id Type
/dev/vdb1           2048 41943039 41940992   20G 83 Linux
/dev/vdb2     41943040 62914559 20971520   10G 83 Linux
```

6. Run the following command to format the new partition and create a file system in a desired format, such as EXT2 or EXT3.

```
mkfs.<fstype> <partition path>
```

Taking EXT4 as an example, run the following command:

```
mkfs.ext4 /dev/vdb2
```

The following figure shows the successful creation of the EXT4 file system.

```
[root@VM-2-4-centos ~]# mkfs.ext4 /dev/vdb2
mke2fs 1.44.3 (10-July-2018)
Creating filesystem with 2621440 4k blocks and 655360 inodes
Filesystem UUID: 87445645-4285-49d7-92d1-b0c291a7c710
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632

Allocating group tables: done
Writing inode tables: done
Creating journal (16384 blocks): done
Writing superblocks and filesystem accounting information: done
```

7. Run the following command to create a mounting point:

```
mkdir <new_mount_point>
```

Taking the new mount point `/data1` as an example, execute the following:

```
mkdir /data1
```

8. Execute the following command to manually mount the new partition:

```
mount <new partition path> <new mount point>
```

Assuming the new partition path is `/dev/vdb2` and the new mount point is `/data1`, execute the following:

```
mount /dev/vdb2 /data1
```

9. Run the following command to view the information of the new partition:

```
df -h
```

If the result as shown in the following figure is returned, the mounting is successful, and you can see the data disk.

```
[root@VM-2-4-centos ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
devtmpfs        902M   0  902M   0% /dev
tmpfs           915M  24K  915M   1% /dev/shm
tmpfs           915M 424K  915M   1% /run
tmpfs           915M   0  915M   0% /sys/fs/cgroup
/dev/vda1       9.8G  2.2G  7.2G  24% /
tmpfs           183M   0  183M   0% /run/user/0
/dev/vdb2       9.8G  37M  9.3G   1% /data1
/dev/vdb1       20G  44M  19G   1% /data
```

Note

If you want the CVM to automatically mount the data disk upon reboot or startup, you need to perform [Step 10](#) and [Step 11](#) to add the new partition information to `/etc/fstab`.

10. Run the following command to add information.

```
echo '/dev/vdb2 /data1 ext4 defaults 0 0' >> /etc/fstab
```

11. Run the following command to view the information: Execute the following.

```
cat /etc/fstab
```

If the result as shown in the following figure is returned, the partition has been successfully added.

```
[root@VM-2-4-centos ~]# cat /etc/fstab
#
# /etc/fstab
# Created by anaconda on Tue Nov 26 02:11:36 2019
#
# Accessible filesystems, by reference, are maintained under '/dev/disk/'.
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info.
#
# After editing this file, run 'systemctl daemon-reload' to update systemd
# units generated from this file.
#
UUID=659e6f89-11fa-4438-8442-ccdf2c06e0fe / ext4 defaults 1 1
/dev/vdb2 /data1 ext4 defaults 0 0
```

Documentation

[Extending Partitions and File Systems \(Windows\)](#)

FAQs Overview

If you encounter a problem when using Tencent Cloud CBS, refer to the following documents for troubleshooting as needed:

- [Usage-related issues](#)
- [Feature-related issues](#)

Extending GPT Partitions and File Systems (Larger than 2 TB)

Last updated: 2023-09-26 16:38:42

Scenario

If your cloud disk has a GPT partition that contains the file system, you can use either of the following methods to extend partitions and file systems:

- [Assigning the expanded capacity to an existing GPT partition](#)
- [Formatting the expanded capacity into an independent new GPT partition](#)

Preparations

You can use automatic expansion tools including `e2fsck` and `resize2fs` to add the expanded cloud disk capacity to the existing file system on a Linux CVM. To ensure a successful expansion, the following requirements must be met:

- Ensure that the partition format for expansion has been confirmed. For more information, refer to [Confirming the Expansion Method](#).
- The file system is EXT or XFS.
- The current file system does not have any error.

Instructions

Assigning the expanded capacity to an existing GPT partition

1. Run the following command as the root user to confirm changes in cloud disk capacity:

```
parted <disk path> print
```

Assuming the disk path in this guide is `/dev/vdc`, execute the following:

```
parted /dev/vdc print
```

If the following message appears during the process, please enter `Fix`. See below:

```
[root@VM-2-4-centos ~]# parted /dev/vdc print
Warning: Not all of the space available to /dev/vdc appears to be used, you can fix the GPT to use all of the space
(an extra 3963617280 blocks) or continue with the current setting?
Fix/Ignore? Fix
```

The expanded cloud disk size is 2040 GB, and the existing partition size is 10.7 GB. See below:

```
Model: Virtio Block Device (virtblk)
Disk /dev/vdc: 2040GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start   End     Size    File system  Name  Flags
  1      1049kB  10.7GB 10.7GB  ext4         opt
```

2. Run the following command to check whether the cloud disk has partitions mounted.

```
mount | grep '<disk path>'
```

Assuming the disk path in this guide is `/dev/vdc`, execute the following:

```
mount | grep '/dev/vdc'
```

- If the following result is returned, it indicates that a partition (vdc1) on the cloud disk is mounted on `/data`.

```
[root@VM-2-4-centos ~]# mount | grep '/dev/vdc'
/dev/vdc1 on /data type ext4 (rw,relatime,stripe=64)
```

Run the following command to unmount all partitions on the cloud disk.

```
umount <mount point>
```

Taking the `/data` mount point as an example, execute the following:

```
umount /data
```

- The following result indicates that there is no partition mounted. Proceed to the next step.

```
[root@VM-2-4-centos ~]# mount | grep '/dev/vdc'
[root@VM-2-4-centos ~]#
```

3. Run the following command to enter the parted partition tool:

```
parted <disk path>
```

Assuming the disk path in this guide is `/dev/vdc`, execute the following:

```
parted /dev/vdc
```

4. Run the following command to change the display and operational unit to sector (GB by default).

```
unit s
```

5. Run the following command to view partition information and record the Start value of the existing partition.

```
print
```

Note

Please ensure to record the Start value. After deleting and creating a new partition, the Start value must remain unchanged, otherwise, it may lead to data loss.

```
(parted) unit s
(parted) print
Model: Virtio Block Device (virtblk)
Disk /dev/vdc: 3984588800s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start  End      Size      File system  Name  Flags
 1      2048s  20969471s  20967424s  ext4         opt
```

6. Run the following command to delete the existing partition.

```
rm <Partition Number>
```

For example, run the following command to delete the partition "1" from the cloud disk.

```
rm 1
```

7. Run the following command to confirm the deletion. The returned information is as shown below:

```
print
```

```
(parted) rm 1
(parted) print
Model: Virtio Block Device (virtblk)
Disk /dev/vdc: 3984588800s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:
Number  Start  End  Size  File system  Name  Flags
(parted) █
```

Note

In case of accidental partition deletion, immediately execute the `rescue` command and enter the Start and End values as prompted to confirm partition recovery.

8. Run the following command to create a new primary partition:

```
mkpart primary <Start sector of original partition> 100%
```

100% indicates that this partition extends to the end of the disk. The Start value can be obtained from [Step 5](#). Please fill in according to your actual situation. In this article, the sector of the original partition starts at 2048s, so the Start value is 2048. Execute the following:

```
mkpart primary 2048s 100%
```

If the status shown below appears, please enter `Ignore`.

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

9. Run the following command to check whether the new partition has been created successfully:

```
print
```

If the result as shown in the following figure is returned, the new partition has been created successfully.

```
(parted) mkpart primary 2048s 100%
(parted) print
Model: Virtio Block Device (virtblk)
Disk /dev/vdc: 3984588800s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:
Number  Start  End  Size  File system  Name  Flags
1       2048s  3984586751s  3984584704s  primary
```

10. Run the following command to exit the parted tool:

```
quit
```

11. Execute the following command to synchronize the new partition table changes to the operating system.

```
partprobe
```

12. Execute the following command to inspect the partition after expansion.

```
e2fsck -f <partition path>
```

Assume that you want to create a partition "1" (that is, the partition path is `/dev/vdc1`), then run:

```
e2fsck -f /dev/vdc1
```

The following figure shows the command output.

```
[root@VM-2-4-centos ~]# e2fsck -f /dev/vdc1
e2fsck 1.44.3 (10-July-2018)
Pass 1: Checking inodes, blocks, and sizes
Pass 2: Checking directory structure
Pass 3: Checking directory connectivity
Pass 4: Checking reference counts
Pass 5: Checking group summary information
/dev/vdc1: 11/655360 files (0.0% non-contiguous), 66753/2620928 blocks
```

13. Please extend the file system on the newly created partition according to your actual situation.

- For EXT file systems, execute the following command:

```
resize2fs <partition path>
```

Assuming the partition path is `/dev/vdc1` , execute the following:

```
resize2fs /dev/vdc1
```

If the result as shown in the following figure is returned, the expansion is successful.

```
[root@VM-2-4-centos ~]# resize2fs /dev/vdc1
resize2fs 1.44.3 (10-July-2018)
Resizing the filesystem on /dev/vdc1 to 498073088 (4k) blocks.
The filesystem on /dev/vdc1 is now 498073088 (4k) blocks long.
```

- For XFS file system, execute the following command:

```
xfs_growfs <partition path>
```

Assuming the partition path is `/dev/vdc1` , execute the following:

```
xfs_growfs /dev/vdc1
```

14. Execute the following command to manually mount the new partition:

```
mount <partition path> <mount point>
```

This guide uses the partition path `/dev/vdc1` and the mount point `/data` as examples. Proceed as follows:

```
mount /dev/vdc1 /data
```

15. Run the following command to view the new partition:

```
df -h
```

If the result as shown in the following figure is returned, the mounting is successful, and you can see the data disk.

```
[root@VM-2-4-centos ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
devtmpfs        902M   0  902M   0% /dev
tmpfs           915M  24K  915M   1% /dev/shm
tmpfs           915M 420K  915M   1% /run
tmpfs           915M   0  915M   0% /sys/fs/cgroup
/dev/vda1       9.8G  2.2G  7.2G  24% /
tmpfs           183M   0  183M   0% /run/user/0
/dev/vdc1       1.9T   59M  1.8T   1% /data
```

Formatting the expanded capacity into an independent new GPT partition

1. Run the following command as the root user to confirm changes in cloud disk capacity:

```
parted <disk path> print
```

Assuming the disk path in this guide is `/dev/vdc`, execute the following:

```
parted /dev/vdc print
```

If you receive the following message during the process, please enter `Fix`.

```
[root@VM-2-4-centos ~]# parted /dev/vdc print
Warning: Not all of the space available to /dev/vdc appears to be used, you can fix the GPT to use all of the space
(an extra 209715200 blocks) or continue with the current setting?
Fix/Ignore? Fix
```

The size of the expanded cloud disk is 2147GB, and the size of the existing partition is 2040GB, as shown below:

```
Model: Virtio Block Device (virtblk)
Disk /dev/vdc: 2147GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start   End     Size    File system  Name      Flags
  1      1049kB 2040GB 2040GB  ext4         primary
```

2. Run the following command to check whether the cloud disk has partitions mounted.

```
mount | grep '<disk path>'
```

Assuming the disk path in this guide is `/dev/vdc`, execute the following:

```
mount | grep '/dev/vdc'
```

- If the following result is returned, it indicates that a partition (`vdc1`) on the cloud disk is mounted on `/data`.

```
[root@VM-2-4-centos ~]# mount | grep '/dev/vdc'
/dev/vdc1 on /data type ext4 (rw,relatime,stripe=64)
```

Run the following command to unmount **all partitions** on the cloud disk.

```
umount <mount point>
```

Taking the `/data` mount point as an example, execute the following:

```
umount /data
```

- The following result indicates that there is no partition mounted. Proceed to the next step.

```
[root@VM-2-4-centos ~]# umount /data
[root@VM-2-4-centos ~]#
```

- Run the following command to enter the parted partition tool:

```
parted '<disk path>'
```

Assuming the disk path in this guide is `/dev/vdc`, execute the following:

```
parted '/dev/vdc'
```

Run the following command to view partition information and record the End value of the existing partition. This value will be used as the starting offset for the next partition.

```
print
```

```
(parted) print
Model: Virtio Block Device (virtblk)
Disk /dev/vdc: 2147GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start   End     Size    File system  Name      Flags
  1      1049kB  2040GB  2040GB  ext4         primary
```

- Execute the following command to create a new primary partition. This partition will start from the end of the existing partition and cover all the newly added disk space.

```
mkpart primary start end
```

Based on the End value obtained from [Step 4](#), please fill in according to your actual situation. In this document, the End value is 2040 GB, then execute:

```
mkpart primary 2040GB 100%
```

- Run the following command to check whether the new partition has been created successfully:

```
print
```

If the following output is returned, the partition has been created.

```
(parted) mkpart primary 2040GB 100%
(parted) print
Model: Virtio Block Device (virtblk)
Disk /dev/vdc: 2147GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start   End     Size    File system  Name      Flags
  1      1049kB  2040GB  2040GB  ext4         primary
  2      2040GB  2147GB  107GB   ext4         primary
```

- Run the following command to exit the parted tool:

```
quit
```

- Run the following command to format the newly created partition: You can select a file system in EXT2 or EXT3 format as needed.

```
mkfs.<fstype> <partition path>
```

Taking EXT4 as an example, run the following command:

```
mkfs.ext4 /dev/vdc2
```

9. Execute the following command to manually mount the new partition:

```
mount <partition path> <mount point>
```

Taking the partition path `/dev/vdc2` and the mount point `/data` as an example, execute the following:

```
mount /dev/vdc2 /data
```

10. Run the following command to view the new partition:

```
df -h
```

If the result as shown in the following figure is returned, the mounting is successful, and you can see the data disk.

```
[root@VM-2-4-centos ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
devtmpfs        902M   0  902M   0% /dev
tmpfs           915M  24K  915M   1% /dev/shm
tmpfs           915M 424K  915M   1% /run
tmpfs           915M   0  915M   0% /sys/fs/cgroup
/dev/vda1       9.8G  2.2G  7.2G  24% /
tmpfs          183M   0  183M   0% /run/user/0
/dev/vdc2       98G   61M  93G   1% /data
```

Documentation

[Extending Partitions and File Systems \(Windows\)](#)

FAQs Overview

If you encounter a problem when using Tencent Cloud CBS, refer to the following documents for troubleshooting as needed:

- [Usage-related issues](#)
- [Feature-related issues](#)

Extending Bare Data Disk File Systems

Last updated: 2023-09-19 20:33:05

Scenario

This document describes how to extend a file system after logging in to the CVM instance. This method is suitable for scenarios where the file system is directly created without partitioning the cloud disk.

Instructions

1. Run the following command to determine the file system type.

```
df -ihT
```

- The following result shows an EXT file system.

```
[root@VM-2-4-centos ~]# df -ihT
Filesystem      Type      Inodes  IUsed  IFree  IUse%  Mounted on
devtmpfs        devtmpfs  228K    334    227K   1%     /dev
tmpfs           tmpfs     230K    7      230K   1%     /dev/shm
tmpfs           tmpfs     230K    444    230K   1%     /run
tmpfs           tmpfs     230K    16     230K   1%     /sys/fs/cgroup
/dev/vda1       ext4      640K    58K    583K   10%    /
tmpfs           tmpfs     230K    1      230K   1%     /run/user/0
/dev/vdb        ext4      640K    11     640K   1%     /data
```

- The following result shows a XFS file system.

```
[root@VM-2-4-centos ~]# df -ihT
Filesystem      Type      Inodes  IUsed  IFree  IUse%  Mounted on
devtmpfs        devtmpfs  228K    334    227K   1%     /dev
tmpfs           tmpfs     230K    7      230K   1%     /dev/shm
tmpfs           tmpfs     230K    444    230K   1%     /run
tmpfs           tmpfs     230K    16     230K   1%     /sys/fs/cgroup
/dev/vda1       ext4      640K    58K    583K   10%    /
tmpfs           tmpfs     230K    1      230K   1%     /run/user/0
/dev/vdc        xfs       5.0M    3      5.0M   1%     /data
```

2. Use the file system-specific command to extend the file system.

Note

Note the following capacity limits for an EXT file system:

- EXT3 supports up to 16 TB, and each file cannot exceed 2 TB.
- EXT4 supports up to 1 EB, and each file cannot exceed 16 TB.

Expanding the EXT File System

Run the following command to extend the EXT file system (using `/dev/vdb` as an example).

```
resize2fs /dev/vdb
```

If the following command output is returned, the expansion is successful.

```
[root@VM_0_102_centos ~]# resize2fs /dev/vdb
resize2fs 1.42.9 (28-Dec-2013)
The filesystem is already 5242880 blocks long.  Nothing to do!

[root@VM_0_102_centos ~]#
```

Extending the XFS File System

Run the following command to extend the XFS file system (using `/dev/vdc` as an example).

```
xfs_growfs /dev/vdc
```

If the following command output is returned, the expansion is successful.

```
[root@VM-2-4-centos ~]# xfs_growfs /dev/vdc
meta-data=/dev/vdc          isize=512    agcount=16, agsize=163840 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    =version 2       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 5242880
```

3. Run the following command to view the disk space of the file system.

```
df -h
```

Adjusting cloud disk types

Last updated: 2023-09-19 20:55:12

Scenario

Tencent Cloud CBS supports the adjustment of storage hardware types during business operation. You can upgrade the type while the service is online to meet the higher performance requirements of business development. After the fee is settled, the cloud disk type will take effect immediately. For the billing details of adjusting the cloud disk type, please refer to [Billing Details for Adjusting Cloud Disk Type](#).

Currently, adjusting the cloud disk type only supports upgrades and does not support downgrades. The details are as follows:

- Standard cloud disks can be adjusted to Premium Cloud Disks, Balanced SSDs, and SSD Cloud Disks.
- Premium Cloud Disks can be adjusted to Balanced SSDs and SSD Cloud Disks.
- A SSD cloud disk cannot be upgraded currently.

Preparations

- **Cloud Server Status:** For cloud disks that have been mounted to a cloud server, the type of the cloud disk can only be adjusted when the cloud server is in the **Running** or **Shut Down** state.
- **Cloud Disk Status:** Adjusting the cloud disk type is only supported when the cloud disk is in the **In Use** or **Unmounted** state.
- **Adjustment of cloud disk types is currently not supported for system disks and non-elastic data disks.**
- The adjustment of cloud disk types is not supported in Guangzhou Zone 1.
- Adjusting the cloud disk type is only supported when there are available upgrade types in the current availability zone and the current disk size is within the range supported by the target cloud disk.
- Adjusting the cloud disk type does not alter the disk size. After the adjustment, you can modify the disk size by following the [Expand Cloud Disk](#) procedure.
- Adjusting the cloud disk type does not change the lifecycle, disk ID, disk device name, or mounting point of the cloud disk.

Supports and Limits

- The operation of adjusting the cloud disk type involves copying data from the source cloud disk to the destination cloud disk. Depending on the size of the data and the speed of data transfer, this operation may take a considerable amount of time to complete. Please be patient. The adjustment process may have a minor impact on the performance of the cloud disk. If it is a critical service application, it is recommended to perform the adjustment operation during off-peak business hours.
- Downgrading of cloud disk types is currently not supported.
- After the adjustment, we recommend you to start up and log in to the CVM to check the data integrity.

Instructions

1. Log in to the [Cloud Disk Console](#) and enter the Cloud Disk list.
2. In the row of the target elastic cloud disk, select **More > Adjust Cloud Disk Type**.
3. In the "Adjust Cloud Disk Type" dialog box, select the target cloud disk type that needs to be adjusted, check the agreement box, and click **Convert Now**.
4. Make necessary payment if applicable, and wait for the operation to complete.

Detaching cloud disks

Last updated: 2023-09-21 14:32:43

Scenario

When you need to mount an elastic cloud disk with the disk attribute set to **Data Disk** to another CVM, you can proactively unmount the elastic cloud disk from the CVM and mount it to another CVM. **Unmounting an elastic cloud disk will not erase the data on the disk.**

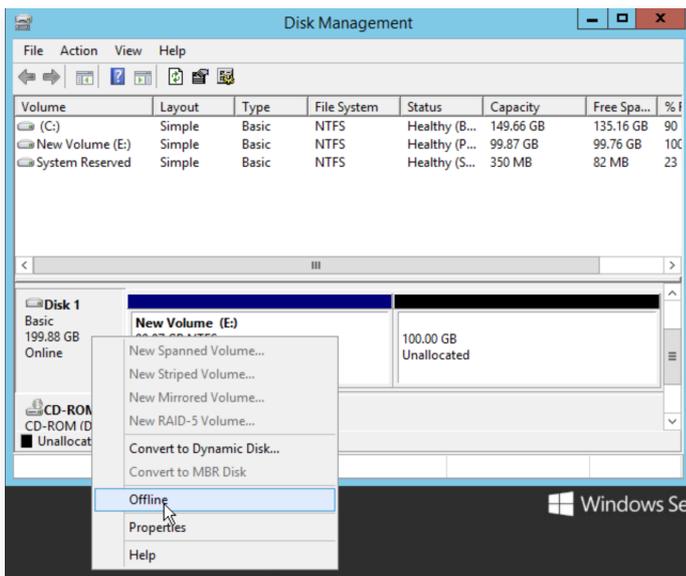
Currently, only elastic cloud disks with the disk attribute set to **Data Disk** can be unmounted. System disks and non-elastic cloud disks cannot be unmounted. **Before unmounting a cloud disk, you must first execute the `umount` (Linux) or `offline` (Windows) operation, otherwise, the CVM may not recognize the elastic cloud disk when it is mounted again.**

Preparations

Before unmounting the data disk, make sure you understand the following prerequisites:

Windows OS

- To prevent data loss, we recommend that you suspend read and write operations on all file systems of the disk. Otherwise, data that has not been read or written will be lost.
- Ensure that no processes are occupying the cloud disk that needs to be unmounted (such as Task Manager processes). Before unmounting an elastic cloud disk, you need to set the disk to offline status. Otherwise, without rebooting the CVM, you may not be able to mount the elastic cloud disk again. See below:



Linux OS

- You need to first **log in** to the instance and perform the `umount` operation on the elastic cloud disk that needs to be unmounted. If it is forcibly unmounted without performing the `umount` operation, problems as shown in the following

figure may occur during shutdown and startup:

```
Checking filesystems
/dev/vda1: clean, 35630/524288 files, 335690/2096474 blocks
fsck.ext3: Unable to resolve 'UUID=dabe8ee8-221b-44c7-9074-4d3f8fc4ae44'
fsck.ext3: No such file or directory while trying to open /dev/disk/by-id/virtio-disk-ezy5q516-part5
/dev/disk/by-id/virtio-disk-ezy5q516-part5:
The superblock could not be read or does not describe a correct ext2
filesystem. If the device is valid and it really contains an ext2
filesystem (and not swap or ufs or something else), then the superblock
is corrupt, and you might try running e2fsck with an alternate superblock:
    e2fsck -b 8193 <device>

[FAILED]

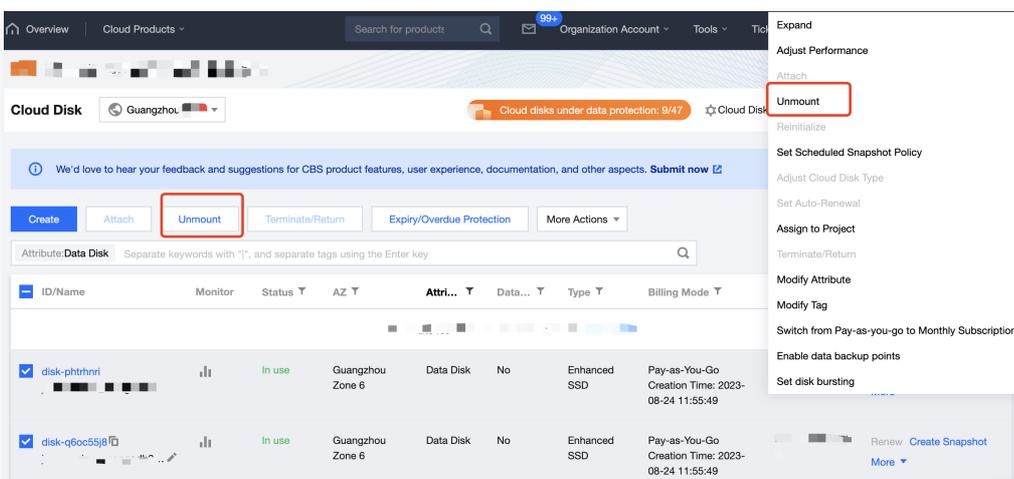
*** An error occurred during the file system check.
*** Dropping you to a shell; the system will reboot
*** when you leave the shell.
Give root password for maintenance
(or type Control-D to continue):
```

- If you have created a Logical Volume Manager (LVM) within the CVM, directly unmounting the disk from the console will leave some device data in the CVM's memory. If an application within the CVM tries to traverse or access the device, a system error will occur. Therefore, you need to perform the following operations in advance (this example assumes that a logical volume `/dev/test/lv1` is created based on `/dev/vdb1` and mounted in the `/data` directory):
 - a. Run the `umount /data` command to unmount the corresponding disk mount point within the CVM.
 - b. Run the `lvremove /dev/test/lv1` command to remove the Logical Volume (LV). If there are multiple LVs, remove all LVs in sequence.
 - c. Run the `vgremove test` command to remove the volume group.
 - d. Run the `pvremove /dev/vdb1` command to remove the physical volume.
 - e. Modify the `/etc/fstab` file to prevent the corresponding LV from being mounted again at the next startup.

Instructions

Unmounting Cloud Disks via the Console

1. Log in to the [Cloud Disk Console](#).
2. You can unmount the cloud disk in the following ways:
 - a. Single unmount: Click on **More > Unmount** in the row of the target cloud disk with the status **Mounted**.
 - b. Batch unmount: Select multiple target cloud disks with the status **Mounted** and click **Unmount** at the top of the list.



3. In the pop-up **Unmount Cloud Disk** dialog box, confirm the warning items and click **Confirm** to complete the unmounting process.

Using the API to Unmount Cloud Disks

You can use the DetachDisks interface to unmount the cloud disk. For more details, please refer to [Unmounting Cloud Disk](#).

FAQs

If you cannot unmount the cloud disk from your Windows CVM via the console, please refer to [Windows CVM Unable to Unmount Cloud Disk](#) to troubleshoot and resolve the issue.

Viewing Cloud Disk Information

Last updated: 2023-09-19 21:01:39

Scenario

This document describes how to view your cloud disk information in the console in the following ways:

- View the information of all cloud disk instances in a specific region on the [Cloud Disk](#) console.
- On the cloud disk details page, view the detailed information of a cloud disk.

Instructions

Viewing Cloud Disk List Information

1. Log in to the [Cloud Disk Console](#) and navigate to the cloud disk list page.
2. On this page, you can view information such as ID/name, monitoring, status, availability zone, attribute, cloud disk type, cloud disk capacity, associated instance, billing mode, projects, tags, and custom fields.
3. You can select  in the upper right corner, and in the pop-up "Customize List Field" window, choose the list field information you want to display.

Viewing Detailed Cloud Disk Information

1. Log in to the [Cloud Disk Console](#) and select the region where the cloud disk is located at the top of the page.
2. Find the target cloud disk in the list and click **ID/Name** to enter its details page.
3. On the cloud disk details page, you can view detailed information such as parameters, operation logs, historical data backup points, snapshot chain, and disk monitoring.

Terminating Cloud Disks

Last updated: 2023-09-22 09:24:08

Scenario

When a cloud disk is no longer in use and **important data has been backed up**, you can release virtual resources by terminating the cloud disk. After the cloud disk is terminated, no fees will be charged for it. **When a cloud disk is terminated, all data in the cloud disk will be deleted and cannot be recovered. A terminated cloud disk cannot be restored, so please proceed with caution.**

The lifecycle of a non-elastic cloud disk follows that of the CVM and can only be terminated when the CVM is terminated. For specific operations, please refer to [Terminating Instances](#).

The lifecycle of an elastic cloud disk is independent of the CVM, so it can be terminated independently of the CVM. This document mainly introduces the related operations of terminating an elastic cloud disk.

Elastic cloud disks support the following termination methods:

Manual Termination

- Monthly subscription cloud disks can be manually terminated before expiry. After being terminated, the cloud disks will be kept in the recycle bin for 7 days, and they can be permanently terminated in the recycle bin.
Each account is entitled to a 5-day unconditional return for one monthly subscription cloud disk, and up to 199 monthly subscription cloud disks can be returned for a refund per year under a single account. For refund details, please refer to [Refund Instructions](#). After exceeding the returnable quantity, you will not be able to manually terminate the monthly subscription cloud disks.
- Manual termination is supported for pay-as-you-go cloud disks, and takes effect immediately.

Automatic Termination

- Monthly subscription cloud disks that are not recovered from the recycle bin within 7 days will be automatically terminated. You can continue to use them if you complete the [renewal](#) within the stipulated time.
- Pay-as-you-go cloud disks will be automatically terminated if the balance remains below zero for 24 hours. You can continue to use them if you [recharge](#) within the stipulated time.

Data Wipe

The data deleted when terminating a cloud disk cannot be accessed by anyone through any means, and the data that has been deleted in the cloud disk system will definitely be completely wiped. The integrity of data wiping is ensured mainly through the following mechanisms:

- The operation of deleting the logical space of a cloud disk is recorded as metadata and initiates a write-zero operation on the physical disk space, forcing permanent deletion. Any read operation on this logical space will be ensured by the storage system to return all zeros.
- When a cloud disk is completely released, the storage system immediately destroys the metadata to ensure that the data can no longer be accessed. At the same time, the physical storage space corresponding to the cloud disk is reclaimed, and it is guaranteed to be zeroed before being allocated again.

Preparations

- The cloud disk is in the **To be attached** status. For cloud disks that are already in use, you need to [unmount](#) them first.
- You have **backed up important data** according to business requirements.

Instructions

Manually Terminating Monthly Subscription Cloud Disks

Manually terminating unexpired cloud disks in the console

When you no longer need a monthly subscription cloud disk, you can manually terminate it. Once the status of the cloud disk changes to **pending recycle**, no related fees will be incurred, **the system will suspend the cloud disk (the cloud disk will be unavailable, only data will be retained) and it will be moved to the recycle bin for 7 days. If you [renew](#) within 7 days, you can restore and continue to use the cloud disk. If you do not renew within 7 days, it will be automatically terminated. You can also log in to the recycle bin to permanently terminate it.**

Each account is entitled to a 5-day unconditional return for one monthly subscription cloud disk, and up to 199 monthly subscription cloud disks can be returned for a refund per year under a single account. For refund details, please refer to [Refund Instructions](#). After exceeding the returnable quantity, you will not be able to manually terminate the monthly subscription cloud disks.

1. Log in to the [Cloud Disk Console](#).
2. You can terminate cloud disks in the following ways:
 - a. **Individual Termination:** In the row where the target cloud disk with the **Ready to Mount** status is located, select **More > Terminate/Return**.
 - b. **Batch Termination:** Select multiple target cloud disks with the **Ready to Mount** status, and click **Terminate/Return** at the top of the list.

Note

When the cloud disk is terminated, all data on the cloud disk will also be deleted, and cannot be restored. Please note that cloud disks that have been terminated cannot be recovered.

3. In the **Terminate Cloud Disk** pop-up window, select **Read and agree to [Refunding Policy](#)**, then click **Submit**. The target cloud disk will stop being billed, be suspended (the cloud disk will be unavailable, only data will be

Terminate Cloud Disk
✕

You've selected 1 cloud disk, [Collapse](#)

Cloud Disk Name	Cloud Disk ID	Billing Mode	Expiration/Creation Time
■	disk-rl0	Monthly Subscription	Expiration Time: 202... ■■■

Are you sure you want to terminate the selected cloud disk?

i Terminated cloud disk will be kept in **Recycle Bin** for 7 days, please back up the data in advance.
After the deletion of resources, the amount falling under the refund within five days without reasons is returned to the original payment account. The ordinary refund amount is returned to your account according to the ratio of cash and voucher amount of your purchase.

Read and agree to [Refunding Policy](#) [↗](#)

Submit
Close

retained), and moved to the recycle bin.

Permanently terminate the monthly subscription cloud disk in the recycle bin

You can permanently terminate the monthly subscription cloud disks located in the [recycle bin](#).

1. Log in to the [Cloud Disk Recycle Bin](#).
2. You can permanently terminate a cloud disk in the following ways:

- a. **Individual Termination:** In the row where the target cloud disk with the **Pending Recycle** status is located, select **Release**.
- b. **Batch Termination:** Select multiple target cloud disks with the **Pending Recycle** status, and click **Batch Release** at the top of the list.

Note

When the cloud disk is terminated, all data on the cloud disk will also be deleted, and cannot be restored. Please note that cloud disks that have been terminated cannot be recovered.

3. Enter the verification code in the pop-up window and click **OK** to complete the termination. The target cloud disk will be **permanently terminated and cannot be recovered**.

Manual Termination of Pay-as-you-go Cloud Disks

1. Log in to the [Cloud Disk Console](#).
2. You can terminate cloud disks in the following ways:
 - a. **Individual Termination:** In the row where the target cloud disk with the **Ready to Mount** status is located, select **More > Terminate/Return**.
 - b. **Batch Termination:** Select multiple target cloud disks with the **Ready to Mount** status, and click **Terminate/Return** at the top of the list.

Note

When the cloud disk is terminated, all data on the cloud disk will also be deleted, and cannot be restored. Please note that cloud disks that have been terminated cannot be recovered.

3. In the pop-up **Terminate Cloud Disk** prompt box, click **Submit** to complete the termination. The target cloud disk will stop charging and will be **permanently terminated and cannot be recovered**.

Adjusting Cloud Disk Performance

Last updated: 2023-09-19 21:05:05

Cloud disk performance is generally related to its capacity. You can adjust the capacity to achieve higher performance when the disk has not reached its maximum performance. Specifically, Enhanced SSD cloud disks support breaking through the basic performance limit by configuring extra performance after reaching the maximum value of basic performance. You can configure and adjust the extra performance as needed when the conditions are met. For more information, please refer to [Enhanced SSD Cloud Disk Performance Description](#).

Note

- Currently, only **Enhanced SSD Cloud Disks** support independent performance adjustments.
- **Basic Performance** must reach its maximum value before **Additional Performance** can be independently adjusted.
- The performance adjustment will not affect your cloud disks or businesses.

Billing Methods for Performance Adjustment

Performance Upgrade

- For monthly subscribed cloud disks: The performance upgrade cost is calculated based on the remaining time of the lifecycle and the price difference between the new and old configurations. The specific cost will be calculated according to the actual situation, and you can view it on the payment page.
- For cloud disks featuring pay-as-you-go billing: the performance downgrade takes effect immediately, and the cloud disks start to be charged according to the new configuration price.

Performance Downgrade

- For monthly subscribed cloud disks: Performance downgrades will be refunded by calculating the price difference between the remaining value of the lifecycle and the value of the new configuration. The specific situation will be calculated based on the actual circumstances, and you can view it on the payment page.
- For cloud disks featuring pay-as-you-go billing: the performance downgrade takes effect immediately, and the cloud disks start to be charged according to the new configuration price.

Performance Upgrade

Upgrade performance using the console

When prerequisites are met, you can upgrade a disk as instructed below in the console:

1. Log in to the [Cloud Disk Console](#).
2. Select the region and the cloud disk that requires performance adjustment.
3. Select **More > Adjust Performance** for the target cloud disk.
4. Select the target configurations you need to adjust in the pop-up “Adjust Performance” window.
5. Select descriptions and start the adjustment.

Upgrade performance using API

You can use the `ModifyDiskExtraPerformance` interface to upgrade the performance of a specified cloud disk. For detailed instructions, please refer to [Adjusting Extra Performance of Cloud Disks](#).

Performance Downgrade

Downgrade performance using the console

When prerequisites are met, you can downgrade performance in the following ways:

1. Log in to the [Cloud Disk Console](#).
2. Select the region and the cloud disk that requires performance adjustment.
3. Select **More > Adjust Performance** for the target cloud disk.
4. Select the target configurations you need to adjust in the pop-up “Adjust Performance” window.
5. Select descriptions and start the adjustment.

Downgrade performance using API

You can use the ModifyDiskExtraPerformance interface to downgrade the performance of a specified cloud disk. For specific operations, please refer to [Adjusting Additional Performance of Cloud Disk](#).

Cloud Disk Data Protection Configuration

Last updated: 2023-09-26 14:34:52

Scenario

Cloud disks will be repossessed by the system due to expiration or overdue payment. To avoid data loss in these cases, you can strengthen the data protection of cloud disks in the following two ways:

- **Cloud Disk Expiration/Overdue Payment Protection**: After enabling this feature, if your cloud disk is repossessed by the system due to expiration or overdue payment, the system will automatically create a snapshot to protect your data.
- **Snapshot Protection**: Once enabled, a scheduled snapshot policy will be associated by default when creating a cloud disk.

Note

- For information on the cloud disk repossession mechanism and overdue payment handling, please see [Arrears Reminder](#).
- Users in domestic regions currently enjoy a free snapshot quota of 80GB. Snapshot capacity exceeding the quota or not complying with the free quota policy will be billed. For more details, please refer to [Snapshot Billing Overview](#).

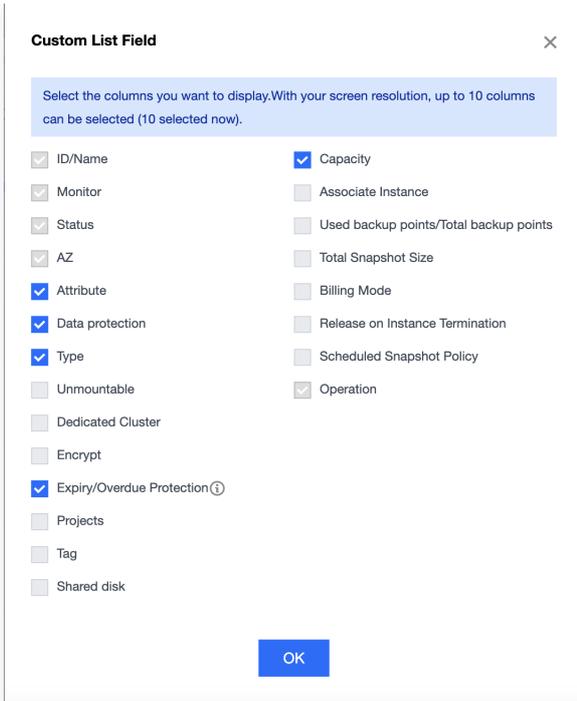
Instructions

Setting Up Cloud Disk Expiration/Overdue Payment Protection

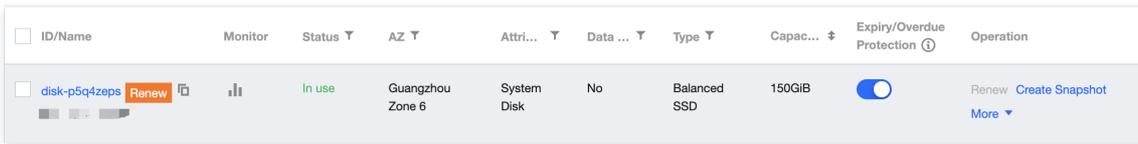
1. Log in to the CVM console and select **Cloud Disk** from the left sidebar.
2. At the top of the **Cloud Disk** page, select the region where the cloud disk is located.
3. Configure the parameter in the following ways as needed:

Setting a Single Cloud Disk

1. Select  in the top-right corner of the list to open the Display Settings window.
2. In the "Customize List Fields" pop-up window, check **Expiry/Overdue Protection** and click **Confirm**, as shown below:
You may select the remaining fields based on your actual situation.

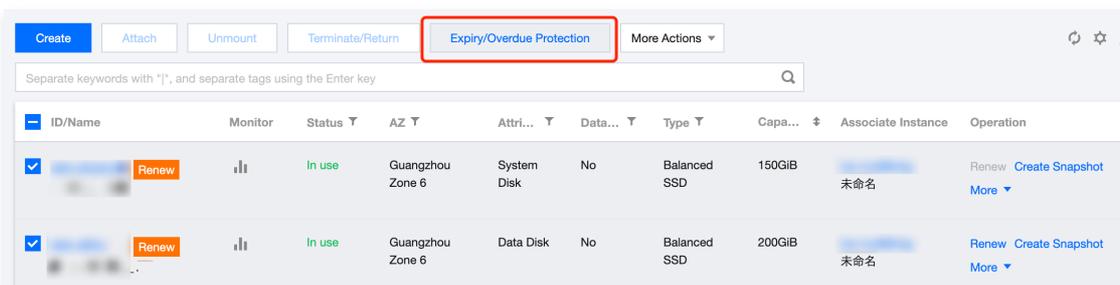


3. To enable this feature, simply toggle the **Expiry/Overdue Protection** switch in the row of the desired cloud disk. As shown in the figure below:



Setting up Multiple Cloud Disks

1. After selecting a cloud disk on the Cloud Disk List page, choose **Expiry/Overdue Protection** at the top of the list, as shown below:



2. In the **Expiration/Overdue Payment Protection** pop-up, click **OK** to complete the setup.

Setting all cloud disks in the availability zone

1. Select **Cloud Disk Data Protection Configuration** in the top-right corner of the page.
2. In the pop-up **Cloud Disk Data Protection Configuration** window, select **Expiry/Overdue Protection** and choose the available zone from the dropdown list, as shown below:

Cloud Disk Data Protection Configuration ✕

i Snapshot can help prevent the risk of data loss caused by mis-operations and viruses, so it is recommended that you enable it. In each region within China, there is a free quota of 80GiB available. The monthly price for the Guangzhou region is 0.13 CNY/GiB. For details, please see [Snapshot Billing Overview](#).

Snapshot Protection for New Cloud Disks

CBS purchase page [Enable Scheduled Snapshot by Default](#)

If it's enabled, Scheduled Snapshot will be enabled by default when you purchase a new cloud disk. You can **disable it manually**.

Expiry/Overdue Protection i

Enable for all cloud disks in this availability zone

You can use "Display Settings" to display this feature, so as to enable and disable protection easily.

Please select the AZ. ▼

Save
Close

3. Click **Save** to complete the setup.

Setting Snapshot Protection

1. Log in to the CVM console and select **Cloud Disk** on the left sidebar.
2. At the top of the **Cloud Disk** page, select the region where the cloud disk is located.
3. Configure the parameter in the following ways as needed:

Associating a scheduled snapshot policy when creating a cloud disk

1. On the **Cloud Disk** page, click **Create** at the top of the list.
2. In the pop-up **Purchase Data Disk** window, check **Scheduled Snapshot** and select the scheduled snapshot policy from the drop-down list, as shown below:

Scheduled Snapshots

Configure scheduled snapshot for the purchased cloud disk Recommended

██████|Tuesday|15:00|Permanent Retention
▼

↻

Policy Details
i

[Create Scheduled Snapshot Policy](#)

[Snapshot Policy](#)

Snapshot can help prevent the risk of data loss caused by mis-operations and viruses. Tencent Cloud currently provides each user with a free tier of 80 GB per region in China. For more information, please see [Snapshot Billing Overview](#).

i Note

For explanations of other parameter settings, please refer to [Creating Cloud Disks](#).

3. After the cloud disk is created successfully, it can be associated with the selected scheduled snapshot policy.

The cloud disk purchase page defaults to associating a scheduled snapshot policy.

1. On the **Cloud Disk** page, select **Cloud Disk Data Protection Configuration** in the upper right corner of the page.
2. In the pop-up **Cloud Disk Data Protection Configuration** window, check **Snapshot Protection for New Cloud Disks**

and select the available zone from the drop-down list, as shown below:

Cloud Disk Data Protection Configuration

Snapshot Protection for New Cloud Disks

CBS purchase page [Enable Scheduled Snapshot by Default](#)

If it's enabled, Scheduled Snapshot will be enabled by default when you purchase a new cloud disk. You can [disable it manually](#).

Please select the AZ.

Expiry/Overdue Protection

Enable for all cloud disks in this availability zone

You can use ["Display Settings"](#) to display this feature, so as to enable and disable protection easily.

[Save](#) [Close](#)

3. Click **Save** to complete the setup.

Note

When creating a cloud disk in the selected availability zone, **Scheduled Snapshot** will be selected by default, meaning a scheduled snapshot policy will be associated. You can deselect **Scheduled Snapshot** as needed.

Documentation

- [Creating Cloud Disk](#)
- [Scheduled Snapshots](#)

Cloud Disk Data Backup Point

Adjusting Data Backup Point Quota

Last updated: 2023-09-19 21:11:45

Scenario

For cloud disks that have enabled data backup points, Tencent Cloud will automatically create data backup points to back up their data. You can set the maximum number of data backup points to be retained (hereinafter referred to as "number of data backup points") when purchasing a new cloud disk, or modify and return the number of data backup points for existing cloud disks as needed.

This document provides an explanation of the costs associated with modifying the maximum number of data backup points retained on a cloud disk, as well as detailed operational steps.

Note

- The billing for data backup points is solely related to the size of the cloud disk and the maximum number of points retained, and is not associated with any other factors.
- When reducing or returning the number of data backup points, your existing data backups may be synchronously deleted. If you need to retain important backup point data, you can convert it into a snapshot. For more details, please refer to [Converting Data Backup Points into Snapshots](#).

Fees after the adjustment

Increasing the Number of Data Backup Points

Billing Rule

- For cloud disks in different billing modes:
 - **Monthly Subscribed Cloud Disks:** The number of data backup points is increased according to the remaining time of the lifecycle to make up the price difference between the new and old configurations. The specific situation is subject to the actual situation, which you can view on the payment page.
 - **Pay-as-you-go cloud disks:** The changes take effect immediately and billing commences based on the newly configured price.
- Billing rules for monthly subscribed cloud disks:

The price difference should be made up based on the number of days. Configuration upgrade fees = price difference for upgrade per month *number of months for upgrade* applicable discount.

 - Price difference for the upgrade per month: Unit price difference between the new and old configurations.
 - Number of months for upgrade: Upgrade fees are calculated by month based on the number of days.
 - Number of days for upgrade = resource expiration time - current time
 - Number of months for upgrade = number of days for upgrade / (365/12)
 - Applicable discount: The applicable discount is matched based on the number of months for upgrade as effective at the official website.

Note

- This operation doesn't affect the resource expiration time.
- This operation allows you to use vouchers and free credits for fees deduction.

Billing Example

Note

The prices listed below are for illustrative purposes only and do not reflect the actual prices on the official website. For the pricing of data backup points, please refer to [Cloud Disk Data Backup Point Pricing Overview](#).

Suppose a user purchased a 200GB monthly subscribed SSD cloud disk for a duration of one month in the Guangzhou region on May 1, 2022, without setting any data backup points. On May 5, 2022, the user adjusted the number of data backup points to 1. The cost of this configuration upgrade is calculated as follows:

- Price difference for the upgrade per month = $220 - 200 = 20$ CNY/month
- Number of days for upgrade = 27 days
- Applicable discount = 50% off

Final total upgrade cost = $20 \times (27 / (365/12)) \times 0.5 = 8.88$ USD

Reducing the number of data backup points/Returning data backup points

Billing Rule

- For cloud disks in different billing modes:
 - **Monthly Subscribed Cloud Disks:** Reducing the number of data backup points will result in a refund calculated based on the difference between the value of the remaining time in the lifecycle and the value of the new configuration. The actual situation may vary, and you can view it on the payment page.
 - **Pay-as-you-go cloud disks:** The changes take effect immediately and billing commences based on the newly configured price.
- Billing rules for monthly subscribed cloud disks:

The refund amount for downgrading = refund cost of the original specification – purchase cost of the new specification. For details on refund costs, please refer to [Refund Policy](#).

 - If the refund amount is greater than zero and the number of data backup points is reduced, the refund will be returned to your Tencent Cloud account using the original payment method.
 - If the refund amount is ≤ 0 , the number of data backup points will decrease, but no refund will be issued.
- If you use a discount or voucher for purchase, the discount amount and voucher will not be refunded.

Billing Example**Note**

The prices listed below are for illustrative purposes only and do not reflect the actual prices on the official website. For the pricing of data backup points, please refer to [Cloud Disk Data Backup Point Pricing Overview](#).

A user in the Guangzhou region purchases a new 200GB SSD cloud disk on May 1, 2022, with a duration of one month and one data backup point (the monthly subscription price is 210 RMB/month, and the pay-as-you-go price is 0.514 RMB/hour). On May 5, 2022, the number of data backup points is reduced to zero. The cost of this downgrade is calculated as follows:

- The original specification has been used for 4 days, the refund amount = $210 - (4 \times 24 \times 0.528) = 159.31$ CNY.
- The new specification's monthly price is 200 CNY/month. With 27 days remaining until the expiration date, the cost for purchasing the new specification would be: $200 \times (27/30) = 180$ CNY.
- Downgrade refund amount = $159.31 - 180 < 0$

At this point, reducing the number of data backup points will result in a refund amount of zero.

Instructions**Enable/Add Data Backup Points**

You can enable or increase the number of data backup points in the following ways:

Increasing the Number of Data Backup Points via the Console

Provided that the [usage restrictions](#) are met, you can increase the number of data backup points through the following methods:

1. Log in to the [Cloud Disk Console](#) and select the region where the cloud disk is located at the top of the page.
2. In the list, locate the cloud disk for which you need to increase the number of data backup points, and select **More > Enable Data Backup Points/Increase Data Backup Points** on the right side of its row.
3. In the pop-up "Enable/Add Data Backup Points" window, select the maximum number of data backup points you wish to retain.
4. Click **Next**, confirm the cost details, and select "I have read and agree to the Fee for Adjusting the Number of Data Backup Points".
5. Click **OK** to initiate the adjustment.

Increasing the Number of Data Backup Points via API

You can use the ModifyDiskBackupQuota interface to adjust the number of data backup points for a specified cloud disk.

Reducing/Returning the Number of Data Backup Points

You can reduce or return the number of data backup points in the following ways:

Reducing or returning the number of data backup points via the console

In compliance with the [Usage Restrictions](#), you can reduce or return the number of data backup points through the following methods:

1. Log in to the [Cloud Disk Console](#) and select the region where the cloud disk is located at the top of the page.
2. In the list, locate the cloud disk for which you want to reduce/return the number of data backup points, and select **More > Reduce/Return Data Backup Points** on the right side of its row.
3. In the pop-up "Reduce/Return Data Backup Points" window, select the maximum number of data backup points you wish to retain.
4. Click **Next**, confirm the cost details, and select "I have read and agree to the Fee for adjusting time-specific data backup points".
5. Click **OK** to initiate the adjustment.

Reducing or returning the number of data backup points via API

You can use the ModifyDiskBackupQuota interface to reduce or return the number of data backup points for a specified cloud disk.

Policy for eliminating and deleting existing backup points when reducing/returning the number of data backup points

When you reduce or return the number of data backup points, the operation may result in the updated number of data backup points being less than the current number of occupied data backup points. In this case, existing data backup points will be eliminated and deleted according to the following principles:

- Tencent Cloud will prioritize the deletion of automatically created data backup points. However, unless the data backup points are completely returned, at least one automatic backup point will be retained for Tencent Cloud to automatically create and rotate backups.
- If the number of automatic backup points still exceeds the maximum number of data backup points retained after downgrading, according to the above rules, the manually created backup points will be eliminated and deleted in the order of their creation time, from oldest to newest.

The specific examples are as follows:

Assuming you purchased a cloud disk on October 1st and configured the number of backup points to be 5, with no additional operations thereafter, the usage of data backup points on October 16th would be as follows:

1	2	3	4	5
October 11: A (Automatically created by Tencent Cloud)	October 12: B (Automatically created by Tencent Cloud)	October 13: C (Automatically created by Tencent Cloud)	October 14: D (Automatically created by Tencent Cloud)	October 15: E (Automatically created by Tencent Cloud)

On October 16th, you performed important operations on the data, so you deleted the backup points for October 11th and 12th, and manually created data backups for two time points: 08:00 and 18:00 on October 16th.

1	2	3	4	5
October 13: C (Automatically created by Tencent Cloud)	October 14: D (Automatically created by Tencent Cloud)	October 15: E (Automatically created by Tencent Cloud)	October 16: F (Manually created at 08:00)	October 16: G (Manually created at 18:00)

On October 17th, the status of data backup points occupancy during Tencent Cloud's automatic backup rotation is as follows:

1	2	3	4	5
October 14: D (Automatically created by Tencent Cloud)	October 15: E (Automatically created by Tencent Cloud)	October 16: F (Manually created at 08:00)	October 16: G (Manually created at 18:00)	October 16: H (Automatically created by Tencent Cloud)

At this point, you choose to reduce the number of backup points:

- Assuming the number of backup points is reduced from 5 to 3: automatic backup points are eliminated first, but at least the most recent one is retained, followed by the elimination of manual backup points. The retained backup points are: F, G, H.
- Assuming the number of backup points is reduced from 5 to 2: The automatic backup points are eliminated first, but at least the most recent one is retained, followed by the elimination of manual backup points. The retained backup points are: G, H.
- Assuming the number of backup points is reduced from 5 to 1: The system will prioritize eliminating automatic backup points but will retain at least the most recent one, followed by manual backup points. The remaining backup point will be: H.
- Assuming the number of backup points is reduced from 5 to 0: no data backup points are retained.

Using Data Backup Point to Restore Cloud Disk Data

Last updated: 2023-09-22 09:33:30

Scenario

The Cloud Disk Data Backup Point provides a continuous data backup service for the cloud disk. Once you have a data backup point, you can restore the cloud disk data to the historical data backup point to mitigate data loss caused by viruses, intrusions, and misoperations. This document explains how to use an existing data backup point to restore cloud disk data to the time of that backup.

Supports and Limits

Keep the following in mind when using the data backup point to restore data:

- The Cloud Virtual Machine (CVM) will need to be shut down, which will impact business continuity. It is advisable to prepare in advance and perform the operation during off-peak business hours.
- The operation will clear cloud disk data after the data backup point. Evaluate the impact in advance.

Instructions

1. Log in to the [Cloud Disk Console](#) and select its region at the top of the page.
2. Find the target cloud disk in the list and click **ID/Name** to enter its details page.
3. In the details page, select the **Auto data backup** tab, locate the data backup point you wish to use, and click **Rollback auto backup** on the right side of the row, as shown in the figure below:

The screenshot shows the 'Auto data backup' tab in the Cloud Disk Console. At the top, it displays 'Used backup points/Maximum backup points retained' as '1/1'. Below this is a table with columns: ID/Name, Historical data time, Status, Disk Attribute, Associate Disk, Attach Disk to Instance, and Operation. A single row is visible with the following data:

ID/Name	Historical data time	Status	Disk Attribute	Associate Disk	Attach Disk to Instance	Operation
dbp-jxh2uqwh	[Progress bar]	Normal	Data Disk 200GIB	disk-q9g236rs	-	Delete Rollback auto backup More ▾

At the bottom of the table, it shows 'Total items: 1' and a pagination control set to '20 / page' with '1 / 1 page'.

4. In the pop-up **Restore Backup Point Data** window, confirm the information. If the cloud server where the current cloud disk is mounted is running, you also need to check the shutdown prompt.
5. After confirming, click **OK** to initiate the restoration process.

Converting Data Backup Point into Snapshot

Last updated: 2023-09-22 09:37:39

Scenario

The lifecycle of the Cloud Disk Data Backup Point is entirely tied to the lifecycle of the cloud disk instance. Upon expiration or proactive return of the cloud disk, the data backup point will also end its lifecycle. You can convert the data backup point of important dates into [Snapshots](#) in advance, making their lifecycle independent and allowing for long-term retention of crucial data backups.

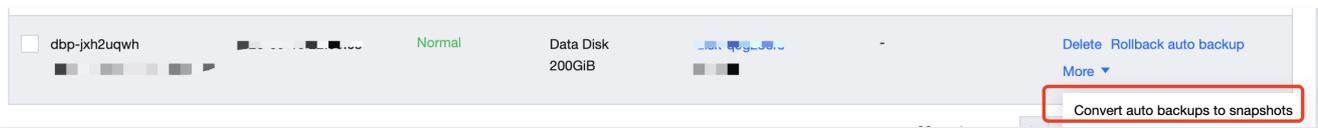
This document describes how to convert a data backup point into a snapshot in the console.

Supports and Limits

- Snapshots have been [commercialized](#). Converting a data backup point into a snapshot will generate a new snapshot resource, which will be billed as a regular snapshot. For billing details, please refer to [Snapshot Pricing Overview](#).
- After being converted into a snapshot, the data backup point will be deleted automatically and no longer occupy the quota.

Instructions

1. Log in to the [Cloud Disk Console](#) and select its region at the top of the page.
2. Find the target cloud disk in the list and click **ID/Name** to enter its details page.
3. In the details page, select the **Data Backup Point** tab, find the data backup point to be converted, and select **More > Convert auto backups to snapshots** on the right side of its row, as shown below:



4. In the pop-up **Convert auto backups to snapshots** window, enter the **Snapshot Name**.
5. After confirming the information, click **OK** to start the conversion. Once completed, you can go to the [Snapshot List](#) page to view the results.

Manual Creation of Data Backup Points

Last updated: 2023-09-20 09:53:51

For cloud disks that have enabled the data backup point, Tencent Cloud will automatically create data backup points for data backup. Under normal circumstances, you do not need to worry as Tencent Cloud will automatically carry out data backup and rotation. If you wish to manually create a data backup for the cloud disk at a specified time, you can refer to this document for manual creation operations.

Notes

- At least one data backup point must be retained for Tencent Cloud's automatic backup and rotation. Therefore, manual creation of backup points is only possible when the number of data backup points is configured to be **more than one**.
- If the current number of used backup points has reached the maximum retention number, it is not possible to manually create new data backup points. Deletion and release of the used backup points are required before manual creation can be performed.

For instance, if a user purchases a cloud disk A with five backup points on October 1st, under normal circumstances, the usage of data backup points on October 16th would be as follows:

1	2	3	4	5
October 11: A (Automatically created by Tencent Cloud)	October 12: B (Automatically created by Tencent Cloud)	October 13: C (Automatically created by Tencent Cloud)	October 14: D (Automatically created by Tencent Cloud)	October 15: E (Automatically created by Tencent Cloud)

At this point, if you wish to manually create a new data backup point, based on the aforementioned notes:

- The number of data backup points is more than one, which is satisfactory.
- The current number of used backup points has reached the maximum retention number, which is not satisfactory.

Hence, you need to delete at least one backup point before you can manually create one. Assuming the backup point automatically created by Tencent Cloud on October 14th is deleted, the usage of data backup points would be as follows:

1	2	3	4
October 11: A (Automatically created by Tencent Cloud)	October 12: B (Automatically created by Tencent Cloud)	October 13: C (Automatically created by Tencent Cloud)	October 15: E (Automatically created by Tencent Cloud)

At this point, you can manually create a data backup point. After successful creation, the usage of backup points would be as follows:

1	2	3	4	5
October 11: A (Automatically created by Tencent Cloud)	October 12: B (Automatically created by Tencent Cloud)	October 13: C (Automatically created by Tencent Cloud)	October 15: E (Automatically created by Tencent Cloud)	October 16th (Manually created)

Instructions

Manually creating data backup points via the console

When the [Notes](#) are met, you can increase the number of data backup points through the following methods:

- Log in to the [Cloud Disk Console](#) and select the region where the cloud disk is located at the top of the page.
- Locate the cloud disk for which you need to manually create a data backup point in the list, and click on the cloud

disk ID to enter the details page.

3. On the **Data Backup Point** tab, click **Create New**.

4. In the pop-up **Create Backup Point** dialog box, confirm the information, optionally enter a backup point name, and click **OK** to create a data backup point.

Manually Creating Data Backup Points via API

You can use the [CreateDiskBackup](#) interface to manually create new data backup points for the specified cloud disk.

Deleting Data Backup Point

Last updated: 2023-09-22 09:41:47

Scenario

For cloud disks that have a maximum number of data backup points configured, Tencent Cloud will automatically create and manage these data backup points to ensure the rotation of backup data updates. For more information, see [Cloud Disk Data Backup Point](#).

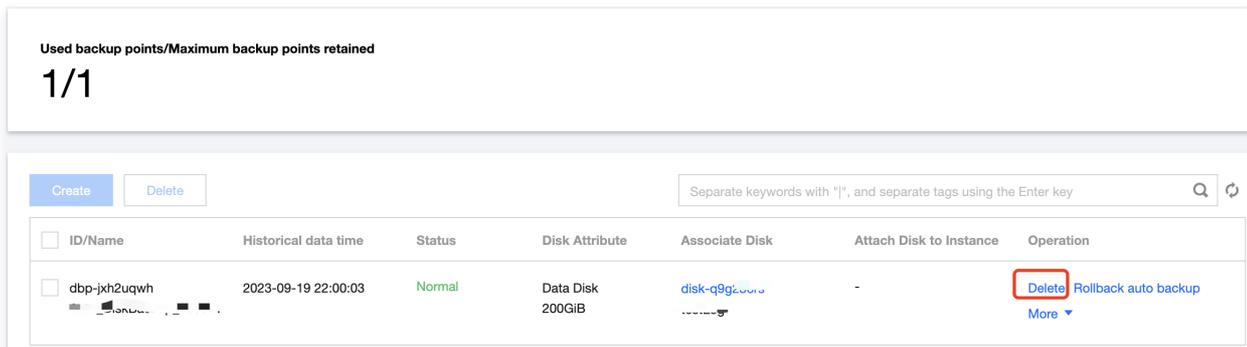
For cloud disks that have a maximum number of data backup points configured, retaining these data backup points will not incur additional charges. If you wish to delete data backup points manually, you can follow the instructions in this document.

Note

Deleting the data backup point cannot be undone, and deleted data backups cannot be recovered.

Instructions

1. Log in to the [Cloud Disk Console](#) and select its region at the top of the page.
2. Find the target cloud disk in the list and click **ID/Name** to enter its details page.
3. On the details page, select the **Data Backup Point** tab, locate the data backup point you wish to delete, and click **Delete** on the right side of the corresponding row. As shown in the figure below:



4. In the **Delete Data Backup Point** pop-up window, click **Confirm** to delete.

Managing Snapshots

Creating Snapshot

Last updated: 2023-09-22 09:48:43

You may refer to the following videos and documents to create snapshots for preserving data at a specified moment on the cloud disk:

[Watch video](#)

Scenario

You have the ability to create snapshots for your cloud disk, thereby preserving data at a specific moment in time. Tencent Cloud employs an incremental approach to snapshot creation, meaning that only new data changes compared to the last snapshot are created. As such, in situations where data changes are minimal, snapshot creation can be completed in a relatively short time. Despite snapshots being created incrementally, deleting a snapshot will not affect your use of any snapshot data, and any undeleted snapshots can restore the cloud disk to the state of that snapshot.

You can create a snapshot under any state of the cloud disk, but the snapshot can only preserve data that has been written at the current time point. If an application or process is writing data, this portion of data may not be saved to the snapshot created at that time point. Depending on your actual business situation, you can choose to temporarily stop all writing and promptly create a snapshot, or first [dismount](#) the cloud disk from the cloud server, create a snapshot, and then [remount](#) it, in order to obtain a snapshot with complete data.

Preparations

- You have successfully [created a cloud disk](#).
- The number and total capacity of your snapshots within the current region have not reached the maximum. For detailed information, please refer to [Snapshot Usage Limits](#).

Supports and Limits

- The size of the created snapshot may be larger than the actual data volume of the disk. For specific reasons, please refer to [Why does the used disk capacity displayed in the file system differ from the snapshot size?](#)
- Snapshots only retain data that has been written to the disk at that moment, and do not preserve data in memory that has not been written to the disk at that time (for example, files under the `/run` directory in the Linux system). It is strongly recommended that you shut down or ensure that memory data has been written to the disk and disk read/write operations are paused before creating a snapshot. Actions should be taken from the following two aspects.

Database level

For database-type businesses, it is recommended to first lock all tables in the database in a read-only state to prevent new data from being written during snapshot creation, which could result in the new data not being captured by the snapshot. This article uses the MySQL database as an example, performing the following operations:

1. Execute the `FLUSH TABLES WITH READ LOCK` command to close all open tables and lock all tables in all databases with a global read lock, as shown in the figure below:

```
mysql> flush tables with read lock;  
Query OK, 0 rows affected (0.01 sec)
```

2. Create a snapshot for the cloud disk.

3. Execute the `UNLOCK TABLES` command to remove the lock. As shown in the figure below:

```
mysql> unlock tables;  
Query OK, 0 rows affected (0.00 sec)
```

System level

From a system perspective, data is typically first stored in the memory buffer during processing, and then written to the cloud disk at the appropriate time to enhance system performance. Therefore, when creating a snapshot, data in the buffer that has not yet been written to the cloud disk cannot be written into the snapshot or recovered from the snapshot, which can affect data consistency.

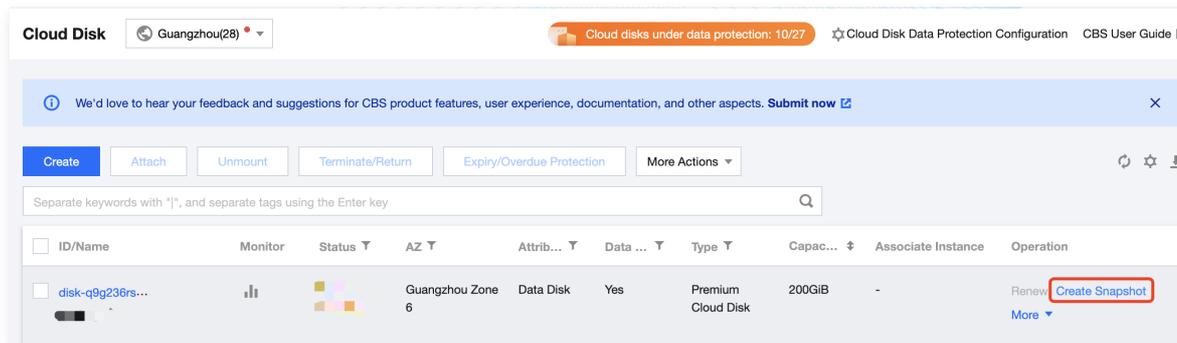
You can execute the `sync` command to force immediate writing of data from the file system memory buffer to the cloud disk, and avoid writing new data before creating a snapshot. If no error message is returned after executing the command, it means that the data in the cache has been written to the cloud disk. As shown in the figure below:

```
ubuntu@VM-30-151-ubuntu:~$ sync
ubuntu@VM-30-151-ubuntu:~$
```

Instructions

Creating a snapshot via the console

1. Log in to the [Cloud Disk Console](#).
2. Click on **Create Snapshot** in the **Operation** column of the row where the target cloud disk is located, as shown in the image below:



3. In the pop-up **Create Snapshot** dialog box, input the snapshot name and click **OK**.

Create Snapshot ✕

1. A snapshot captures only data in the cloud disk but not the data in the memory. To ensure that the snapshot includes the complete data, please do the following before creating a snapshot:

- Database service: flush and lock the table data to save it to the disk
- File system: Perform synchronization to save the memory data to the disk

2. Please note that snapshot service has been commercialized. You need to pay for the snapshots generated. The total capacity of your snapshots in the current region: 16326.88 GiB, available free quota: 80 GiB. For more information about snapshots, see [Snapshot Overview](#) and [Commercialization FAQs](#)

Cloud Disk ID: disk-██████████

Cloud Disk Name: ██████████

Disk Capacity: 150 GiB

Disk Type: Balanced SSD

Tag ⓘ [+ Add](#)

Snapshot Name:
You can enter 60 more characters.

Permanent Retention: Yes No

[OK](#) [Close](#)

Creating a snapshot via an API

You can create a snapshot using the CreateSnapshot API. For specific operations, please refer to [CreateSnapshot](#).

Rolling back Data from Snapshot

Last updated: 2023-09-22 09:53:08

Scenario

When a data error occurs or data is lost due to a change, you can roll back the snapshot data to the corresponding cloud disk, so that data in the cloud disk can be restored to the status when the snapshot is created.

- The snapshot data can be rolled back only to the source cloud disk but not other cloud disks.
- You can roll back the data in the following scenarios:
 - If the source cloud disk is in **To be attached** status (that is, it has not been attached to a CVM instance), you can directly perform the rollback.
 - If the source cloud disk has been attached to a CVM instance, you can perform the rollback when the CVM instance is shut down.

Instructions

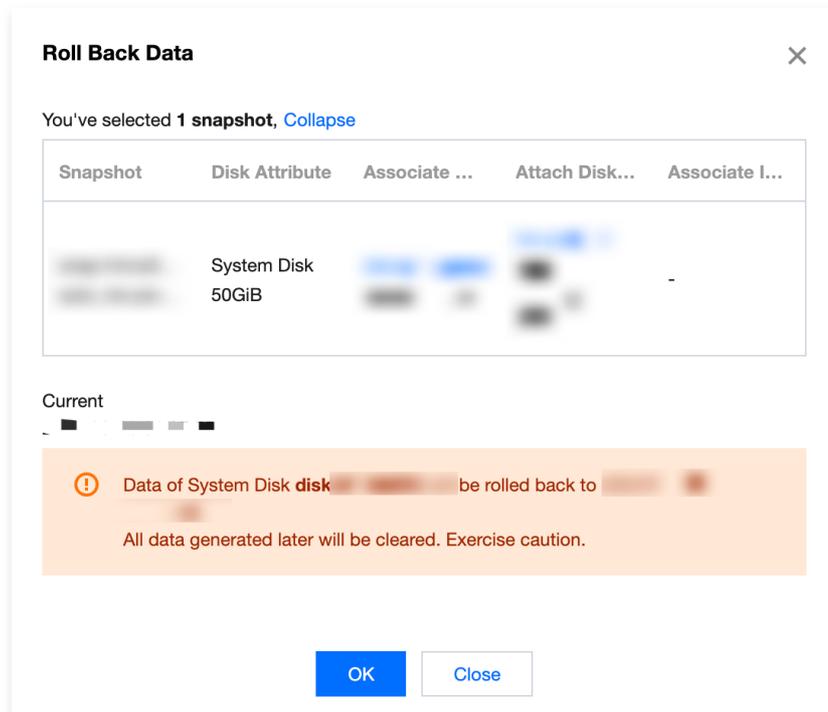
Rolling back the data in the console

1. Log in to the [Snapshot List](#) page.
2. Click **Rollback** in the row of the target snapshot.

Note

The data of the source cloud disk will be rolled back to the status when the snapshot is created, and all the data after this time will be cleared.

3. On the **Roll Back Data** page, confirm the rollback information and click **OK** to start the rollback. As shown below:



Rolling back the data via an API

You can use the ApplySnapshot API to perform snapshot rollback. For detailed operations, please refer to [Rollback Snapshot](#).

Creating Cloud Disks Using Snapshots

Last updated: 2023-09-19 21:22:50

Scenario

Making snapshot is an important method for data sharing and migration. Cloud disks created using a snapshot own all data in the snapshot. You can use a snapshot to create a cloud disk whose capacity is greater than or equal to that of the snapshot.

- When creating a data disk of the same size using a snapshot, the new data disk does not need to be initialized. Simply [attach](#) it and right-click  **Disk Management**. After setting it **online** in "Disk Management", it can be read and written normally on the cloud server.
- When creating a data disk with a capacity larger than the snapshot using a snapshot, the system only completes the block device-level disk expansion, and does not implement the expansion of the file system or the automatic conversion of the partition form. After the new data disk is [attached](#), it can only use the file system and data of the source snapshot, and cannot directly use the new disk space. You need to manually expand the file system or even convert the partition form.

For instance, when using a snapshot of a data disk with an MBR partition and a capacity of 1TB to create a new 3TB data disk. Since the maximum disk space supported by MBR is 2TB, you need to format the data disk and re-partition it using GPT, this operation will **delete the original data**. Therefore, **please operate carefully according to your actual needs**.

This document guides you to create a cloud disk through a snapshot on the [Snapshot List](#) page. In addition, you can also specify the corresponding snapshot to create a cloud disk by configuring the "Snapshot" parameter when [creating a cloud disk](#).

Instructions

Creating a cloud disk with a snapshot in the console

1. Log in to the [Snapshot List](#) page.
2. In the row of the target snapshot, select **More > Create Cloud Disk**.
3. In the **Purchase Data Disk** pop-up window, configure the following parameters:

Configuration Item	Description
Availability Zones	Required. The availability zone of the cloud disk cannot be modified after the cloud disk has been created.
Cloud Disk Type	Required. For information on cloud disk types, see Cloud Disk Types .
Capacity	Required. For the capacity and specifications of the cloud disk, please refer to Performance Metrics . When creating a cloud disk from a snapshot, the capacity cannot be less than the size of the snapshot. If you do not specify the capacity of the cloud disk, the capacity will default to the same size as the snapshot.
Snapshot	Optional parameter. When creating a cloud disk using a snapshot, you need to check "Create Cloud Disk with Snapshot" and select the snapshot you need to use. <ul style="list-style-type: none"> • You can adjust the capacity of a cloud disk created from a snapshot to be greater than the default value (snapshot size). • When you create a cloud disk using a snapshot, the disk type is the same as that of the snapshot's source disk. You can change the disk type.
Disk Name	Optional.

	<p>Supports up to 20 characters, starting with a letter (upper or lower case) or Chinese character. It can consist of letters (upper or lower case), Chinese characters, numbers, and special symbols <code>._-`</code>. It can be modified after the cloud disk is created.</p> <ul style="list-style-type: none"> • Creating a Single Cloud Disk: The disk name is the name of the cloud disk. • Batch creation of cloud disks: When creating multiple cloud disks at once, the disk name is the prefix of the cloud disk name. The final cloud disk name is composed of "Disk Name_Number", ranging from "Disk Name_0" to "Disk Name_49".
Project	<p>Required.</p> <p>When creating a cloud disk, you can set the project to which the cloud disk belongs. The default project is "Default Project".</p>
Tag	<p>Optional.</p> <p>You can bind tags to the cloud disk when creating it. Tags are used to identify cloud resources, and you can categorize and search for cloud resources through tags. For more detailed information about tags, please refer to the Tag Product Documentation.</p>
Billing	<p>Required.</p> <p>The cloud disk supports the following two billing types:</p> <ul style="list-style-type: none"> • Monthly Subscription. If you choose this mode, you must set the "Validity Period". • Pay as you go.
Scheduled Snapshot	<p>Optional.</p> <p>You can choose to schedule backups when creating a cloud disk. Based on the created scheduled snapshot policy, you can regularly create snapshots for this cloud disk. For more detailed information about scheduled backups, please refer to Scheduled Snapshots.</p>
Purchase quantity	<p>Optional.</p> <p>The default quantity is "1", indicating the creation of only one cloud disk. Currently, up to 50 cloud disks can be created in bulk.</p>
Duration	<ul style="list-style-type: none"> • If the "Billing Mode" is set to Monthly Subscription, this parameter is mandatory. The selectable time range is from 1 month to 5 years. • If the "Billing Mode" is set to pay-as-you-go, this parameter is not applicable.
Auto-Renewal	<ul style="list-style-type: none"> • If the "Billing Mode" is set to Monthly Subscription, this is an optional parameter. When auto-renewal is selected, the device will be automatically renewed on a monthly basis after expiration, provided that the account balance is sufficient. • If the "Billing Mode" is set to pay-as-you-go, this parameter is not applicable.

4. Click OK.

- If the **Pay-as-You-Go** option is selected for "Billing Mode", then the creation is complete.
- If you select **Monthly Subscription** for "Billing Mode", you will be directed to the "Review Information" page.

a. After confirming the specifications are correct, decide whether to use a voucher based on the actual situation, then click **Submit Order**.

b. Complete the payment.

5. You can view the created cloud disks on the [Cloud Disk List](#) page. The newly created elastic cloud disk is in the "To Be Mounted" state. You can refer to [Mounting Cloud Disks](#) to attach the cloud disk to a cloud server within the same availability zone.

Using an API to create a cloud disk from snapshot

You can use the CreateDisks interface to create a snapshot. For specific operations, please refer to [Creating a Cloud Disk](#).

Scheduled Snapshots

Last updated: 2023-09-22 10:33:12

Overview

Tencent Cloud's [Snapshot](#) feature now includes a **Scheduled Snapshot** function, which allows developers to flexibly set backup task policies. It is recommended to adopt different scheduled snapshot policies for different businesses, as suggested in the following table:

Scenario	Snapshot Policy	Recommended Snapshot Retention
Core product/service	Use scheduled snapshots, with the policy set to once per day.	7 to 30 days
Non-core and non-data product/service	Use scheduled snapshots, with the policy set to once per week.	7 days
Archive	Scheduled snapshot is not required. You can create snapshots manually whenever needed.	One month to several months
Test	Scheduled snapshot is not required. You can create snapshots manually whenever needed.	Deleted after being used

Policy description

The following table describes the content and features of scheduled snapshot policies, which help you better use snapshots in your businesses.

Item	Note
Objects	All cloud disks, including system disks and data disks.
Execution policy	The point in time for scheduled snapshot creation can be accurate to every hour or every day. A scheduled snapshot policy is valid permanently after being set. If you modify a scheduled snapshot policy, it takes effect immediately.
Scheduled Deletion (Important)	Scheduled snapshots can be terminated periodically. After you set a snapshot lifecycle (1-30 days), scheduled snapshots are automatically deleted upon expiration, which reduces the backup costs. If you do not set a scheduled termination policy, scheduled snapshots will be stored permanently.
Batch	You can apply the same scheduled snapshot policy to multiple cloud disks.
Naming rule	The naming convention for automatic snapshots is snap_yyyyMMdd_HH, where yyyyMMdd represents the current date and HH represents the hour. You can also modify snapshot names. For instance, snap_20180418_11 denotes an automatic snapshot created at 11 o'clock on April 18, 2018.
Lifecycle (Critical)	Snapshot lifecycles vary depending on the snapshot creation method: <ul style="list-style-type: none"> Manually created snapshots are by default set for long-term storage, and as long as the account balance is sufficient, they can be retained indefinitely. Scheduled snapshots, according to the creation rules, can be set to be regularly deleted at a specific time, or they can be configured for long-term retention.
Snapshot conflict	Scheduled snapshots do not conflict with custom snapshots in use. However, they may conflict with each other on the creation time. <ul style="list-style-type: none"> You cannot create a custom disk snapshot until the automatic disk snapshot has finished being created, and vice versa.

	<ul style="list-style-type: none"> If the disk data size is large and the snapshot creation time exceeds the interval between two automatic snapshot time points, the next automatic snapshot will be skipped. For example, if a user sets 9:00, 10:00, and 11:00 as automatic snapshot time points, and the 9:00 snapshot takes 70 minutes to complete (finishing at 10:10), the 10:00 automatic snapshot will be skipped, and the next snapshot time point will be 11:00.
Snapshot quota	Each disk has a certain snapshot quota. If the number of snapshots for a disk reaches the quota limit, the automatic snapshot task will be suspended and blocked. The snapshot quota is primarily designed to prevent developers from forgetting about an automatic snapshot policy, which could lead to an endless increase in storage costs.
ASP	Indicates the scheduled snapshot policy, that is, Auto Snapshot Policy.
ASP quota	Under a single Tencent Cloud account, a maximum of 30 ASP policies can be set for each region. A single ASP can be associated with up to 200 disks.
Retention period	<ul style="list-style-type: none"> The console displays the repossession countdown for scheduled snapshots. You can manually change the retention period of scheduled snapshots to permanent. Manually created snapshots are permanently stored.
ASP pause	The ASP automatic snapshot policy provides a manual pause function. Once paused, no new snapshots will be automatically created. However, the lifecycle of already generated automatic snapshots is not affected by the pause function and will still be periodically destroyed or stored long-term according to the set rules.
Operations log	Show the creation process of all scheduled snapshots, same as that of manually created snapshots.

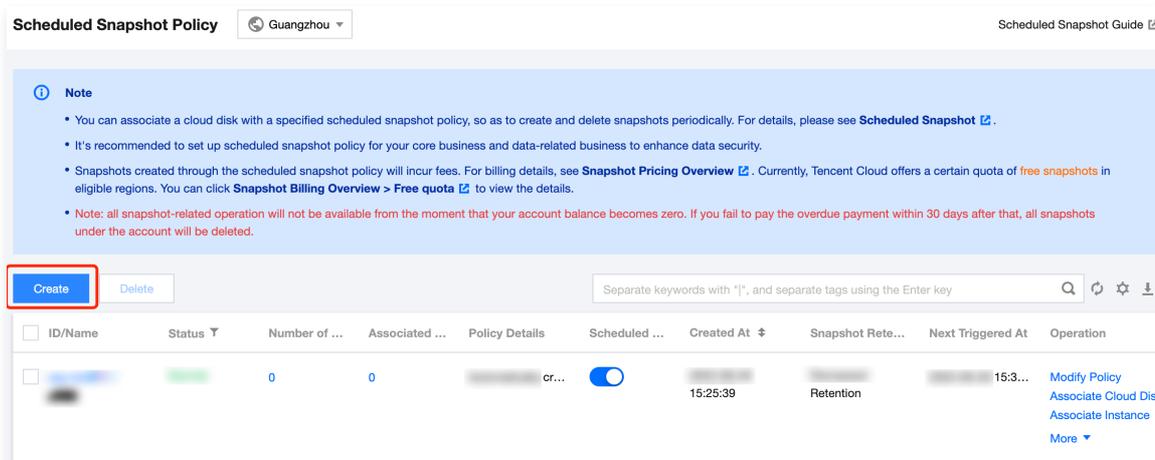
Directions

Creating a scheduled snapshot policy

Note

A single Tencent Cloud account can set a maximum of 30 scheduled snapshot policies in each region.

1. Log in to the [Scheduled Snapshot Policy](#) page.
2. Select a region.
3. Click **Create**, as illustrated below:



4. In the **Create Snapshot Policy** page, configure the following parameters and click **OK**, as shown in the figure below:

Configuration Item	Description

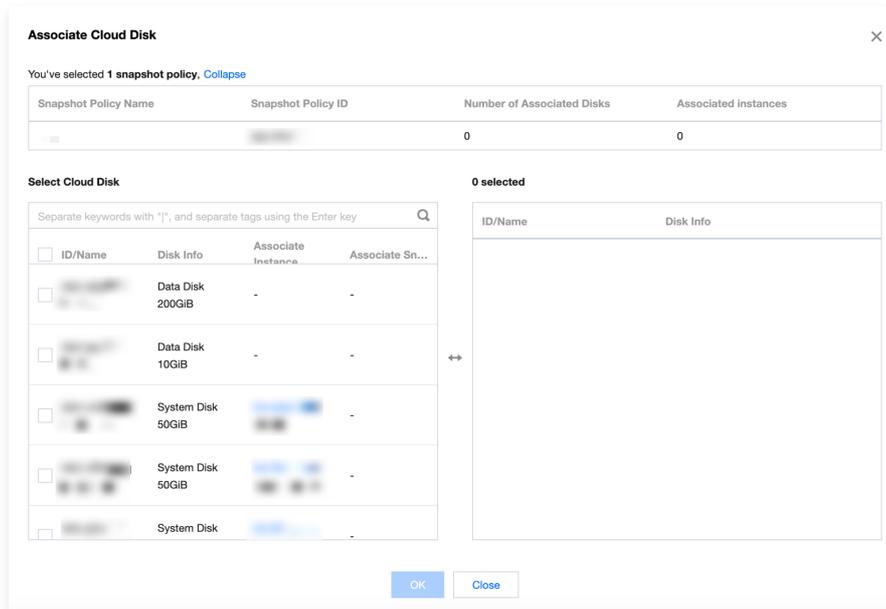
Name	Required. The name of a scheduled snapshot policy. It can contain up to 60 characters.
Region	This is a required parameter. It cannot be modified on the current page. For specific setting methods, please refer to Step 2 .
Backup Day	Required. The day to execute the scheduled snapshot, with selectable range: Sunday – Saturday.
Backup Time	Required. The time to automatically create snapshots. Value range: 00:00 – 23:00. Due to server performance issues, the actual scheduled snapshot creation time may be slightly different from the time set here. In this case, data in a snapshot is based on the actual creation time.
Recommended Snapshot Retention	Required. <ul style="list-style-type: none"> Automatically deleted after a fixed number of days, with a selectable range of 1 – 365 days. The default retention period is 30 days.

Associating cloud disks

Note

A scheduled snapshot policy can be associated with a maximum of 200 cloud disks.

1. Log in to the [Scheduled Snapshot Policy](#) page.
2. Select a region.
3. Click on **Associate Cloud Disk** in the row of the target policy.
4. On the **Associate Cloud Disk** page, select the cloud disks you wish to associate, as shown below:



5. Click **OK**.

Enable/Disable scheduled snapshot policies

1. Log in to the [Scheduled Snapshot Policy](#) page.
2. Select a region.
3. Find the target policy and click the button in the **Scheduled Snapshot** column to enable or disable the scheduled snapshot policy.

Modifying scheduled snapshot policies

1. Log in to the [Scheduled Snapshot Policy](#) page.
2. Select a region.

- Click on **Modify Policy** in the row of the target policy.
- On the **Edit Snapshot Policy** page, modify the relevant parameters (for parameter descriptions, see [Step 4](#)) and click **OK**.

Deleting scheduled snapshot policies

- Log in to the [Scheduled Snapshot Policy](#) page.
- Select a region.
- You can delete a scheduled snapshot policy in the following ways:
 - Single deletion: Select **More > Delete** on the right side of the target policy row.
 - Batch deletion: Select the scheduled snapshot policies you want to delete and click **Delete** at the top of the list.

Turning automatic snapshots into long-term stored snapshots

Note

If the **Snapshot Retention Period** in the automatic snapshot policy is set to **Permanent**, there is no need to perform the following operations on the snapshots generated by this policy.

- Log in to the [Snapshot List](#) page.
- Select a region.
- Click the ID of the target automatic snapshot.
- In the details page, click on **Modify expiry time** to set automatic snapshots to be permanently retained, as shown in the figure below:

The screenshot shows the 'Basic Info' tab of a snapshot details page. The 'Retention Period' is currently set to 'Automatically delete in 4 days', and a red box highlights the 'Modify expiry time' button next to it. The 'Associated Images' section is empty, showing 'No data yet'.

- Returning to the snapshot list, you can see that the **retention time** of this snapshot has been changed to **permanent retention**.

The screenshot shows a list of snapshots. The 'Retention' column for the selected snapshot is highlighted in red and shows 'Permanent Retention'. Other columns include 'Normal', 'Data Disk', '200GiB', and '0/50'.

Cross-region snapshot replication

Last updated: 2023-09-22 10:40:25

Cloud disk snapshots now support cross-region replication. This feature allows you to conveniently migrate data and services to other regions, and you can also build your business's cross-region disaster recovery system based on this cross-region replication capability.

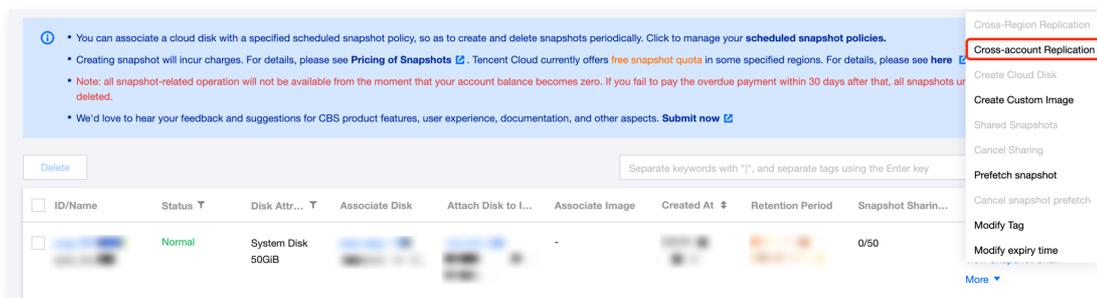
The snapshot cross-region replication is currently in the beta testing phase. You can [apply](#) to use Tencent Cloud's cloud disk snapshot cross-region replication.

Usage Limits

- **Beta Application:** Snapshot cross-region replication is currently in the beta phase, and an [application](#) is required to use this feature.
- **Supported Regions:** Please refer to [Regions and Availability Zones](#).
- **Financial Special Zone:** Replication is only supported between financial special zones.
- System disk snapshots are not supported.

Instructions

1. Log in to the [Snapshot List](#) page.
2. Click on **More > Cross-account Replication** in the **Operation** column of the target snapshot row.



3. Configure the following parameters:
 - **New Snapshot Name:** Enter the name for the new snapshot, which can contain up to 60 characters. This is an optional parameter.
The default naming will include the source snapshot ID and region information, formatted as Copied <Source Snapshot ID> from <Source Snapshot Region>, for example, Copied snap-oi5spwt2 from ap-shanghai.
 - **Region:** Select the target region to which you need to replicate, this is a mandatory parameter.
Check the snapshot quota and geographical restriction when you select the region.
4. Click **OK** to initiate replication. The source snapshot will display a status prompt, and a new snapshot will be added in the target region.
5. Once the replication is completed, you can view the new snapshot in the snapshot list of the target region.

Note

The source snapshot cannot be deleted during the cross-region replication of the snapshot.

During the process of cross-region replication:

- **Source Snapshot Status:** You can view it in the status bar of the source snapshot row in the [Snapshot List](#) of the source region.
- **Status of target snapshot:** You can view it by going to the snapshot list page of the target region.

Deleting Snapshots

Last updated: 2023-09-19 21:24:44

Scenario

When there is no need to use the snapshot again, you can delete the snapshot to release virtual resources.

Notes

- When you delete a snapshot, only the data exclusive to the snapshot will be deleted, and the cloud disk for which the snapshot is created will not be affected.
- You can use a snapshot to recover a cloud disk to the data status when the snapshot is created. Deleting a snapshot created earlier for a cloud disk will not affect the continued use of snapshots created later.
- If the snapshot has an associated image, please [delete the image](#) before deleting the snapshot.
- **When a snapshot is deleted, all data within the snapshot will be simultaneously erased and cannot be retrieved. Once deleted, snapshots cannot be restored. Please proceed with caution.**

Instructions

Deleting a snapshot via the console

1. Log in to the [Snapshot List](#) page.
2. You can delete snapshots in the following ways:
 - a. **Single deletion:** Click **Delete** in the row of the snapshot to be deleted.
 - b. **Batch deletion:** Select all snapshots to be deleted (ensure that the snapshots are not in task status), and click **Delete** above the list.
3. Click **OK**.

Using API to delete snapshots

You can use the DeleteSnapshots API to delete a snapshot. For specific operations, please refer to [Delete Snapshot](#).

Shared Snapshots

Last updated: 2023-09-26 14:39:57

Scenario

Snapshot Sharing involves sharing an already created [Data Disk Snapshot](#) with **other users**. Users can easily share snapshots with each other and utilize the acquired snapshots to swiftly create and mount cloud disks.

Supports and Limits

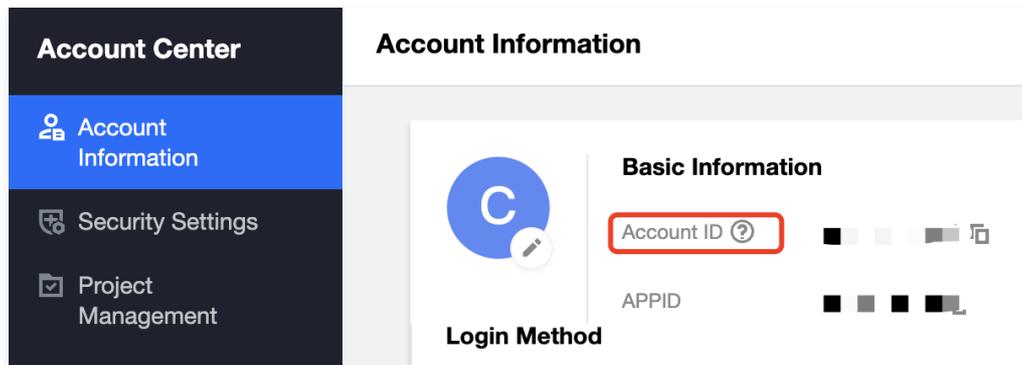
- Each snapshot can be shared with a maximum of 50 users.
- Snapshots created from system disks and encrypted cloud disks do not support sharing.
- **Snapshots that have been shared with other users must be unshared before they can be deleted. You can click on the target snapshot ID to enter the snapshot details page and cancel the share.**
- Snapshots can be shared within the same region of the recipient's account. If you need to share to a different region, you must first replicate the snapshot to the corresponding region before sharing. Please refer to [Cross-Region Snapshot Replication](#) for more details.
- Sharing of acquired shared snapshots with other users is not supported.
- The Snapshot Sharing feature is free of charge, and the recipient of the shared snapshot is not charged for it.

Instructions

Obtain the unique ID of the destination account

Tencent Cloud identifies shared snapshots through the unique ID of the root account of the recipient. That is, you first need to obtain the unique ID of the root account of the Tencent Cloud account used by the other party. The acquisition method is as follows:

1. Log in to the [CVM console](#).
2. Click the account name in the upper right corner and select **Account Information**.
3. On the **Account Information** management page, view and record the **Account ID** of the primary account.



4. Send his/her account ID to you.

Implementing Snapshot Sharing

You can share snapshots with other users in the following two ways:

Sharing a Snapshot via the Console

1. Navigate to the [Snapshot List](#) page.
2. Select **More > Shared Snapshots** on the right of the row where the snapshot to be shared is located.



3. In the **Shared Snapshots** pop-up window, input the unique ID of the destination account.
4. Click **OK** to share the snapshot.
5. Inform the recipient to navigate to the [Snapshot List](#) page and select the **Shared Snapshots** tab to view the shared snapshots. If you need to share with multiple users, please repeat the above steps.

Sharing Snapshots via API

You can use the `ModifySnapshotsSharePermission` API to share snapshots. For more details, refer to [Modify Snapshot Share Information](#).

Multi-mounting of Cloud Disk

Last updated: 2023-09-26 14:42:48

During the creation of a cloud disk, the multi-mount feature can be enabled. A cloud disk with the multi-mount feature enabled is referred to as a shared cloud disk. A shared cloud disk can be simultaneously mounted to multiple CVM instances within the same availability zone, supporting concurrent read and write access to the same cloud disk by multiple CVM instances.

Note:

The multi-mount feature is currently in beta testing. If you wish to utilize this feature, please contact your account manager or [submit a ticket](#) to apply.

Reminder:

1. Directly mounting a shared cloud disk to multiple cloud servers does not achieve traditional file sharing functionality. It is recommended to use **File Storage** for multi-instance file sharing.
2. To share files between multiple cloud servers, a dedicated shared file system or a similar cluster management system needs to be established.
3. Using common file systems such as ext2, ext3, ext4, xfs, ntfs, etc., cannot achieve data synchronization between multiple cloud servers, which may lead to data inconsistency issues.

Scenarios

The multi-mount feature of Cloud Block Storage is primarily used in key enterprise application scenarios that require support for High Availability (HA) cluster capabilities. To prevent business interruptions caused by the failure of one or more computing nodes, users can use shared cloud disks to ensure the continuity and high availability of business access provided by the cluster. Typical applications include high-availability database clusters commonly used in government, enterprise, and financial industries, such as Oracle RAC (Real Application Cluster). These require unified management of shared block storage through a cluster file system.

Supports and Limits

The essence of the multi-mount feature of cloud disks is to mount the same cloud disk to multiple cloud servers, similar to mounting a physical disk to multiple physical servers. Each server can read and write data from any area of the disk. If these servers do not have mutually agreed rules for reading and writing data, it may lead to interference between servers when reading and writing data, or other unpredictable errors.

Shared cloud disks do not come pre-installed with a cluster file system. Directly mounting a shared cloud disk to multiple cloud servers does not achieve file sharing functionality and poses a risk of data overwriting. To implement file sharing, conventional file systems (such as EXT, XFS, etc.) cannot be used. Instead, a shared file system or similar cluster management system needs to be set up, such as Windows MSCS cluster, Linux RHCS cluster, NFS cluster, etc. For typical Oracle RAC business scenarios, it is recommended to use Oracle ASM (Automatic Storage Management) to manage storage volumes and file systems. If a shared cloud disk is used without management through a cluster system, it may lead to the following issues.

Read/write conflict leading to data inconsistency.

When a shared cloud disk is simultaneously mounted to two cloud servers, Server A and Server B cannot perceive the storage space used by the other. Server A may allocate space on the cloud disk that has already been used by Server B, leading to allocation conflicts and potential data errors.

For instance, after formatting a shared cloud disk into an ext3 file system and mounting it to Cloud Server A and Cloud Server B, if Cloud Server A writes metadata to area a on the cloud disk at a certain moment, and Cloud Server B writes metadata to area a at the next moment, the data written by Cloud Server A will be overwritten, resulting in errors in the file system.

Data inconsistency due to caching

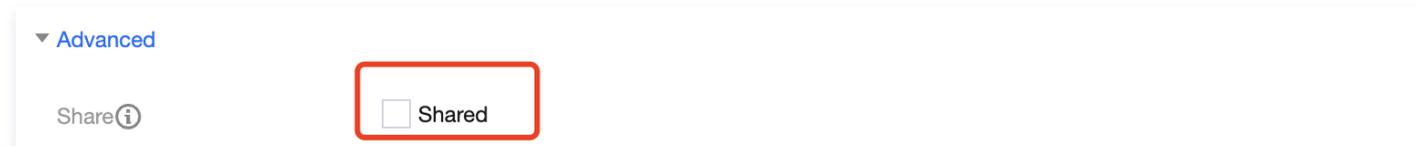
When a shared cloud disk is simultaneously mounted to two cloud servers, the traditional file system's file caching mechanism will not recognize the data modifications made to the cloud disk by the other cloud server. That is, the data modifications made to the shared cloud disk by cloud server B will not be readable by cloud server A.

How to Use

Note:

When using the multi-mount cloud disk, it is strongly recommended to combine it with the instance [spread placement group](#) feature to construct a highly available business architecture.

1. Once you have enabled the multi-mount capability, log in to the Tencent Cloud [Cloud Disk Console](#) and click **Create** to create a new cloud disk.
2. In the **Purchase Data Disk** window, expand **Advanced**, and check the **Shared** checkbox to create a cloud disk that supports multi-mounting.



3. Once successfully created, select the instance from the cloud disk list and click **Mount** to attach the cloud disk to multiple cloud server instances within the same availability zone.
4. Construct a shared file system or other cluster management systems on multiple cloud server instances to achieve data sharing.

Monitoring and Alarming

Last updated: 2023-09-20 09:54:53

Maintaining a robust monitoring environment for Cloud Block Storage (CBS) is a crucial aspect of ensuring high data reliability. You can use the [Tencent Cloud Observability Platform](#) to monitor **attached to CVM instances** CBS, view CBS metric data, and analyze and set alerts related to CBS. The [Tencent Cloud Observability Platform](#) collects raw data from the hard drives of running CVM instances and presents the data in easily readable charts, with statistical data saved by default for one month. You can observe the related conditions of the disk during different periods within this month to better understand usage, read-write operations, and other information.

You can obtain data through the [Tencent Cloud Observability Platform Console](#) or the [Tencent Cloud Observability Platform API](#). For more details, please refer to [Obtaining Monitoring Data for Specific Metrics](#) and [Obtaining Monitoring Views and Reports](#).

Currently, the Tencent Cloud Observability Platform provides the following monitoring metrics for CBS:

Description	Name	Calculation	Description	<p>	Statistical Granularity (Period)
DiskReadIOPS	Disk read IOPS	Average of read IOPS of the disk during the statistical period	Number of I/O reads from the disk to the memory every second	Time s	10s、60s、300s
Disk Read Traffic	Disk Read Traffic	Average of read throughput of the disk during the statistical period	Speed of data reads from the disk to the memory	KB/s	10s、60s、300s
DiskWriteIOPS	Disk write IOPS	Average of write IOPS of the disk during the statistical period	Number of I/O writes from the memory to the disk every second	Time s	10s、60s、300s
Disk Write Traffic	Disk Write Traffic	Average of write throughput of the disk during the statistical period	Speed of data writes from the memory to the disk	KB/s	10s、60s、300s
DiskAwait	Disk I/O Wait Time	Average of I/O wait time of the disk during the statistical period	Percentage of time during which the CPU is free and there are unfinished I/O requests during the statistical period	ms	10s、60s、300s
DiskSvcTm	Disk I/O Service Time	Average of service time of the disk during the statistical period	I/O service time	ms	10s、60s、300s
DiskUtil	Disk I/O Utilization	Average of I/O utilization of the disk during the statistical period	Percentage of time during which the disk has I/O operations (i.e., non-idle time)	%	10s、60s、300s

For detailed explanations of the monitoring metrics, please refer to the [Tencent Cloud Observability Platform Product Documentation](#).