

Cloud Block Storage Operation Guide Product Documentation





Copyright Notice

©2013-2024 Tencent Cloud. All rights reserved.

Copyright in this document is exclusively owned by Tencent Cloud. You must not reproduce, modify, copy or distribute in any way, in whole or in part, the contents of this document without Tencent Cloud's the prior written consent.

Trademark Notice

STencent Cloud

All trademarks associated with Tencent Cloud and its services are owned by Tencent Cloud Computing (Beijing) Company Limited and its affiliated companies. Trademarks of third parties referred to in this document are owned by their respective proprietors.

Service Statement

This document is intended to provide users with general information about Tencent Cloud's products and services only and does not form part of Tencent Cloud's terms and conditions. Tencent Cloud's products or services are subject to change. Specific products and services and the standards applicable to them are exclusively provided for in Tencent Cloud's applicable terms and conditions.

Contents

Operation Guide

Operation Overview

Usage 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

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

Deleting Data Backup Point

Managing Snapshots

Creating Snapshot

Rolling back Data from Snapshot

Creating Cloud Disks Using Snapshots

Cross-region snapshot replication Deleting Snapshots Monitoring and Alarms

Operation Guide Operation Overview

Last updated : 2023-12-21 21:01:26

This document describes common cloud disk operations, including monitoring, alarms, managing cloud disks and snapshots.

Cloud Disk

Creating cloud disks Mounting Cloud Disks Initializing cloud disks Initialization Scenarios Initializing Cloud Disks (<2 TB) Initializing Cloud Disks (≥2 TB) Expanding cloud disks Cloud Disk Expansion Scenarios Expanding Cloud Disk Capacity Extending Partitions and File Systems (Windows) Extending Partitions and File Systems (Linux) Unmounting cloud disks Terminating cloud disks

Snapshots

Creating Snapshots Rolling Back Snapshots Creating cloud disks from snapshots Scheduled Snapshots Cross-region snapshot replication Deleting Snapshots

Monitoring and Alarms



Monitoring and alarms

Usage Limits

Last updated : 2024-04-30 11:02:54

Use Limits on Cloud Disk

Limit Type	Description	
Enhanced SSD usage	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. The performance of Enhanced SSD is only guaranteed when it's attached to S5, M5, and SA2 models created after August 1, 2020, and all later generation models, for example, Standard S5, S6, and later. Enhanced SSD cannot be used as the system disk. Enhanced SSD cloud disks cannot be encrypted. Enhanced SSD cannot be upgraded from other disk types.	
ulTra SSD usage	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 SSD can only be purchased together with Standard Storage Optimized S5se instances. 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	Starting from May 2018, all data disks purchased upon creation of CVMs are elastic cloud disks. You can detach them from and re- attach them to CVMs. This feature is supported in all availability zones.	
Cloud disk performances	I/O specification applies to both input and output performance at the same time.For example, if a 1-TB SSD has a maximum random IOPS of 26,000, it means that both its read and write performance can reach this value. Due to performance limits, if the block size in this example is 4 KB or 8 KB, the maximum IOPS can be reached. If the block size is 16 KB, the maximum IOPS cannot be reached (throughput has already reached the limit of 260 MB/s).	
Max attachable elastic cloud disks per	Up to 20.	

CVM	
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	Cloud disks and instances (CVM or Lighthouse instances) to which they are attached must be in the same AZ. Cloud disks of CVM instances are independent of cloud disks of Lighthouse instances. Cloud disks of CVM instances cannot be attached to Lighthouse instances, and vice versa.
Repossession of cloud disks in arrears	 When a monthly-subscribed cloud disk expires and you do not renew it within 7 days after the expiry, it will be forcibly unmounted from the CVM (if any), and moved to the recycle bin. For details about the repossession mechanism, see Arrears. Currently, when you mount a monthly-subscribed cloud disk to monthly-subscribed CVM, the following renewal method are available: Renew the cloud disk when the associated CVM expires Renew the cloud disk automatically on a monthly basis after it expires. Mount directly without a renewal policy.

Use Limits on Snapshot

Limit Type	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	Snapshot data can only be rolled back to the source cloud disk from which the snapshot was created. If you want to create a cloud disk with data in an existing snapshot, see Creating Cloud Disks Using Snapshots.
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



Limit Type	Description
Scheduled snapshot policy quota in one region	A single Tencent Cloud account can set a maximum of 30 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 : 2023-12-21 21:12:14

Overview

CBS allows you to create a cloud disk and attach it to any CVM in the same availability zone. The cloud disk is identified and used by the CVM through block storage device mapping. Once created, the cloud disk can reach its maximum performance without prefetch.

You can create different types of cloud disks as needed as detailed in Cloud Disk Types.

Prerequisite

Before creating a cloud disk, you need to sign up for Tencent Cloud and complete identity verification.

Directions

Creating a cloud disk through the console

Creating a cloud disk using a snapshot

Creating a cloud disk using the API

1. Log in to the CBS console.

2. Select a region and click **Create**.

3. In the **Purchase Data Disk** dialog box, configure the following parameters:

Parameter	Configuration
AZ	Required. The availability zone where the created cloud disk resides. It cannot be modified after the cloud disk is created.
Cloud Disk Type	Required. Cloud disk types include: Premium Cloud Storage Balanced SSD SSD Enhanced SSD ulTra SSD. This type can only be purchased together with the Storage Optimized CVM instance.

Quick Disk Creation	Optional. To create a cloud disk using a snapshot, you need to tick Create a cloud disk with a snapshot and select the snapshot you want to use. The capacity of a cloud disk created using a snapshot is equal to that of the snapshot by default. You can adjust the disk capacity. 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 cloud disk capacity and specifications are as follows: Premium Cloud Storage: 10 to 32,000 GB Balanced SSD: 20 to 32,000 GB SSD: 20 to 32,000 GB Enhanced SSD: 20 to 32,000 GB When you create a cloud disk with a snapshot, the disk capacity cannot be smaller than that of the snapshot. If you don't specify the capacity of the cloud disk, the capacity is equal to that of the snapshot by default.
Scheduled Snapshot	Optional. You can associate scheduled snapshot policies when creating a cloud disk to regularly manage your cloud disk snapshots. Currently, Tencent Cloud provides a 50 GB free tier for each region in the Chinese mainland. For more information, see Billing Overview.
Disk Name	 Optional. A maximum of 128 characters are supported. It must start with a letter and can be a combination of letters, digits, and special symbols (. , _ , : , -). This parameter can be modified after the cloud disk is created. Creating a single cloud disk: The disk name entered is the name of the cloud disk you create. Batch creating cloud disks: The disk name entered is the prefix of your cloud disk names, in the format of disk name_number.
Projects	Required. When creating a cloud disk, you can configure the project to which the cloud disk belongs. The default value is DEFAULT PROJECT .
Label	Optional. When creating a cloud disk, you can bind a tag to it. A tag is used for identification, helping you categorize and search cloud resources. For more information on tags, see Tag.
Billing mode	Required. CBS is pay-as-you-go.
Quantity	Optional. The default value is 1, meaning only one cloud disk is created. Currently, up to 50 cloud disks can be created at one time.
Validity Period	The pay-as-you-go billing mode does not involve this parameter.

4. Click OK.

If Billing Mode is Pay-as-you-go, the creation is completed.

5. You can view the cloud disk(s) you created in the Cloud Block Storage list page. The new elastic cloud disk is in the

To be attached status. To attach it to a CVM instance in the same availability zone, see Attaching Cloud Disks.

If you want to create a cloud disk that contains all data upon creation, you can create cloud disks using snapshots.

You can use the CreateDisks API to create a cloud disk. For more information, see CreateDisks.

Attaching Cloud Disks

Last updated : 2023-12-21 21:15:02

Overview

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.

This document describes how to attach cloud disks to a CVM instance in the console and by using an API.

Note:

Some Linux CVMs may not recognize an elastic cloud disk. You must first enable the disk hot swapping feature in the CVM. For more information, see Enabling the disk hot swapping feature.

Directions

Attaching cloud disks in the CBS console

1. Log in to the CBS console.

2. On the cloud disk list page, select and attach the cloud disks. Note that you can only attach cloud disks in **To be attached** status.

Attach a single cloud disk: Find the target cloud disk, and click More > Attach under the Operation column.

Attach multiple cloud disks: Select the cloud disks, and click Attach above the cloud disk list to attach them in a batch.

3. In the Attach to instance pop-up window, attach the disks to the instance.

Select the target instance, and complete the parameters **Attaching options**.

Unified expiry time with the instance (XXX)

Enable monthly auto-renewal of the cloud disks (recommended)

Attach directly

4. Click Next.

After attaching the cloud disk, you need to log in to the instance and initialize the disk.

5. Click Attach now.

If the status of the cloud disk changes to Attached, the attachment is successful.

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

Creation mode	Cloud disk capacity	Subsequent operations
Manually created	Cloud disk capacity < 2 TB	Initializing cloud disks (<2 TB)



	Cloud disk capacity \geq 2 TB	Initializing cloud disks (≥2 TB)
Created from a snapshot	Cloud disk capacity = Snapshot capacity	Attach to Windows CVMs Log in to the instance, and go to Server ManagerStorageDisk Management to make the disk online. Attach to Linux CVMs: Log in to the instance, and run he command mount such as mount /dev/vdb /mnt , to make the disk usable.
	Snapshot capacity < Cloud disk capacity ≤ 2 TB or 2 TB < Snapshot capacity < Cloud disk capacity	Attach to Windows CVMs: Extending Partitions and File Systems (Windows) Attach to Linux CVMs: Extending Partitions and File Systems (Linux)
	Snapshot capacity ≤ 2 TB < Cloud disk capacity	If the snapshot uses a MBR partition, see Initializing cloud disks (≥2 TB) to repartition the disk with GPT. Note that this operation will delete the existing data. If the snapshot uses a GPT partition, Attach to Windows CVMs: [Extending Partitions and File Systems (Windows)] Attach to Linux CVMs: [Extending Partitions and File Systems (Linux)]

Using the API to attach cloud disks

You can use the AttachDisks API to attach cloud disks. For more information, see AttachDisks.

Related Operations

Enabling the disk hot swapping feature

All existing images already support attaching and detaching elastic cloud disks. **To detach a cloud disk, you must** first run the umount command in Linux CVM or perform offline operations in Windows CVM. Otherwise, the reattached elastic cloud disk may not be recognized.

Note:

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

Operating system	Version
CentOS	5.11 64-bit
	5.11 32-bit
	5.8 64-bit



	5.8 32-bit
Debian	6.0.3 32-bit
Libuatu	10.04 64-bit
Obuntu	10.04 32-bit
openSUSE	12.3 64-bit
opensose	12.3 32-bit

1. Log in to the Linux instance using Web Shell as the root user.

2. Run the following command to add the driver.





modprobe acpiphp

Note:

If you still need to load the acpiphp module after shutting down or re-starting the CVM, we recommend you run Step 3 to set the acpiphp module to autoload.

3. (Optional) If you need to load the acpiphp module automatically after shutting down or re-starting the CVM, run the following command according to the operating system.

CentOS 5 Series

Debian 6 Series, Ubuntu 10.04 Series



openSUSE 12.3 Series

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



vi /etc/sysconfig/modules/acpiphp.modules

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





#!/bin/bash
modprobe acpiphp >& /dev/null

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





chmod a+x /etc/sysconfig/modules/acpiphp.modules

1. Run the following command to modify the file.





vi /etc/modules

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





acpiphp

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"

Attaching data disks automatically upon instance creation

If you specify a custom image and data disk snapshot while creating a CVM instance, the cloud disk created at the same time is automatically attached to the CVM and can be read and written without being initialized by partitioning, formatting and other operations. However, the specified custom images and data disk snapshots must meet the following requirements.

Windows instances



Linux instances

If you use a custom image to create a Windows CVM instance, the cloud disk created from the corresponding data disk snapshot will be automatically attached. The custom image and the data disk snapshot must meet the following requirements:

The data disk **must** be formatted to ntfs or fat32 before you create a snapshot.

The SAN policy in the custom image is onlineAll .

Checking configurations

Tencent Cloud provides pre-configured public images for Windows by default, but we still recommend that you perform the following steps to check configurations before creating a custom image. Run the following commands in order and check the returned results.





diskpart





san

See below:

PS C:\Users\Administrator>			
PS C:\Users\Administrator> diskpart			
Microsoft DiskPart Version 6.1.7601			
Copyright (C) 1999-2008 Microsoft Corporation.			
On computer : 10_135_6_116			
DISKPART> san			
SAN Policy: OfflineShared			
DISKPART> san policy=onlineall			
DiskPart You have modified the SAN policy for the operating system.			
DISKPART> san			
SAN Policy: OnlineAll			

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

The data disk **must** be formatted and attached to the source CVM.

Add the following commands to the /etc/rc.local file to configure the mount point of the data disk before creating a system disk on the system disk.





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

Note:

Replace <mount-point> with the mount point of the file system, such as /mydata .

Replace<device-id>with the partition path of the file system. For example, enter/dev/vdbif the file systemhas no partition, and/dev/vdb1if the file system has partitions.

Initializing Cloud Disks Initialization Scenarios

Last updated : 2023-12-21 21:16:06

The cloud disks created through the console are manually mounted to your CVM and used as data disks in Online status by default. To use the disks, you need to initialize them first, including formatting, partitioning, and creating file systems. The initialization method varies by actual use scenario as shown below:

If the entire disk is presented as one independent partition (that is, there are not multiple logical disks such as D disk/vdb1 and E disk/vdb2), we strongly recommend you not use partitions, and directly create the file system on bare devices.

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.

Common disk partition styles are Main Boot Record (MBR) and Guid Partition Table (GPT). If the disk partition format is changed after the disk is put into use, original data on the disk will be erased. Therefore, select an appropriate partition style according to actual needs.

Partition style	Maximum supported disk capacity	Number of partitions supported	Partition tool
MBR	2TB	4 primary partitions 3 primary partitions and 1 extended partition	Windows operating system: Disk management Linux operating system: fdisk tool parted tool
GPT	18EB Currently, cloud disk supports a maximum capacity of 32TB	No limit on number of partitions	Windows operating system: Disk managementLinux operating system: parted tool

Basics of the two partition styles are shown in the following table :

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

For disk capacity less than 2TB:

Initializing cloud disks (Windows)

Initializing cloud disks (Linux)

For disk capacity larger than or equal to 2TB:

Initializing cloud disks (Windows)

Initializing cloud disks (Linux)

Initializing Cloud Disks (<2 TB)

Last updated : 2023-12-21 21:18:42

Overview

This document describes how to initialize a cloud disk with a capacity less than 2 TB. For more cases, see Initialization Scenarios.

Prerequisite

You have attached the cloud disk to your CVM.

Notes

Read FAQs about cloud disk usage before working with cloud disks.

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 prevent service exceptions, ensure that the CVM has stopped providing external services before formatting.

Directions

Initializing cloud disks (Windows)

Initializing cloud disks (Linux)

Note:

This example uses a Lighthouse instance using Windows Server 2012 R2 operating system. Note that the steps may vary according to the operating system version.

1. Log in to the Windows CVM instance.

2. On the desktop, right-click

in the lower-left corner.
3. Select **Disk Management** in the pop-up menu.
Note:



If the newly added disk is in offline status, execute Step 4 before Step 5 to perform initialization. Otherwise, you can directly execute Step 5.

4.

Disks are listed on the right-side pane. Right-click disk 1 area, and select **Online** to bring it online. The status of disk 1 changes from **Offline** to **Not Initialized**.

5.

Right-click disk 1 area and select Initialize Disk in the menu.

6. In the Initialize Disk pop-up window, the disk to be initialized is displayed. Select MBR (Master Boot Record) or GPT (GUID Partition Table) and click OK.

Note:

If the disk partition format is changed after the disk is put into use, the original data on the disk will be erased. Select an appropriate partition format based on actual needs.

7. Right-click the free space of the disk, and select **New Simple Volume**.

8. In the **New Simple Volume Wizard** pop-up window, click **Next**.

- 9. Specify the volume size as needed, which is the maximum value by default. Click Next.
- 10. Assign a drive letter and click **Next**.

11. Select **Format this volume with the following settings**, configure parameters as needed, format the partition, and click **Next** to complete the partition creation.

12. Click **Complete** to complete the wizard. Wait for the completion of initialization. When the volume status becomes **Healthy**, the disk initialization is successful.

After the initialization is complete, enter the **PC** interface to view the new disk.

Select the initialization method according to your actual use cases:

If the entire disk is presented as one independent partition (there is no logical disks such as vdb1 and vdb2), we

strongly recommend that you not use partition, and directly create the file system on bare devices.

If the entire disk needs to be presented as multiple logical partitions (there are multiple logical disks), you need to first partition the disk, and then create the file system on a partition.

Creating file systems on bare devices

1. Log in to the Linux CVM instance.

2. Run the following command as the root user to view the disk name.





fdisk -l

If the returned result is as shown below, the current CVM has two disks, where /dev/vda is the system disk and /dev/vdb is the newly added data disk.

[root0VM_16_14_centos ~]# fdisk -1 Disk /dev/vda: 53.7 GB, 53687091200 bytes, 104857600 sectors Units = sectors of $1 \times 512 = 512$ 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 Start Device Boot End Blocks Id System 2048 104857599 83 Linux /dev/vda1 × 52427776 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

The partition size supported by different file systems varies. Select an appropriate file system as needed. The following example takes EXT4 as the file system:





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>

Take creating a new 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>

Take creating a new 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 don't need to configure disk automount at startup, skip the following steps.

7. Confirm the mount 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:

🔗 Tencent Cloud

method		
Use the soft link of the elastic cloud disk (recommended)	Pros: The soft link of an elastic cloud disk is fixed and unique. It does not change with operations such as mounting, unmounting, and formatting partitions.Cons: Only an elastic cloud disk can use the soft link, which operates imperceptibly for the partition formatting operation.	=Run the following command to obtain the soft link of the elastic cloud disk.
Use the UUID of the file system	Auto-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 obtain the UUID of the file system.
Use device name	Automatic mounting 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 obtain the device name.





8. Run the following command to back up the /etc/fstab file to the /home directory, for example:





cp -r /etc/fstab /home

9. Execute the following command to use VI editor to open the /etc/fstab file.





vi /etc/fstab

10. Press i to enter edit mode.

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





<Device information> <Mount point> <File system format> <File system installation o

(**Recommended**) Take automatic mounting using the soft link of an elastic cloud disk as an example. Add the following content:





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

Note:

If you have multiple elastic cloud disks, you can compare disk-xxxx 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 0 0

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





/dev/vdb /data ext4 defaults 0 0

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

Save the configuration and exit the editor.

13. Run the following command to check whether 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 example uses the parted partition tool in the CentOS 7.5 operating system to configure data disk /dev/vdc as the primary partition. MBR is used as the default partition format, EXT4 format as the file system, and

/data/newpart as the mount point. Disk automount at startup is configured. Note that the formatting operation may vary according to the operating system.

- 1. Log in to the Linux CVM instance.
- 2. Run the following command as the root user to view the disk name.



fdisk -l

If the returned result is as shown in the following figure, the current CVM has two disks, where /dev/vda is the system disk and /dev/vdb is the newly added data disk.

[root@VM_16_14_0	entos ~]# f	°disk −1			
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 /dev/vda1 *	Start 2048	End 104857599	Blocks 52427776	Id 83	System 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 use the fdisk tool to partition the newly added data disk.





fdisk <Newly added data disk>

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





fdisk /dev/vdb

The returned information is similar to what is shown below:

[root@VM_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 Øxcb7892a8.
Command (m for help):

4. Enter n and press **Enter** to start creating the new partition.

The returned information is similar to what is shown below:



This indicates that the disk has two partition types:

p: Primary partition.

e: Extended partition.

5. Take creating a primary partition as an example. Enter p and press **Enter** to start creating a primary partition.

The returned information is similar to what is shown below:



Partition number indicates the number of the primary partition. Valid range: 1-4.

6. Take selecting partition number 1 as an example. Enter the primary partition number 1 and press Enter.

The returned information is similar to what is shown below:

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

First sector indicates the start sector. Valid range: 2048 (default value) - 125829119.

7. Take selecting the default start sector number 2048 as an example. Press Enter.

The returned information is similar to what is shown below:

First sector (2048-125829119, default 2048): Last sector, +sectors or +size{K,M,G} (2048-125829119, default 125829119): Last sector indicates the end sector. Valid range: 2048 - 125829119 (default value).

8. Take selecting the default end sector number 125829119 as an example, press Enter.

The returned information is similar to what is shown below:



This indicates that partitioning is completed. A new partition has been created on the 60 GB data disk.

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

The returned information is similar to what is shown below:

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			

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

Note:

If an error occurs during the partitioning operation, enter q to exit the fdisk partition tool, and the previous partition result will not be retained.

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

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



11. Run the following command to sync the partition table to the operating system.





partprobe

12. Run the following command to configure the file system of the newly created partition to that required by the system.





mkfs -t <File system format> /dev/vdb1

The partition size supported by different file systems varies. Select an appropriate file system as needed. The following example takes EXT4 as the file system:





mkfs -t ext4 /dev/vdb1

The returned information is similar to what is shown below:

[rootQVM_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

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>





mkdir /data/newpart

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





mount /dev/vdb1 <Mount point>

Take creating a new mount point /data/newpart as an example:





mount /dev/vdb1 /data/newpart

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





df -TH

The returned information is similar to what is shown below:

[root@VM_16_	14_centos ~]	# mkdin	r ∕dat	ta∕n <mark>ew</mark> p	part
[root@VM_16_	14_centos ~]	# mount	t ∕dev	∥vdb1	/data/newpart
[root@VM_16_	14_centos ~]	# df -'	ΓH		
Filesystem	Туре	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
[root@VM_16_	14_centos ~]	#			

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

Note:

If you don't need to configure disk automount at startup, skip the following steps.

16. Confirm the mount 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:

Mounting method	Pros and cons	Information acquisition method
Use the soft link of the elastic cloud disk (recommended)	Pros: The soft link of an elastic cloud disk is fixed and unique. It does not change with operations such as mounting, unmounting, and formatting partitions.Cons: Only an elastic cloud disk can use the soft link, which operates imperceptibly for the partition formatting operation.	Run the following command to obtain the soft link of the elastic cloud disk.
Use the UUID of the file system	Auto-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 obtain the UUID of the file system.





17. Run the following command to back up the /etc/fstab file to the /home directory:





cp -r /etc/fstab /home

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





vi /etc/fstab

19. Press i to enter edit mode.

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





<Device information> <Mount point> <File system format> <File system installation o

(**Recommended**) Take automatic mounting using the soft link of an elastic cloud disk as an example. Add the following content:





/dev/disk/by-id/virtio-disk-drkhklpe-part1 /data/newpart ext4 defaults 0 2

Note:

If you have multiple elastic cloud disks, you can compare disk-xxxx 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 0 2

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





/dev/vdb1 /data/newpart ext4 defaults 0 2

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

Save the configuration and exit the editor.

22. Run the following command to check whether 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 Operations

Initializing Cloud Disks (\geq 2 TB).

Initializing Cloud Disks (≥2 TB)

Last updated : 2023-12-21 21:22:45

Overview

This document describes how to initialize a cloud disk with a capacity greater than or equal to 2 TB. For more cases, see Initialization Scenarios.

MBR supports disk with a maximum capacity of 2 TB. When you partition disk with a capacity greater than 2 TB, we recommend that you use the GPT partition format. When GPT is used on Linux operating system, fdisk can no longer be used and the parted tool must be used.

Prerequisite

You have attached the cloud disk to your CVM.

Notes

Read FAQs about cloud disk usage before working with cloud disks.

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 prevent service exceptions, ensure that the CVM has stopped providing external services before formatting.

Directions

Initializing cloud disks (Windows)

Initializing cloud disks (Linux)

Note:

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

1. Log in to the Windows CVM instance.

2. On the desktop, click

÷

3. Go to the Server Manager page and click File and Storage Services on the left sidebar.

4. On the left sidebar, select **Volume** > **Disk**.

Note:

If the newly added disk is in offline status, execute Step 5 before Step 6 to perform initialization. Otherwise, you can directly execute Step 6.

5.

Disks are listed on the right pane. Right-click the row where 1 is located, and select **Online** to bring it online. Then its status becomes **Online**.

6.

Right-click the row where 1 is located, and select **Initialize** in the menu.

- 7. Click **Yes** as prompted.
- 8. After the initialization, 1 changes from the **Unknown** to **GPT** partition. Click the row where 1 is located and select **New Simple volume** in the menu.
- 9. In the welcome page of New Volume Wizard pop-up window, click Next.
- 10. Select the server and disk and click Next.
- 11. Specify the volume size as needed, which is the maximum value by default. Click **Next**.
- 12. Assign a drive letter and click **Next**.
- 13. Configure parameters as needed, format the partition, and click **Next** to complete the partition creation.
- 14. After confirming that everything is correct, click Create.
- 15. Wait for the completion of the volume creation and click **Close**.
- After the initialization is completed, go to My Computer to view the new disk.

Select the initialization method according to your actual use cases:

If the entire disk is presented as one independent partition (there is no logical disks such as vdb1 and vdb2), we

strongly recommend that you not use partition, and directly create the file system on bare devices.

If the entire disk needs to be presented as multiple logical partitions (there are multiple logical disks), you need to first partition the disk, and then create the file system on a partition.

Creating file systems on bare devices

- 1. Log in to the Linux CVM instance.
- 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.

[root0VM_16_14_centos ~]# fdisk -1 Disk /dev/vda: 53.7 GB, 53687091200 bytes, 104857600 sectors Units = sectors of $1 \times 512 = 512$ 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 Start Device Boot End Blocks Id System 2048 104857599 83 Linux /dev/vda1 × 52427776 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

The partition size supported by different file systems varies. Select an appropriate file system as needed. The following example takes EXT4 as the file system:





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>

Take creating a new 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>

Take creating a new 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 don't need to configure disk automount at startup, skip the following steps.

7. Confirm the mount 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:

Mounting	Pros and cons	Information acquisition
----------	---------------	-------------------------

🔗 Tencent Cloud

method		method
Use the soft link of the elastic cloud disk (recommended)	Pros: The soft link of an elastic cloud disk is fixed and unique. It does not change with operations such as mounting, unmounting, and formatting partitions.Cons: Only an elastic cloud disk can use the soft link, which operates imperceptibly for the partition formatting operation.	Run the following command to obtain the soft link of the elastic cloud disk.
Use the UUID of the file system	Auto-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 obtain the UUID of the file system.
Use device name	Automatic mounting 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 obtain the device name.





8. Run the following command to back up the /etc/fstab file to the /home directory, for example:





cp -r /etc/fstab /home

9. Execute the following command to use VI editor to open the /etc/fstab file.





vi /etc/fstab

10. Press i to enter edit mode.

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





<Device information> <Mount point> <File system format> <File system installation o

(**Recommended**) Take automatic mounting using the soft link of an elastic cloud disk as an example. Add the following content:





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

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





/dev/vdb /data ext4 defaults 0 0

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

Save the configuration and exit the editor.

13. Execute the following command to check whether 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 document uses the parted partition tool in the CentOS 7.5 operating system to configure the /dev/vdc data disk as the primary partition. GPT is used as the default partition format, EXT4 format as the file system, and

/data/newpart2 as the mount point. Disk automount at startup is configured. Note that the formatting operation may vary according to the operating system.

- 1. Log in to the Linux CVM instance.
- 2. Run the following command as the root user to view the disk name.



lsblk

If the returned result is as shown below, the current CVM instance has two disks, where /dev/vda is the system disk and /dev/vdc is the newly added data disk.

[root@VM_16_14_centos ~]#					lsbll	ĸ
Name	MAJ:MIN	RM	SIZE	RO	TYPE	MOUNTPO INT
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	ЗT	0	disk	

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



parted <Newly added data disk>

Take the newly mounted data disk /dev/vdc as an example:





parted /dev/vdc

The returned information is similar to what is shown below:





4. Enter p and press Enter to view the current disk partition format.

The returned information is similar to what is shown below:

```
(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) _
```

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 whether the disk partition format has been configured successfully. The returned information is similar to what is shown below:

(parted) mklabel gpt					
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 unit of the disk to sector.

8. Take creating one partition for the entire disk as an example, enter mkpart opt 2048s 100% and press

Enter.

Where, 2048s indicates the initial disk capacity and 100% indicates the final disk capacity. This is for reference only. You can choose the number of disk partitions and their capacities based on business needs.

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

The returned information is similar to what is shown below:



This indicates the detailed information 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

The returned information is similar to what is shown below. You can now see the new partition /dev/vdc1 .

[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	ЗT	0	disk
└-vdc1	253:33	0	ЗT	0	part

12. Run the following command to configure the file system of the newly created partition to that required by the system.



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

The partition size supported by different file systems varies. Select an appropriate file system as needed. The following example takes EXT4 as the 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>

Take creating a new 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>

Take creating the /data/newpart2 mount point as an example:





mount /dev/vdc1 /data/newpart2

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




df -TH

The returned information is similar to what is shown below:

[root@VM_16_14	_centos ~]	#mkdi	ir /dat	ta∕newp	part2	
[root@VM_16_14	_centos ~]	# moun	nt Zdev	∥vdc1	/data	a/newpart2
[root@VM_16_14	_centos ~]	# df -	-TH			
Filesystem	Туре	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

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

Note:

If you don't need to configure disk automount at startup, skip the following steps.

16. Confirm the mounting method and get 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:

Mounting method	Pros and cons	Information acquisition method
Use the soft link of the elastic cloud disk (recommended)	Pros: The soft link of an elastic cloud disk is fixed and unique. It does not change with operations such as mounting, unmounting, and formatting partitions.Cons: Only an elastic cloud disk can use the soft link, which operates imperceptibly for the partition formatting operation.	Run the following command to obtain the soft link of the elastic cloud disk.
Use the UUID of the file system	Auto-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 obtain the UUID of the file system.





17. Run the following command to back up the /etc/fstab file to the /home directory, for example:





cp -r /etc/fstab /home

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





vi /etc/fstab

19. Press i to enter edit mode.

Move the cursor to the end of the file, press Enter, and add the following content.





<Device information> <Mount point> <File system format> <File system installation o

(**Recommended**) Take automatic mounting using the soft link of an elastic cloud disk as an example. Add the following content:





/dev/disk/by-id/virtio-disk-bm42ztpm-part1 /data/newpart2 ext4 defaults 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 0 2

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





/dev/vdc1 /data/newpart2 ext4 defaults 0 2

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

Save the configuration and exit the editor.

21. Run the following command to check whether 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 Operations

Initializing Cloud Disks (<2 TB).

Expanding Cloud Disks Overview

Last updated : 2023-11-20 15:04:39

Expanding Cloud System Disks

To expand a system disk, the following three methods are provided. Expanding via the CVM console Expanding via the CBS console Expanding via API

Expanding Cloud Data Disks

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

Expanding via the CVM console

Expanding via the CBS console

Expanding via API

Based on the **attaching** status of CBS data disks, you can expand their capacity by different methods.

If the current CBS data disk can be detached, you can expand its capacity in the CBS console or via the

ResizeDisk API.

If the current CBS data disk **cannot be detached**, you can expand its capacity in the CVM console or via the ResizeDisk API.

Note

If the maximum capacity of the cloud disk cannot meet your business needs, try building up RAID groups or building 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	Subsequent operations
File system not	Disk capacity < 2 TB	Initializing cloud disks (< 2 TB)
created	Disk capacity ≥	Initializing cloud disks (≥ 2 TB)



	2 TB	
	Disk capacity < 2 TB	Cloud disks attached to a Windows CVM: extending partitions and file systems (Windows) Expanded cloud disk attached to a Linux CVM: extending partitions and file systems (Linux)
File system created	Disk capacity ≥ 2 TB	GPT partition format: extending partitions and file systems (Windows) or extending partitions and file systems (Linux) MBR partition format: not supported. MBR partition supports disk with a maximum capacity of 2 TB. When you partition disk with a capacity greater than 2 TB, we recommend that you create and attach a new data disk and use the GPT partition format to copy data.

Billing Description

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

Monthly subscribed cloud disks: To expand the capacity, you need to make up the price difference between new and old configurations for the rest of the lifecycle, subject to the details as displayed on the payment page.

Pay-as-you-go cloud disks: The adjustment takes effect immediately, and the cloud disks will be billed by the new configuration.

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 Examples



Note

The following prices are only used as examples and are not the actual prices listed on the official website. For accurate CBS pricing, refer to Price Overview.

A user in the Guangzhou region, purchased a 200 GB general-purpose yearly/monthly subscription SSD cloud disk

with a duration of 1 month as the data disk on May 1, 2022. The user expanded the disk capacity to 300 GB on May 5, 2022. The billing calculation is as follows:

Monthly upgrade price difference = 150 - 100 = 50 USD/month

Number of upgrade days = 27 days

Applicable discount = No discount

Total upgrade fee = 50 * (27/(365/12)) * 1 = 44.38 USD

Expanding Cloud Disks

Last updated : 2023-12-21 21:30:23

Overview

A cloud disk is an expandable storage device on cloud. You can expand its capacity at any time without losing any data in it.

After expanding the cloud disk, you need to expand the partitions and file systems. You can allocate the capacity of the expanded part to an existing partition or format it into a new partition.

Note:

MBR partition supports a disk with a maximum capacity of 2 TB. If you need to extend for over 2 TB, we recommend you create and attach a new data disk and use the GPT partition format.

Expanding Data Disks

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

Note:

If multiple cloud disks of the same capacity and type are attached to the CVM, you can identify them using the method shown in Distinguishing data disks. Select a data disk and expand its capacity as instructed below.

Expanding via the CVM console (recommended)

Expanding via the CBS console

Expanding via API

1. Log in to the CVM console.

- 2. Choose More > Resource adjustment > Expand cloud disks in the Operation column.
- 3. Select the data disk to be expanded in the pop-up window, and click Next.
- 4. Select a new capacity, which must be greater than or equal to the current capacity, and click Next.
- 5. Read the notes and click **Adjust Now**.

6. Assign its expanded capacity to an existing partition, or format it into an independent new partition. Depending on the operating system of the CVM, see Extending Partitions and File Systems (Windows) or Determining the Expansion Method.

- 1. Log in to the CBS console.
- 2. Choose **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. Assign its expanded capacity to an existing partition, or format it into an independent new partition. Depending on the operating system of the CVM, see Extending Partitions and File Systems (Windows) or Determining the Expansion Method.

You can use the ResizeDisk API to expand the specified cloud disks. For more information, see ResizeDisk.

Expanding System Disks

Expanding via the CVM console (recommended) Expanding via the CBS console Expanding via API

1.

Log in to the CVM console, and choose More > Resource adjustment > Expand cloud disks in the row of the target CVM.

2. In the pop-up Expand Cloud Disks window, select the system disk that needs expansion and click Next.

3. In the **Adjust Capacity** step, set the target capacity (which must be greater than the current capacity) and click **Next**.

4. In the Extend Partition and File System step, view the precautions and click OK.

5. After completing the expansion operation on the console, log in to the instance to confirm whether the file system is automatically expanded. If not, refer to Extending System Disk Partitions and File Systems Online for steps on expanding partitions and file systems.

1. Log in to the CBS console.

- 2. Choose More > Expand for the target cloud disk.
- 3. Select the desired new capacity size (which must be greater than the current capacity).
- 4. Complete the payment.

5. After completing the expansion operation on the console, log in to the instance to confirm whether the file system is automatically expanded. If not, refer to Extending System Disk Partitions and File Systems Online for steps on expanding partitions and file systems.

You can use the ResizeInstanceDisks API to expand the specified non-elastic cloud disk. For specific operations, see ResizeInstanceDisks.

Related Operations

Distinguishing data disks

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

Linux



Windows

- 1. Log in to a 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 following information will appear:

[root@VM_63_126_centos	~]# ls -1	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-jwz431pg ->//vde
lrwxrwxrwx 1 root root	9 Mar 1	17:31 virtio-disk-punhzcju ->//vdd

Note that disk-xxxx is the ID of a cloud disk. You can use it to view cloud disk details on the CBS console.

1. Log in to a Windows instance.

2. Right-click



3. Enter cmd in the popup window and press Enter.

4. Run the following command to view the relationship between the elastic cloud disks and the device name.





wmic diskdrive get caption, deviceid, serialnumber

You can also run the following command.





wmic path win32_physicalmedia get SerialNumber,Tag

The following information will appear:

C:\Users\Administrator>wmic diskdrive get caption,deviceid,serialnumber Caption DeviceID SerialNumber Red Hat VirtIO SCSI Disk Device \\.\PHYSICALDRIVEØ Red Hat VirtIO SCSI Disk Device \\.\PHYSICALDRIVE1 disk-hmvcmqrm Note that disk-xxxx is the ID of a cloud disk. You can use it to view cloud disk details on the CBS console.

Checking the cloudinit configuration

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

Checking the cloudinit configuration for Linux instances

Checking the cloudinit configuration for Windows instances

After the system disk is expanded, log in to the Linux instance and check whether the/etc/cloud/cloud.cfgfile contains thegrowpartandresizefsconfiguration items.

If yes, ignore other operations.



growpart: Expands the partition to the disk size.

resizefs: Expands or adjusts the file system in the / partition to the partition size.

If no, manually extend partitions and file systems of the Linux instance, and assign its extended capacity to an existing partition, or format it into an independent new partition.

After the system disk is expanded, log in to the Windows instance and check whether the

ExtendVolumesPlugin configuration item exists under plugin in C:\\Program Files\\Cloudbase Solutions\\Cloudbase-Init\\conf\\cloudbase-init.conf .

If yes, ignore other operations.

If no, manually extend partitions and file systems (Windows) according to the operating system, and assign its extended capacity to an existing partition, or format it into an independent new partition.

Extending System Disk Partitions and File Systems Online

Last updated : 2023-12-21 21:31:55

Overview

Expanding a data disk in the console only increases its storage capacity. You need to expand the partition or file system of the cloud disk. This document describes how to expand partitions and file systems online.

Prerequisite

Before expanding a partition or file system, create a snapshot of the system disk to back up data. For detailed directions, see Creating Snapshots.

This practice helps you roll back snapshot to recover data in case of data loss due to maloperation.

You have expanded the cloud disk in the console. For detailed directions, see Expanding Cloud Disks.

The Linux CVM kernel should be version 3.6.0 or later. You can use the uname -a command to check the kernel version.

If the kernel version is earlier than 3.6.0, see Extending MBR Partitions and File Systems (Smaller than 2 TB).

Operating Environment

Linux instances

Windows instances

Resource	Description
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.

Resource	Description
Operating system	Windows Server 2012 R2 IDC 64-bit Chinese



Cloud disk (system disk) C drive: Uses the MBR partition and the NTFS file system, and expands from 50 GB to 100 GB online in the console.

Directions

Perform the following steps based on your operating system:

Linux instances

Windows instances

1. Log in to a Linux instance using standard login method.

2. Run the following command to query partitions of the cloud disk.





fdisk -l

The 60 GB data disk dev/vda contains a 50 GB MBR partition /dev/vda1.

[root@VM-5-1	17-ce	ntos ~]# fdisk -	-1			
Disk /dev/va	da: 6	0 GiB,	644245094	40 bytes,	12582	9120	sectors
Units: secto	ors o	f 1 *	512 = 512	bytes			
Sector size	(log	ical/p	hysical):	512 bytes	/ 512	byte	3
I/O size (mi	inimu	m/opti	.mal): 512	bytes / 51	2 byt	es	
Disklabel ty	ype:	dos					
Disk identif	fier:	0x89e	e0607				
Device H	Boot	Start	End	Sectors	Size	Id Ty	pe
/dev/vda1 *	*	2048	104857566	104855519	50G	83 Li	nux

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





df -TH

The /dev/vda1 file system type is EXT4.

[root@VM-5-17-0	centos ~] #	df -TH	I			
Filesystem	Туре	Size	Used	Avail	Use %	Mounted on
devtmpfs	devtmpfs	937м	0	937M	08	/dev
tmpfs	tmpfs	953M	25k	953M	18	/dev/shm
tmpfs	tmpfs	953M	447k	953M	18	/run
tmpfs	tmpfs	953M	0	953M	08	/sys/fs/cgroup
/dev/vda1	ext4	53G	3.2G	48G	78	/
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 instance. 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 /dev/vda1 partition. In the command, you need to separate /dev/vda and 1 by space.





growpart /dev/vda 1

The following information appears:

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

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





resize2fs /dev/vda1

The following information appears:

```
[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-5-17-0	centos ~] ‡	df -TH	I			
Filesystem	Туре	Size	Used	Avail	Use%	Mounted on
devtmpfs	devtmpfs	937M	0	937M	08	/dev
tmpfs	tmpfs	953M	25k	953M	18	/dev/shm
tmpfs	tmpfs	953M	447k	953M	18	/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

After the successful expansion, check the data integrity and whether the business in the CVM instance runs normally.

1. Log in to the CVM instance as instructed in Logging in Using Standard Method (Recommended).

2. On the desktop, right-click



in the lower-left corner and click **Disk management** in the pop-up menu.

3. In the **Disk Management** pop-up window, select **Operation** > **Rescan the disk** at the top of the page.

After the scan, you can view the newly added space.

4. Right-click the area of the C drive and select **Expand Volume** in the pop-up menu.

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

After the successful expansion, check the data integrity and whether the business in the CVM instance runs normally.

Extending Data Disk Partitions and File Systems (Windows)

Last updated : 2023-12-21 21:32:23

Overview

After expanding a cloud disk in the console, you need to either assign its expanded capacity to an existing partition, or format it into an independent new partition.

If you expand a cloud disk that is attached to a running CVM, you need to rescan the disk to recognize the disk capacity 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:

Extending the file system may affect existing data. We strongly recommend that you manually create a snapshot to back up your data before performing the operation.

To extend the file system, you need to restart the instance or rescan the disk, which will lead to business interruption for a certain period. We recommend you choose an appropriate time for this operation.

After extension of the file system, we strongly recommend you rescan the disk to recognize the capacity. If you **refresh** the system or perform other operations, the expanded capacity may not be recognized.

Prerequisite

You have expanded the cloud disk via the console.

You have attached the cloud disk to a Windows CVM via the console and created a file system. You have logged in to a Windows CVM instance on which you want to extend partitions and the file system. See Logging in to Windows Instance Using RDP (Recommended).

Note:

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

Directions

Note:

If you expand a cloud disk that is attached to a running CVM instance in the console, you must rescan the disk to recognize the capacity before extending the file system of an existing partition or creating a partition. If the cloud disk to be expanded is not attached to a CVM instance or the attached CVM instance is shut down, you can directly proceed with extending the file system of an existing partition or creating a partition. If the Virtio driver of the CVM is earlier than version 58003, restart the instance before performing the operation. You can first check the Virtio driver version.

Rescanning the disk

1. Right-click

and select Computer Management.

2. On the left sidebar of the **Computer Management** window, select **Storage > Disk Management**.

3. Right-click **Disk Management** and select **Rescan Disks**.

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 10 GB to 50 GB).

Extending the file system of an existing partition or creating a partition

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

Extending the file system of an existing partition

Creating a partition

1. Right-click the blank space of the disk and select **Extend Volume**.

2. Follow the **Extend Volume Wizard** to extend the volume.

The new data disk capacity will be added to the original volume.

1. Right-click the free space of the disk and select **New Simple Volume**.

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 Operations

Viewing the Virtio driver version

1. Right-click



and select Device Manager.

2. Unfold Storage Controller items and double-click Tencent VirtIO SCSI controller.

3. Select the **Driver** tab to view the current version, which is 58005 here.

References

Expanding Cloud Disk Capacity Extending Partitions and File Systems (Linux)

Extending Data Disk Partitions and File Systems (Linux) Extending Data Disk Partitions and File Systems Online

Last updated : 2023-12-21 21:34:54

Overview

Expanding a data disk via the console only increases its storage space. You need to extend the partition or file system of the cloud disk to the larger size. This document describes how to extend partitions and file systems online.

Prerequisite

Before extending partition or file system, create a snapshot of the cloud disk to back up data. For more information, see Creating Snapshots.

This practice helps you roll back snapshot to recover data in case of data loss due to maloperation.

The cloud disk has been expanded and attached to a CVM via the console. For more information, see Expanding Cloud Disk Capacity.

The Linux CVM kernel should be version 3.6.0 or later. You can use the uname -a command to check the kernel version.

If the kernel version is earlier than 3.6.0, see Determining the Expansion Method.

Operating Environment

Resource	Description
Operating system	CentOS 8.0 64-bit
Cloud disk (data disk)	 /dev/vdb : Uses the MBR partition and the EXT4 file system, and expands from 50 GB to 60 GB via the console. /dev/vdc : Uses the GPT partition and the XFS file system, and expands from 50 GB to 60 GB via the console.

Directions

Viewing partitions of the cloud disk

- 1. Log in to a Linux instance using standard login method.
- 2. Run the following command to query partitions of the cloud disk.



fdisk -l

The following information appears:
[root@VM-8-61-centos ~] # fdisk -1 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 60 GB $\rm /dev/vdb$ data disk contains a 50 GB MBR partition $\rm /dev/vdb1$.

The 60 GB $\,$ /dev/vdc $\,$ data disk contains a 50 GB GPT partition $\,$ /dev/vdc1 $\,$.

3. Run the following command to determine the file system type of existing partitions.





df -TH

The following information appears:

[root@VM-8-61-centos ~]								
Filesystem	туре	Size	Used	Avail	Use%	Mounted on		
devtmpfs	devtmpfs	946M	0	946M	0%	/dev		
tmpfs	tmpfs	960M	25k	960M	18	/dev/shm		
tmpfs	tmpfs	960M	443k	959M	18	/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	18	/mnt/disk2		
tmpfs	tmpfs	192M	0	192M	0%	/run/user/0		

As shown in the figure,

The	/dev/vdb1	partition is on an EXT4 file system that has been attached to	/mnt/disk1	
The	/dev/vdc1	partition is on an XFS file system that has been attached 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

Ubuntu or Debian





yum install gdisk -y





apt-get install gdisk -y

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

CentOS

Ubuntu or Debian





yum install -y cloud-utils-growpart





apt-get install -y cloud-guest-utils

3. Run the following command to extend partitions using growpart.

 Take extending the /dev/vdb1
 partition as an example. Note that there is a space between /dev/vdb
 and

 1
 in the command. Replace with your actual values.





```
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
7039,end=125829087
```

Extending a file system



1. Use the file system-specific command to resize a file system based on the type obtained in step 3.

Extending an EXT file system

Extending an XFS 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.
```

Run the following command to extend the XFS file system.



xfs_growfs <Mount point>

Take mounting the /dev/vdc1 file system to '/mnt/disk2` as an example, then run the following command:





xfs_growfs /mnt/disk2

The following information will appear:

[root@VM-8-61-centos ~] # xfs_growfs /mnt/disk2							
meta-data	a=/dev/vdc1	isize=512	agcount=4, agsize=3276672 bl				
	=	sectsz=512	attr=2, projid32bit=1				
	=	crc=1	<pre>finobt=1, sparse=1, rmapbt=0</pre>				
	=	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	<pre>sunit=0 blks, lazy-count=1</pre>				
realtime	=none	extsz=4096	<pre>blocks=0, rtextents=0</pre>				

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 ~]								
Filesystem	Туре	Size	Used	Avail	Use∜	Mounted on		
devtmpfs	devtmpfs	946M	0	946M	0%	/dev		
tmpfs	tmpfs	960M	25k	960M	18	/dev/shm		
tmpfs	tmpfs	960M	443k	959M	18	/run		
tmpfs	tmpfs	960M	0	960M	08	/sys/fs/cgroup		
/dev/vda1	ext4	22G	2.1G	19G	11%	/		
/dev/vdb1	ext4	64G	55M	60G	18	/mnt/disk1		
/dev/vdc1	xfs	65G	484M	64G	18	/mnt/disk2		
tmpfs	tmpfs	192M	0	192M	08	/run/user/0		

3. Check data integrity and CVM running status after expansion.

You can roll back the snapshot to recover data in case of exceptions. For more information, see Rolling Back Snapshots.

Extending Data Disk Partitions and File Systems Offline Determining the Expansion Method

Last updated : 2023-12-21 21:36:32

Overview

A cloud disk is an expandable storage device on cloud. After a cloud disk is created, you can expand its capacity at any time to increase its storage capacity without losing any data in it.

After expanding cloud disk capacity on the console, you need to log in to the CVM instance to assign its expanded capacity to an existing partition using an appropriate method as needed. This document describes how to determine the expansion method on a Linux CVM.

Note:

Extending the file system may affect the existing data. We recommend you manually create a snapshot to back up your data before the operation.

Prerequisite

You have expanded the cloud disk via the console. You have mounted the cloud disk to a Linux CVM and created a file system.

You have logged in to the Linux instance that requires extending partitions and file systems.

Directions

1.

Run the following command as the root user to view the partition format of the cloud disk.





fdisk -l

If the result only shows /dev/vdb without a partition, you need to extend the file system.



If the result is as shown in the following two figures (which may vary according to the operating system), the GPT partition format should be used.

Disk /dev/vdb: 32.2 GB, 32212254720 bytes, 7864320 sectors Units = sectors of 1 * 4096 = 4096 butes							
Sector size (logical/physical): 4096 butes / 4096 butes							
I/O size (minimum	voptimal):	4096 bytes ⁄	4096 bytes				
Disk label type:	dos	7	_				
Disk identifier:	0×000000000						
Device Boot	Start	End	Blocks	Id	Sustem		
/dev/vdb1	1	2621439	10485756	ee	GPT		
/dev/vab1	1	2621439	10485756	ee	GPT		

WARNING: GPT (GUID Partition Table) detected on '/dev/vdb'! The util fdisk does

If the result is as shown in the following figure (which may vary according to the operating system), the MBR partition format should be used.

[root@VM_20_49_tlinux ~]# fdisk -1								
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/xvdal *	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. Choose the expansion method corresponding to the partition format obtained in step 1.

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	Expansion method	Remarks
N/A	Extending File Systems on Bare Devices	Applicable to scenarios where a file system is created directly on a bare device and no partition is created.
CDT	Extending GPT Partitions and File Systems (Larger than 2 TB)	Applicable to scenarios of direct formatting when no partition is created.
GPT	Extending GPT Partitions and File Systems (Larger than 2 TB)	Applicable to scenarios where the original partitions remain unchanged and a new GPT partition is created for expansion.
MBR	Extending MBR Partitions and File Systems (Smaller than 2 TB)	Applicable to scenarios of direct formatting when no partition is created.
	Extending MBR Partitions and File Systems (Smaller than 2	Applicable to scenarios where the original partitions remain unchanged and a new MBR partition is created for expansion.



TB)		

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

Last updated : 2023-12-21 21:48:19

Overview

If your cloud disk has a MBR partition that contains the file system, with a disk size of less than 2 TB after expansion, you can use either of the following methods to extend partitions and file systems: Assigning the expanded capacity to an existing MBR partition Formatting the expanded capacity into an independent new MBR partition

Prerequisite

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: The way to expand and partition has been confirmed. For more information, see Determining 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.

Directions

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

The following output indicates there is only one partition. In this case, you can perform the automatic expansion using tools.

[root@	7M-8-57-0	cent	:os ~] #	lsk	olk	
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	

The following output indicates there are two partitions: vdb1 and vdb2 . In this case, you need to choose a partition to be extended as instructed in manual 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				

Automatic Expansion

Manual Expansion

Note:

This method is only applicable to the scenario where there is only one partition. If you have two or more partitions,

choose manual expansion.

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





umount <Mount point>

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





umount /data

2. Run the following command to download an expansion tool.





wget -0 /tmp/devresize.py https://raw.githubusercontent.com/tencentyun/tencentcloud

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





python /tmp/devresize.py <Disk path>

Taking the disk path `/dev/vdb` and the file system `vdb1` as an example, run the following command:





python /tmp/devresize.py /dev/vdb

If The filesystem on /dev/vdb1 is now XXXXX blocks long. is output as follows, the expansion is successful. Then, perform step 4.

[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/ 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

If [ERROR] - e2fsck failed!! is output, perform the following steps:

a. Run the following command to fix the partition where the file system resides.





fsck -a <Partition path>

Taking the disk path /dev/vdb and the file system `vdb1` as an example, run the following command:





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
```

4.

Run the following command to manually mount the extended partition. This document uses the mount point /data as an example.





mount <Partition path><Mount point>

If a partition at the partition path /dev/vdb1 exists before expansion, run the following command:





mount /dev/vdb1 /data

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





df -h

If the result similar to 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	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			

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





ll /data

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



umount <Mount point>

2. Run the following command to extend the vdb2 partition. Replace vdb2 with your actual partition when using the command.




growpart /dev/vdb 2

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





resize2fs /dev/vdb2

If the following output is returned, the file system has been extended.

```
[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 extended partition. This document uses the mount point

/data as an example.



mount <Partition path> <Mount point>

If a partition at the partition path `/dev/vdb2` exists before expansion, run the following command:





mount /dev/vdb2 /data

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





df -h

If the result similar to the following figure is returned, the mounting is successful, and you can see the data disk.

[root@VM-8-	-57-centos	~]# di	E –h		
Filesystem	Size	Used	Avail	Use%	Mounted on
devtmpfs	909M	0	909M	08	/dev
tmpfs	919M	24K	919M	18	/dev/shm
tmpfs	919M	472K	919M	18	/run
tmpfs	919M	0	919M	08	/sys/fs/cgroup
/dev/vda1	9.8G	2.0G	7.4G	21%	/
tmpfs	184M	0	184M	08	/run/user/0
/dev/vdb2	20G	44M	19G	18	/data

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





ll /data

Formatting the expanded capacity into an independent new MBR partition

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





df -h

As shown in the following figure, the mounted partition of the data disk is 20 GB.

[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 data disk that has no partition after expansion:





fdisk -l

As shown in the following figure, the data disk has been expanded to 30 GB.

```
[root@VM-2-4-centos ~] # fdisk -1
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
                2048 41943039 41940992 20G 83 Linux
/dev/vdb1
```

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





umount <Mount point>

Taking the mount point /data as an example, run the following command:





umount /data

Note:

After all the partitions of the cloud disk are unmounted, perform step 4 again.

4. Run the following command to create a partition.





fdisk <Disk path>

Taking the disk path /dev/vdb as an example, run the following command:





fdisk /dev/vdb

Perform the following steps in sequence when prompted.

- 4.1 Enter **p** to check existing partitions, such as /dev/vdb1 in this document.
- 4.2 Enter **n** to create a partition.
- 4.3 Enter **p** to create a primary partition.
- 4.4 Enter 2 to create the second primary partition.
- 4.5 Press **Enter** twice to use the default partition size.



4.6 Enter w to save the partition table and start partitioning.

See 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
                2048 41943039 41940992 20G 83 Linux
/dev/vdb1
Command (m for help): n
Partition type
      primary (1 primary, 0 extended, 3 free)
  p
      extended (container for logical partitions)
   e
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 as needed.

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





fdisk -l

The following figure shows that the new partition vdb2 has been created.

```
[root@VM-2-4-centos ~] # fdisk -1
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.

7. Run the following command to create a mount point.





mkdir <New mount point>

Taking the new mount point /data1 as an example, run the following command:





mkdir /data1

8. Run the following command to manually mount the new partition.





mount <New partition path> <New mount point>

Taking the new partition path /dev/vdb2 and the new mount point /data1 as an example, run the following command:





mount /dev/vdb2 /data1

9. 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/vdb2	9.8G	37M	9.3G	1%	/data1		
/dev/vdb1	20G	44M	19G	1%	/data		

Note:

To allow the CVM instance to automatically mount a data disk upon restart or startup, perform step 10 and step 11 to add the new partition information to /etc/fstab.

10.

Run the following command to add the partition.





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

11. Run the following command to view the partition.





cat /etc/fstab

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



References

Extending Partitions and File Systems (Windows)

FAQs

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

Usage FAQs

Features FAQs

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

Last updated : 2023-12-21 21:52:43

Overview

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

Prerequisite

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: The way to expand and partition has been confirmed. For more information, see Determining the Expansion Method. The file system is EXT or XFS.

The current file system does not have any error.

Directions

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

Taking the disk path /dev/vdc as an example, run the following command:





parted /dev/vdc print

If a message as shown in the following figure appears in the process, 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
(an extra 3963617280 blocks) or continue with the current setting?
Fix/Ignore? Fix
```

The cloud disk size is 2,040 GB after expansion and the existing partition capacity is 10.7 GB, as shown in the



following figure:

Model: N Disk /de Sector s Partitic Disk Fla	Virtio B ev/vdc: size (loc on Table ags:	lock Dev: 2040GB gical/phy : gpt	ice (virt ysical):	512B/	′512B		
Number 1	Start 1049kB	End 10.7GB	Size 10.7GB	File ext4	system	Name opt	Flags

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





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

Taking the disk path /dev/vdc as an example, run the following command:



mount | grep '/dev/vdc'

The following result indicates that the cloud disk has one partition (vdc1) mounted to /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** from the cloud disk.



umount <Mount point>

Taking the mount point /data as an example, run the following command:





umount /data

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



3. Run the following command to use the parted partition tool.




parted <Disk path>

Taking the disk path /dev/vdc as an example, run the following command:





parted /dev/vdc

4. Run the following command to change the unit from the default "GB" to "sector" for display and operation.





unit s

5. Run the following command to view partitions and record their Start values.





print

Note:

Do record Start values. After a partition is deleted and a new one is created, the Start value must remain unchanged. Otherwise, data may be lost.

(parted	(parted) unit s										
(parted	(parted) print										
Model:	Virtio	Block Device	e (virtblk)								
Disk /d	ev/vdc:	3984588800s	3								
Sector	size (l	ogical/physi	ical): 512B/	/512B							
Partiti	on Tabl	e: gpt									
Disk Fl	ags:										
Number	Start	End	Size	File system	Name	Flags					
1	2048s	20969471s	209674243	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: V	irtio H	Block	Device	e (vin	tblk)					
Disk /de	v/vdc:	3984	588800	3						
Sector s	ize (lo	gical	L/phys:	ical):	512B/5	12B				
Partitio	n Table	e: gpt	t							
Disk Fla	Disk Flags:									
Number Start End Size File system Name Flags										
(parted)	(parted)									

Note:

You can immediately run the rescue command, and enter Start and End values as prompted to restore a partition being accidentally deleted.

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





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

The 100% in the command indicates this partition goes to the end of the disk. Enter the Start value obtained in step 5. In this document, the start sector of the original partition is 2048s (that is, the Start value is 2048s), run the following command:





mkpart primary 2048s 100%

If a status as shown in the following figure appears, enter 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.

<pre>(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:</pre>										
Number 1	Start 2048s	End 3984586751s	Size 3984584704s	File system	Name primar y	Flags				

10. Run the following command to close the parted tool.





quit

11. Run the following command to sync the partition table to the operating system.



partprobe

12. Run the following command to check the extended partition.





e2fsck -f <Partition path>

Taking the new partition "1" (its partition path is /dev/vdc1) as an example, run the following command:





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. Use a file system-specific command to resize each file system on the new partition.

Run the following command on the **EXT file system**.





resize2fs <Partition path>

Taking the partition path `/dev/vdc1 as an example, run the following command:



resize2fs /dev/vdc1

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



Run the following command on the XFS file system.



xfs_growfs <Partition path>

Taking the partition path /dev/vdc1 as an example, run the following command:





xfs_growfs /dev/vdc1

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





mount <Partition path> <Mount point>

Taking the partition path /dev/vdc1 and the mount point /data as an example, run the following command:





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	18	/dev/shm				
tmpfs	915M	420K	915M	18	/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

Taking the disk path /dev/vdc as an example, run the following command:





parted /dev/vdc print

If a message as shown in the following figure appears in the process, 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
(an extra 209715200 blocks) or continue with the current setting?
Fix/Ignore? Fix
```

The cloud disk size is 2,147 GB after expansion and the existing partition capacity is 2,040 GB, as shown in the



following figure.

Model: Virtio Block Device (virtblk) Disk /dev/vdc: 2147GB Sector size (logical/physical): 512B/512B								
Partition Table: gpt Disk Flags:								
Number 1	Start 1049kB	End 2040GB	Size 2040GB	File system ext4	Name primar y	Flags		

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





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

Taking the disk path /dev/vdc as an example, run the following command:



mount | grep '/dev/vdc'

The following result indicates that the cloud disk has one partition (vdc1) mounted to /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** from the cloud disk.



umount <Mount point>

Taking the mount point /data as an example, run the following command:





umount /data

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



3. Run the following command to use the parted partition tool.





parted '<Disk path>'

Taking the disk path /dev/vdc as an example, run the following command:





parted '/dev/vdc'

4. Run the following command to view partitions and record their End values, which will be used as the start offset of the next partition.





print

(parted) print								
Model: V	Virtio B	lock Dev	ice (virt	tblk)				
Disk /de	ev/vdc: 2	2147GB						
Sector s	size (log	gical/ph	ysical):	512B/	/512B			
Partitio	on Table	gpt						
Disk Fla	ags:							
Number	Start	End	Size	File	system	Name	Flags	
1	1049kB	2040GB	2040GB	ext4		primary		

5. Run the following command to create a primary partition. This partition starts at the end of existing partitions, and covers all the new space on the disk.





mkpart primary start end

Obtain the End value in step 4. In this example, the End value is 2,040 GB, run the following command:



mkpart primary 2040GB 100%

6. Run the following command to check whether the new partition has been created.





print

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

(parted) mkpart primary 2040GB 100% (parted) print									
Model: V	/irtio Bl	lock Devi	ice (virt	:blk)					
Disk /de	v/vdc: 2	147GB							
Sector s	size (log	jical/phy	<pre>/sical):</pre>	512B/	512B				
Partitic	on Table:	gpt							
Disk Fla	igs:								
Number	Start	End	Size	File	system	Name	Flags		
1	1049kB	2040GB	2040GB	ext4		primary			
2	2040GB	2147GB	107GB			primary			

7. Run the following command to close the parted tool.




quit

8. Run the following command to format the new partition into EXT2, EXT3, etc. as needed.





mkfs.<fstype> <Partition path>

Taking EXT4 as an example, run the following command:





mkfs.ext4 /dev/vdc2

9. Run 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, run the following command:





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-c	entos ~]# df	-h		
Filesystem	Size	Used	Avail	Use €	Mounted on
devtmpfs	902M	0	902M	0%	/dev
tmpfs	915M	24K	915M	18	/dev/shm
tmpfs	915M	424K	915M	18	/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	18	/data

References

Extending Partitions and File Systems (Windows)

FAQs

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

Usage FAQs

Features FAQs

Extending Bare Data Disk File Systems

Last updated : 2023-12-22 09:58:09

Overview

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.

Directions

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-ce	entos ~]#	df -ih]	1			
Filesystem	Type	Inodes	IUsed	IFree	IUse%	Mounted on
devtmpfs	devtmpfs	228K	334	227K	18	/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	18	/run/user/0
/dev/vdb	ext4	640K	11	640K	1%	/data
devtmpfs tmpfs tmpfs tmpfs /dev/vda1 tmpfs /dev/vdb	devtmpfs tmpfs tmpfs tmpfs ext4 tmpfs ext4	228K 230K 230K 230K 640K 230K 640K	334 7 444 16 58K 1 11	227K 230K 230K 230K 583K 230K 640K	1% 1% 1% 10% 10% 1%	/dev /dev/shm /run /sys/fs/cg / /run/user/ /data

The following result shows an XFS file system.

[root@VM-2-4-centos ~] # df -ihT						
Filesystem	Туре	Inodes	IUsed	IFree	IUse≹	Mounted on
devtmpfs	devtmpfs	228K	334	227K	1%	/dev
tmpfs	tmpfs	230K	7	230K	18	/dev/shm
tmpfs	tmpfs	230K	444	230K	18	/run
tmpfs	tmpfs	230K	16	230K	18	/sys/fs/cgroup
/dev/vda1	ext4	640K	58K	583K	10%	/
tmpfs	tmpfs	230K	1	230K	18	/run/user/0
/dev/vdc	xfs	5.OM	3	5.OM	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

Expanding the XFS file system

Run the following command to extend the EXT file system /dev/vdb .





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 ~]#
```



Run the following command to extend the XFS file system $~/{\rm dev}/{\rm vdc}$.



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	a=/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	<pre>sunit=0 blks, lazy-count=1</pre>	
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

Detaching Cloud Disks

Last updated : 2023-12-22 10:01:52

Overview

When you need to mount an elastic cloud disk that is a **data disk** on another CVM, you can unmount this elastic cloud disk from a CVM, and then mount it to other CVMs. **Unmounting an elastic cloud disk does not erase data on this disk**.

Currently, unmounting of elastic cloud disks that are **data disk** is supported. You cannot unmount system disks or non-elastic cloud disks. **To unmount a cloud disk, you must execute umount** (Linux) or offline (Windows) operations. Otherwise, the elastic cloud disk may not be recognized by the CVM next time it is mounted.

Prerequisite

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

Windows

Linux

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.

When detaching an elastic cloud disk, you must first set the disk to offline status. Otherwise, you may not be able to reattach the cloud disk unless you restart the CVM instance.

You must first log In to the instance, and perform a umount operation on the elastic cloud disk you want to unmount. If you directly force unmounting without executing the umount operation, the problem shown in the following figure may occur during shutdown and bootup:

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/v -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 create a logical volume (LV) in the CVM instance, detaching the disk directly in the console will cause part of the device data to remain in the CVM memory. If a CVM application attempts to traverse or access this device, a system error will occur. Therefore, you must first execute the following operations (this example assumes that the

/dev/test/lv1 LV is created based on /dev/vdb1 and is attached to the /data directory):

a. Execute the umount /data command to detach the disk from the corresponding mount point in the CVM instance.

b. Execute the lvremove /dev/test/lv1 command to remove the LV. If there are multiple LVs, remove them one by one.

c. Execute the vgremove test command to remove the VG.

d. Execute the pvremove /dev/vdb1 command to remove the PV.

e. Modify the /etc/fstab file to avoid the continuous mounting of the corresponding LV on next bootup.

Directions

Detaching cloud disks in the console

Detaching cloud disks via API

1. Log in to the CBS console.

- 2. You can use the following method to unmount a cloud disk:
- a. Detach one cloud disk: Select More > Detach on the row of the target cloud disk in Attached status.
- b. Batch detach cloud disks: Select multiple target cloud disks in Attached status and click Detach above the list.

3. In the **Detach cloud disk** pop-up window, confirm the warning and click **OK**.

You can use the DetachDisks API to unmount a cloud disk. For more information, see Unmounting cloud disks.

FAQs

If you cannot detach cloud disks from Windows CVM instances in the console, see Unable to Detach Cloud Disks from CVM Instances for Windows for troubleshooting.

Viewing Cloud Disk Information

Last updated : 2023-12-22 10:03:46

Overview

This document describes how to view your cloud disk information in the console in the following ways: In the CBS console, view the information of all the cloud disks in a region. On the cloud disk details page, view the detailed information of a cloud disk.

Directions

Viewing the cloud disk list

1. Log in to the CBS console and go to the Cloud Disk List page.

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.
 Select

ά

in the top-right corner and configure the fields you want to display in the **Display Settings** pop-up window.

Viewing cloud disk details

1. Log in to the CBS console and select the region of the cloud disk 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 : 2023-12-22 10:04:04

Overview

When a cloud disk is no longer in use and **important data has been backed up**, you can release the virtual resources by terminating the cloud disk. You will not be billed for the cloud disk after termination. When the cloud disk is terminated, all data on the cloud disk will be deleted and cannot be restored. Note that cloud disks that have been terminated cannot be recovered.

The lifecycle of a non-elastic cloud disk is the same as the CVM instance. It can only be terminated when the CVM instance is terminated. For more information, see Terminate Instances.

The lifecycle of an elastic cloud disk is independent of that of the CVM. Therefore, it can be terminated independently. This document describes how to terminate elastic cloud disks.

Elastic cloud disks can be terminated with the following methods:

Manual termination

Automatic 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 entity can return one monthly subscribed cloud disk within five days unconditionally. Each account can return 199 monthly subscribed cloud disks every year. For more information, see Refund. When you hit the return limit, you will not be able to manually terminate monthly subscribed cloud disks.

Manual termination is supported for pay-as-you-go cloud disks, and takes effect immediately.

A monthly subscription cloud disk in the recycle bin will be automatically terminated if it is not recovered within 7 days. It can be continuously used if you renew it within the specified time.

A pay-as-you-go cloud disk will be automatically terminated if your balance becomes negative for more than 24 hours. You can continue using it if you top up your account within the specified time.

Data Wipe

The data deleted during the termination of the cloud disk is inaccessible to anyone, while data deleted from a CBS cloud disk is completely erased. The following mechanism ensures that all data are wiped.

Deleting the logical space of a cloud disk will be recorded as metadata. The physical disk space will be cleared, which forcibly deletes all data permanently. All reads to the logical space return 0.

The metadata will be immediately destroyed when the cloud disk is released to ensure that the data can no longer be accessed. The physical storage space of the cloud disk will be repossessed and cleared before reassigning.

Prerequisite

The cloud disk is in **To be attached** status. For a cloud disk in use, first detach it. **All important data are already backed up**.

Directions

Manually terminating monthly subscribed cloud disks Manually terminating pay-as-you-go cloud disks

Manually terminating unexpired cloud disks in the console

If you no longer need a monthly subscribed cloud disk, you can terminate it manually. Once the status of the cloud disk becomes **To be repossessed**, no fees will be incurred. **The cloud disk is disabled (the cloud disk becomes unavailable, but data is retained) and kept in the recycle bin for seven days. You can recover the cloud disk and continue to use it after you renew it. The cloud disk will be automatically terminated if you don't renew it within seven days. You can also permanently terminate the cloud disk in the recycle bin.** Each entity can return one monthly subscribed cloud disk within five days unconditionally. Each account can return 199 monthly subscribed cloud disk every year. For more information, see Refund. When you hit the return limit, you will not be able to manually terminate monthly subscribed cloud disks.

1. Log in to the CBS console.

2. Terminate cloud disks as instructed below:

a. Terminate one cloud disk: Select More > Terminate/Return on the row of the target cloud disk in To be attached status.

b. Batch terminate cloud disks: Select multiple target cloud disks in **To be attached** status and click **Terminate/Return** above the list.

Caution:

When the cloud disk is terminated, all its data will be deleted and cannot be restored. Terminated cloud disks cannot be recovered. Proceed with caution.

3. In the **Terminate Cloud Disk** pop-up window, select **I have read and agree to Refund Policy** and click **OK**. The system will stop billing, disable the target cloud disk (the cloud disk will be unavailable, but data is retained), and move it to the recycle bin.

Permanently terminate the monthly subscription cloud disk in the recycle bin

You can permanently terminate the monthly subscription cloud disk in the Recycle Bin.

1. Log in to the CBS Recycle Bin.

2. You can use the following methods to permanently terminate a cloud disk:

a. Terminate one cloud disk: Select Release on the row of the target cloud disk in To be repossessed status.

b. Batch terminate cloud disks: Select multiple target cloud disks in To be repossessed status and click Batch

Release above the list.

Caution:

When the cloud disk is terminated, all its data will be deleted and cannot be restored. Terminated cloud disks cannot be recovered. Proceed with caution.

3. Enter the verification code in the pop-up window and click **OK**.

The target cloud disk is **permanently terminated and cannot be recovered**.

1. Log in to the CBS console.

2. Terminate cloud disks as instructed below:

a. Terminate one cloud disk: Select More > Terminate/Return on the row of the target cloud disk in To be attached status.

b. Batch terminate cloud disks: Select multiple target cloud disks in **To be attached** status and click

Terminate/Return above the list.

Caution:

When the cloud disk is terminated, all its data will be deleted and cannot be restored. Terminated cloud disks cannot be recovered. Proceed with caution.

3. In the Terminate Cloud Disk pop-up window, click Submit.

The target cloud disk will no longer be billed. It is permanently terminated and cannot be recovered.

Adjusting Cloud Disk Performance

Last updated : 2023-12-22 10:05:00

The performance of a cloud disk depends on its capacity. You can improve the performance to the maximum level by adjusting its capacity. After the basic performance of an enhanced SSD reaches the ceiling, you can configure and adjust extra performance to achieve even higher performance anytime when prerequisites are met. For more information, see Enhanced SSD Performance.

Note:

Currently, only **Enhanced SSD** supports independent performance adjustment.

The extra performance can be independently adjusted only after the basic performance reaches the ceiling. The performance adjustment will not affect your cloud disks or businesses.

Performance Adjustment Billing

Performance upgrade

For cloud disks featuring monthly subscription: you need to make up the price difference between new and old configurations for the rest of the lifecycle. You can view the actual details on the payment page. For pay-as-you-go cloud disks: the performance upgrade takes effect immediately, and the cloud disks will be charged by the new configuration.

Performance downgrade

For cloud disks featuring monthly subscription: the price difference between the value of new configurations and the value in the rest of the lifecycle is refunded. You can view the actual details on the payment page.

For pay-as-you-go cloud disks: the performance upgrade takes effect immediately, and the cloud disks will be charged by the new configuration.

Performance Upgrade

Performance upgrade via the console

Performance upgrade via an API

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

- 1. Log in to the CBS 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 a target configuration in the pop-up window.

5. Select descriptions and start the adjustment.

You can use the ModifyDiskExtraPerformance API to upgrade performance of a specified cloud disk. For detailed directions, see ModifyDiskExtraPerformance.

Performance Downgrade

Performance downgrade via the console

Performance downgrade via an API

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

- 1. Log in to the CBS 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 a target configuration in the pop-up window.
- 5. Select descriptions and start the adjustment.

You can use the ModifyDiskExtraPerformance API to downgrade performance of a specified cloud disk. For detailed directions, see ModifyDiskExtraPerformance.

Cloud Disk Data Protection Configuration

Last updated : 2023-12-22 10:09:36

Overview

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:

Expiry/Overdue Protection: After expiry/overdue protection is enabled for cloud disks, the system will automatically protect your data by creating a snapshot if your cloud disk is repossessed due to expiration or overdue payments. Snapshot Protection: After snapshot protection is enabled, the scheduled snapshot policy will be associated with a cloud disk by default when it is created.

Note:

For more information on the repossession mechanism and overdue payment processing of cloud disks, see Overdue Policy.

Currently, users in the Chinese mainland can enjoy a free snapshot tier of 80 GB. The snapshot capacity that exceeds the free tier or doesn't meet the free tier policy will incur fees. For more information, see Billing Overview.

Directions

Setting expiry/overdue protection for cloud disks

1. Log in to the CVM console and select Cloud Block Storage on the left sidebar.

2. At the top of the Cloud Block Storage page, select the region of the cloud disk.

3. Configure the parameter in the following ways as needed:

Setting the parameter for one cloud disk

Configuring the parameter for multiple cloud disks

Setting the parameter for all the cloud disks in an AZ

1. Select

in the top-right corner to open the **Display Settings** window.

2. In the **Display Settings** pop-up window, select **Expiry/Overdue Protection** and click **OK**.

Set other fields as needed.

- 3. Toggle on **Expiry/Overdue Protection** on the row of the target cloud disk.
- 1. On the Cloud Disk List page, select cloud disks and click Expiry/Overdue Protection above the list.
- 2. In the **Expiry/Overdue Protection** pop-up window, click **OK**.

1. In the top-right corner of the page, select Cloud Disk Data Protection Configuration.

2. In the Cloud Disk Data Protection Configuration pop-up window, select Expiry/Overdue Protection and

select the AZ from the drop-down list.

3. Click Save.

Setting snapshot protection

1. Log in to the CVM console and select Cloud Block Storage on the left sidebar.

2. At the top of the **Cloud Block Storage** page, select the region of the cloud disk.

3. Configure the parameter in the following ways as needed:

Associating a cloud disk with a scheduled snapshot policy during creation

Associating a cloud disk with the default scheduled snapshot policy on the cloud disk purchase page

1. On the Cloud Block Storage page, click Create above the list.

2. In the **Purchase Data Disk** pop-up window, select **Scheduled Snapshot** and select a scheduled snapshot policy from the drop-down list.

Note:

For more information on other parameters, see Creating Cloud Disks.

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

1. In the top-right corner of the Cloud Block Storage page, select Cloud Disk Data Protection Configuration.

2. In the Cloud Disk Data Protection Configuration pop-up window, select Snapshot Protection for New Cloud Disks and select the AZ from the drop-down list.

3. Click Save.

Note:

When you create a cloud disk in the selected AZ, **Scheduled Snapshot** will be selected by default, that is, the cloud disk will be associated with the policy by default. You can deselect the option as needed.

References

Creating cloud disks Scheduled Snapshots

Cloud Disk Data Backup Point Adjusting Data Backup Point Quota

Last updated : 2023-12-22 10:14:50

Overview

For cloud disks with the data backup point quota configured, Tencent Cloud will automatically create a data backup point for data backup. You can set the quota when purchasing a cloud disk or adjust/return the quota for an existing cloud disk as needed.

This document describes how to adjust the data backup point quota of a cloud disk and the billing details after the adjustment.

Note:

Currently, only one data backup point is supported.

The billing of the data backup point quota is only related to the size of the cloud disk.

When the number of data backup points becomes 0 after adjustment/return, the data backup point of the existing data disk will be deleted. To retain a backup point, you can convert it into a snapshot as instructed in Converting Data Backup Point into Snapshot.

Billing After the Adjustment

Increasing the data backup point quota Reducing/Returning the data backup point quota

Billing rules

For cloud disks in different billing modes:

Monthly subscribed cloud disks: To increase the data backup point quota, you need to make up the price difference between new and old configurations for the rest of the lifecycle, subject to the details as displayed on the payment page.

Pay-as-you-go cloud disks: The adjustment takes effect immediately, and the cloud disks will be billed by the new configuration.

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 rules

For cloud disks in different billing modes:

Monthly subscribed cloud disks: After the data backup point quota is reduced, the price difference between the value of the new configuration and the value in the rest of the lifecycle will be refunded, subject to the details as displayed on the payment page.

Pay-as-you-go cloud disks: The adjustment takes effect immediately, and the cloud disks will be billed by the new configuration.

Billing rules for monthly subscribed cloud disks:

Refund for downgrade = fees of original specification - fees of new specification. For more information, see Refund. If the refund is greater than 0, the data backup point quota will be reduced, and fees will be returned to your Tencent Cloud account accordingly.

If the refund is equal to or smaller than 0, the data backup point quota will be reduced, but no fees will be returned. If you use a discount or voucher for purchase, the discount amount and voucher will not be refunded.

Directions

Increasing the data backup point quota

You can increase the data backup point quota in the following ways:

Increasing the data backup point quota in the console

Increasing the data backup point quota via an API

If the use limits are met, you can increase the data backup point quota in the following ways:

1. Log in to the CBS console and select the region of the cloud disk at the top of the page.

2. Find the target cloud disk in the list and select More > Adjust Data Backup Point Quota on the right.

3. In the Adjust Data Backup Point Quota pop-up window, select the target configuration.

4. Click Next, confirm the billing details, and select I have read and agree to the billing details for adjusting the data backup point quota.

5. Click OK.

You can use the ModifyDiskBackupQuota API to adjust the data backup point quota of a specified cloud disk.

Reducing/Returning the data backup point quota

You can reduce/return the data backup point quota in the following ways:

Reducing/Returning the data backup point quota in the console

Reducing/Returning the data backup point quota via an API

If the use limits are met, you can reduce/return the data backup point quota in the following ways:

1. Log in to the CBS console and select the region of the cloud disk at the top of the page.

2. Find the target cloud disk in the list and select More > Adjust Data Backup Point Quota on the right.

3. In the Adjust Data Backup Point Quota pop-up window, select the target configuration.

4. Click Next, confirm the billing details, and select I have read and agree to the billing details for adjusting the

data backup point quota.

5. Click OK.

You can use the ModifyDiskBackupQuota API to adjust the data backup point quota of a specified cloud disk.

Using Data Backup Point to Restore Cloud Disk Data

Last updated : 2023-12-22 10:15:54

Overview

The data backup point of a cloud disk provides continuous data backup service for the cloud disk. When you have the data backup point, you can restore cloud disk data to a historical time point to mitigate data loss caused by viruses, intrusions, and maloperations. This document describes how to restore cloud disk data to a backup time point with an existing data backup point.

Notes

Keep the following in mind when using the data backup point to restore data:

You need to shut down the CVM instance. As this will affect business continuity, we recommend you make preparations and perform the operation during off-peak hours.

The operation will clear cloud disk data after the data backup point. Evaluate the impact in advance.

Directions

1. Log in to the CBS console and select the target 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 **Historical Data Backup Point** tab, find the target data backup point, and click **Roll back to Backup Point** on the right.

4. In the **Roll back to Backup Point** pop-up window, confirm the information. If the CVM instance to which the cloud disk is attached is running, you also need to indicate your consent to CVM shutdown.

5. Click **OK**.

Converting Data Backup Point into Snapshot

Last updated : 2023-12-22 10:16:38

Overview

The data backup point of a cloud disk follows the lifecycle of the cloud disk. When the cloud disk expires or is returned, the data backup point reaches the end of its lifecycle. You can convert the data backup point on a specific date into a snapshot, so that you can make it independent of the cloud disk lifecycle and retain a data backup for a long time.

This document describes how to convert a data backup point into a snapshot in the console.

Notes

Snapshots have been commercialized. If the data backup point is converted into a snapshot, a snapshot will be created and billed. For billing details, see Price Overview.

After being converted into a snapshot, the data backup point will be deleted automatically and no longer occupy the quota.

Directions

- 1. Log in to the CBS console and select the target 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 Historical Data Backup Point tab, find the target data backup point, and select

More > Convert Backup Point into Snapshot on the right.

4. In the Convert Backup Point into Snapshot pop-up window, enter the Snapshot Name.

5. After confirming that everything is correct, click **OK**.

After the operations, you can view the configuration on the Snapshot List page.

Deleting Data Backup Point

Last updated : 2023-12-22 10:18:21

Overview

For cloud disks with the data backup point quota configured, Tencent Cloud will automatically create and manage a data backup point for data backup, so that the backed up data can stay up to date. For more information, see Data Backup Point Quota of Cloud Disk.

The data backup point is free of charge, and retaining it doesn't incur fees. This document describes how to delete the data backup point.

Note:

Deleting the data backup point cannot be undone, and deleted data backups cannot be recovered.

Directions

1. Log in to the CBS console and select the target 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 **Historical Data Backup Point** tab, find the target data backup point, and click **Delete** on the right.

4. In the **Delete Data Backup Point** pop-up window, click **OK**.

Managing Snapshots Creating Snapshot

Last updated : 2023-12-22 10:19:47

Overview

You can create snapshots for a cloud disk to save its data at a specific point of time. The incremental snapshot only records the data change compared to the last snapshot. This process is quick if the data changes a little. Moreover, deleting a snapshot will not affect your use of snapshot data. You can also restore your cloud disks by using any remaining snapshots.

You can create a snapshot for a cloud disk in any state, but the snapshot can only capture the written data rather than data being written by an app or process. Before creating a snapshot, you can choose to suspend all disk I/O operations, or detach the cloud disk and attach it later to create a complete snapshot.

Prerequisite

You have created a cloud disk.

You've not reached the upper limits for the number and total size of snapshots in the current region. See Snapshot Use Limits.

Notes

A snapshot can only capture the written data but not cached data in the memory (such as files in the /run directory on a Linux CVM) of the cloud disk. We recommend that you shut down the instance, or write all memory data to disk and suspend disk I/O operations before creating a snapshot. This can be done at the following two levels.

Database level

For database services, we recommend locking all tables in databases as read-only to ensure that the snapshot captures all data. Taking MySQL as an example, the procedure is as follows:

1. Run the command FLUSH TABLES WITH READ LOCK to close all tables and use the global read lock to lock all tables in every database, as shown below:

mysql> flush tables with read lock; Query OK, 0 rows affected (0.01 sec)

- 2. Create a snapshot for the cloud disk.
- 3. Run the command UNLOCK TABLES to unlock the tables, as shown below:



System level

For better system performance, data is stored in the memory buffer before it is written to the cloud disk at the proper moment. Therefore, the snapshot created for the cloud disk does not contain data stored in the memory buffer and not already written to the cloud disk. As a result, data inconsistency occurs.

To resolve this issue, run the sync command to forcibly write the data in the file system memory buffer immediately to the cloud disk, and then prevent new data from being written to the cloud disk. If no error message is returned after the command is executed, the data in the memory buffer has been successfully written to the cloud disk, as shown below:

ubuntu@VM-30-151-ubuntu:~\$ sync ubuntu@VM-30-151-ubuntu:~\$

Directions

Creating a snapshot via the console

- 1. Log in to the CBS console.
- 2. Click Create a snapshot on the right of the target cloud disk.
- 3. In the Create a snapshot pop-up window, enter the snapshot name and click OK.

Creating a snapshot via an API

You can use the CreateSnapshot API to create a snapshot. For more information, see CreateSnapshot.

Rolling back Data from Snapshot

Last updated : 2023-12-22 10:20:44

Overview

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.

Directions

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.

Rolling back the data via an API

You can use the ApplySnapshot API to roll back the data. For detailed directions, see ApplySnapshot.

Creating Cloud Disks Using Snapshots

Last updated : 2023-12-22 10:23:44

Overview

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 you use a snapshot to create a data disk of the same size as the snapshot, the new data disk doesn't need to be initialized. You can directly attach it, right-click

> **Disk Management**, and set it to **Online** to bring it online in the CVM instance, after which it can be normally read and written.

When you use a snapshot to create a data disk whose capacity is greater than that of the snapshot, the system only expand the storage block and does not extend the file system or convert the partition format. After you mount the new data disk, it can only use the file system and data of the source snapshot and cannot use the new disk space. You need to manually extend the file system and convert the partition format.

For example, if you want to a 3 TB data disk by using a data disk snapshot that uses the MBR partition format and has a capacity of 1 TB, you need to format the data disk in GPT partition style because the maximum disk space supported in MBR partition style is 2 TB. Note that this operation will **delete original data**.

This document describes how to use a snapshot to create a cloud disk on the Snapshot List page. When creating a cloud disk, you can configure the **Snapshots** parameter to specify a snapshot for creating the cloud disk.

Directions

Creating a cloud disk with a snapshot in the console

- 1. Log in to the Snapshot List page.
- 2. On the row of the target snapshot, select **More** > **Create Cloud Disk**.
- 3. In the Purchase Data Disk pop-up window, configure the following parameters:

Parameter	Configuration
AZ	Required. The availability zone where the created cloud disk resides. It cannot be modified after the cloud disk is created.



Cloud Disk Type	Required. For more information on cloud disk types, see Cloud Disk Types.
Capacity	Required. For more information on the cloud disk capacity and specifications, see Performance metrics. When you create a cloud disk with a snapshot, the disk capacity cannot be smaller than that of the snapshot. If you don't specify the capacity of the cloud disk, the capacity is equal to that of the snapshot by default.
Snapshots	Optional. To use a snapshot to create a cloud disk, select Create a Cloud Disk with a Snapshot and select the required snapshot. The disk capacity is the same as that of the snapshot by default. You can adjust the capacity to be greater than that of the snapshot. The disk type is the same as that of the snapshot by default. You can change the cloud disk type.
Disk Name	Optional. A maximum of 20 characters are supported. It must start with a letter and can be a combination of letters, digits, and special symbols (`.`, `_`,` and `-`). This parameter can be modified after the cloud disk is created. If you create only one cloud disk, the disk name is the name of the cloud disk you create. If you create multiple cloud disks at one time, the disk name entered will be used as he prefix of the final disk name, in the format of disk name_number , for example, "disk name_0" to "disk name_49".
Projects	Required. When creating a cloud disk, you can configure the project to which the cloud disk belongs. The default value is Default Project.
Label	Optional. When creating a cloud disk, you can bind a tag to it. A tag is used for identification, helping you categorize and search cloud resources. For more information on tags, see Tag.
Billing mode	Required.


	The cloud disk supports the following billing modes: Monthly subscription. If this mode is selected, you must set the Purchase Period . Pay-as-you-go.
Scheduled Snapshot	Optional. When creating a cloud disk, you can select Scheduled Snapshot to create snapshots for the cloud disk periodically based on the created scheduled snapshot policy. For more information, see Scheduled Snapshot.
Quantity	Optional. The default value is 1, meaning only one cloud disk is created. Currently, up to 50 cloud disks can be created at a time.
Validity Period	If the Billing mode is Monthly Subscription, this parameter is required. Value range: One month to five years. If the Billing Mode is Pay-as-you-go, this parameter is not involved.
Automatic Renewal	If the Billing Mode is Monthly Subscription, this parameter is optional. If you select Auto-renewal , auto-renewal will be performed on the device monthly upon expiration if your account has a sufficient balance. If the Billing Mode is Pay-as-you-go, this parameter is not involved.

4. Click OK.

If the Billing Mode is Pay-as-you-go, the creation is completed.

If the Billing Mode is Monthly Subscription, the Information Verification page will be popped up.

4.1.1 After you confirm the specifications, select whether to use a voucher as needed and click **Submit Order**.

4.1.2 Complete the payment.

5. You can view the cloud disk you created in the Cloud Disk List. The new elastic cloud disk is in the **To be attached** status. To attach it to a CVM instance in the same AZ, see Attaching Cloud Disks.

Using an API to create a cloud disk from snapshot

You can use the CreateDisks API to create a cloud disk. For more information, see CreateDisks.

Cross-region snapshot replication

Last updated : 2023-12-22 10:43:54

Cross-region snapshot replication is now in beta. With this feature, you can easily migrate data and services to other regions, or build a cross-region disaster recovery system for your business. Cross-region snapshot replication is now in beta test. To try it out, submit a ticket for application.

Use Limits

Application for beta test eligibility: Cross-region snapshot replication is now in beta test. To try it out, submit a ticket for application.

Supported regions: For more information, see Regions and Availability Zones.

Finance zones: Finance zones support replication between finance zones only.

System disk snapshots are not supported.

Directions

1. Log in to the Snapshot List page.

2. Click Cross-region replication for the target snapshot.

3. Configure the following parameters:

New snapshot name: (optional) enter the new snapshot name with up to 60 characters.

By default, a new snapshot name contains the source snapshot ID and region information in the format of Copied <Source snapshot ID> from <Source snapshot region>, such as Copied snap-oi5spwt2 from ap-shanghai.

Region: (required) a target region to which a snapshot is copied

Check the snapshot quota and geographical restriction when you select the region.

4. Click **OK** to start the replication. Hover over the information icon to view the status of the source snapshot. The new snapshot is added to the target region.

5. Once the replication is completed, you can view the new snapshot in the snapshot list of the target region.

No:

The source snapshot cannot be deleted during the cross-region replication of the snapshot.

During the process of cross-region replication:

Status of source snapshot: You can view it by going to the source region's snapshot list and looking in the status column on the source snapshot's row.

Status of target snapshot: You can view it by going to the snapshot list page of the target region.

Deleting Snapshots

Last updated : 2023-12-22 10:44:46

Overview

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 a snapshot has been associated with an image, delete the image before deleting the snapshot.

Deleting a snapshot will also delete all data in the snapshot, and the data cannot be retrieved. Deleted snapshots cannot be restored, so delete snapshots with caution.

Directions

Deleting a snapshot in the console

Deleting a snapshot via API

1. Log in to the Snapshot List page.

2. Delete a snapshot as follows:

a. Delete one snapshot: Click Delete on the row of the target snapshot.

b. **Batch delete snapshots**: Select all of the snapshots you want to delete (make sure the snapshots are not involved in any tasks) and click **Delete** at the top of the list.

3. Click **Confirm**.

You can use the DeleteSnapshots API to delete a snapshot. For detailed directions, see DeleteSnapshots.

Monitoring and Alarms

Last updated : 2023-12-22 10:46:34

A good monitoring environment is the key to ensuring high availability of cloud disks. You can use Cloud Monitor (CM) to monitor cloud disks attached to CVM instances, view cloud disk metrics, and analyze and configure cloud disk alarms. CM collects raw data of the cloud disk from running CVM instances and displays data through easy-to-read charts. By default, the data is retained for one month, during which you can view the disk conditions in different time periods to better know its usage and read/write operations.

You can get data in the CM console or via API. For more information, see Getting the Monitoring Data of Specified Metrics and Getting Monitoring View and Chart.

Metric	Name	Calculation	Description	Unit	Statistical Granularity (Period)
DiskReadlops	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	Count	10s, 60s, 300s
DiskReadTraffic	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
DiskWritelops	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	Count	10s, 60s, 300s
DiskWriteTraffic	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	Percentage of time during	ms	10s, 60s, 300s

Currently, CM provides the following metrics to monitor cloud disks:



		disk during the statistical period	which the CPU is free and there are unfinished I/O requests during the statistical period		
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 more information on monitoring metrics, see Cloud Monitor.