

CBS

Product Documentation





Document Contents

Product Overview

- Product Overview

- Typical Usage Scenarios

- Introduction to Cloud Disk Types

- Product Functions

Product Functions

Product Advantages

- Reliability

- Strengths

- Ultra-High Performance

- Ease to use

- Snapshot Backup

Usage Guidelines

Beginner's Guide

- Beginner's Guide

- Prerequisites

- Creating a New Cloud Disk

- Attaching the Cloud Disk to a Cloud Server Instance

 - Attaching the Cloud Disk to a Cloud Server Instance

 - Connecting to Windows-type Instances

Lifecycle of Cloud Disks

- Creating a Cloud Disk

 - Creating a Cloud Disk

- xx

- Attaching Cloud Disks

 - Attaching Cloud Disks

 - Attaching Elastic Cloud Disks to Instances via Console

 - Solutions for Elastic Cloud Disks Not Recognized by Some Created

 - When Launching New Instances Using Custom Images and Data Disk S

Initialize Cloud Disks

- Initializing Cloud Disks on Windows Systems

 - Initializing Cloud Disks on Windows Systems

 - Prerequisites

 - Partitioning and Formatting Disks in Windows 2012

 - Partitioning and Formatting Disks in Windows 2012



When Launching a New Instance Using a Custom Image and Data Disk

Expanding Cloud Disks

Expanding Data Disks of the Cloud Disk Type

Expanding System Disks of the Cloud Disk Type

Expanding Cloud Disks

Expanding the File System

Expanding the Linux File System

Expanding Windows File Systems

Unmounting an Elastic Cloud Disk

Unmounting an Elastic Cloud Disk

Snapshots

What is a Snapshot

Monitor Cloud Disks

Best Practices

Best Practices

Frequently Asked Questions



Product Overview

Product Overview

Last Updated At: 2025-08-12 15:58:03

CBS provides efficient and durability data block-level storage services for Cloud Virtual Machines. CBS is a block storage device with high availability, high reliability, low cost and customization, which can be used as an independent and scalable hard disk of Cloud Virtual Machine. It provides data storage at the data block level and adopts a three-copy distributed mechanism to provide data reliability for CVM.

CBS supports automatic replication in the availability zone and multi-replica redundancy storage to back up your data on different machines, thus avoiding problems such as data loss caused by a single machine failure and improving the availability and durability of data.

It is usually used as the main storage device for data (such as file system and database) that needs frequent and fine-grained updates. It provides data storage at the data block level and adopts a three-copy distributed mechanism to provide data reliability for CVM.

The lifecycle of the CBS is independent from the CVM instance and is not affected by the instance running time. In addition, you can connect multiple CBS to the same instance, or disconnect them from the instance and mount them to another instance.

The lifecycle of a cloud disk is independent of the CVM instance and is not affected by the instance's running time. Additionally, users can attach multiple cloud disks to the same instance, or detach a cloud disk from an instance and attach it to another instance. Cloud disks support mounting and detaching on cloud servers in the same availability zone at any time, and can be used as data disks by the cloud server after mounting.



Typical Usage Scenarios

Last Updated At: 2025-08-11 15:29:04

If the hard disk space of a cloud server becomes insufficient during use, you can apply for one or more cloud disks and attach them to the cloud server to meet storage capacity requirements. No additional storage space is needed when applying for a cloud server, and cloud disks can be applied for later to expand the server's storage capacity when storage needs arise. When there is a need for data exchange between multiple servers, this can be achieved by detaching the cloud disk (data disk) and reattaching it to other servers. You can apply for multiple cloud disks and configure LVM logical volumes to break through the maximum storage capacity limit of a single cloud disk. You can apply for multiple cloud disks and configure a RAID policy to break through the maximum IO capability limit of a single cloud disk.



Introduction to Cloud Disk Types

Last Updated At: 2025-08-12 16:28:50

High-performance cloud disks

The Premium Cloud Disk is a CBS product that combines low cost and high performance by analyzing massive access data and adopting hot and cold data partitioning technology based on the operation and use experience of massive distributed storage. It uses the SSD+HDD hybrid hard disk model to improve the hard disk performance at a relatively low cost by means of SSD cache and large IO overwriting.

SSD cloud disks

SSD Cloud Disk is based on all-SSD storage medium and uses CBS three-copy distributed storage technology to provide low-latency, high-random IOPS, high-throughput I/O capabilities and high-performance storage, which can be used as an independent scalable hard disk for CVM.



Product Functions

Last Updated At: 2025-08-12 16:28:50

CBS can be divided into Premium Cloud Disk, SSD Cloud Disk according to different performances. Different types of hard disks have different performance characteristics. Users can select storage performance based on the deployed application requirements.

CBS offers elastic capacity configuration, allowing you to upgrade the capacity as needed. Up to 20 cloud disks, maximum 16TB per disk, can be mounted on a single virtual machine. Scaling can be done without interrupting services

You can back up your data by taking a snapshot of the point in time of CBS to prevent data loss caused by tampering and accidental deletion, and ensure quick rollback in case of business failure.



Product Functions

Last Updated At: 2025-08-12 16:28:50

- CBS offers a variety of storage devices, allowing users to flexibly choose disk types and perform operations such as storing files and setting up databases on the disks themselves.
- flexible Mounting/Unmounting: Standard cloud disks support flexible mounting and unmounting, enabling the creation of large-capacity file systems for a single cloud server.
- Elastic Expansion: A single disk supports up to 16TB of storage, and you can expand the disk capacity at any time.
- Snapshot Backup: Supports snapshot creation and rollback for timely backup of critical data. Snapshots can also be used to create disks, allowing rapid business deployment.



Product Advantages

Reliability

Last Updated At: 2025-08-02 14:35:02

Before each store write request is returned to the user, CBS ensures that data has been successfully written into three cross-rack storage nodes. The data replication mechanism of CBS is higher than the average level in the industry, which can ensure quick data migration and recovery when any replica fails to protect your applications from component failures.



Strengths

Last Updated At: 2025-08-02 14:35:02

You can freely configure the storage capacity and scale out as needed to meet your business data scale-out needs without interrupting services, thus efficiently coping with big data processing scenarios of TB/PB-level data.

The maximum capacity of a single disk is 32000 GiB. A single CVM can mount 20 elastic cloud disks in total, and the storage capacity of a single CBS cluster has no upper limit.



Ultra-High Performance

Last Updated At: 2025-08-02 14:35:02

The SSD Cloud Disk adopts NVMe standard Premium Cloud Disk. A single disk provides a maximum of 26,000 random IOPS and 260MB/s throughput to provide industry-leading IOPS and throughput to realize the integration of ultra-high performance and ultra-high reliability. Premium SSD Cloud Disk helps users easily support high-throughput DB access on the business side.



Ease to use

Last Updated At: 2025-08-02 14:35:02

The cloud disk can be managed and used through simple operations such as creating, mounting, uninstalling and deleting, to save the cost of manual management and deployment.



Snapshot Backup

Last Updated At: 2025-08-02 14:35:02

You can back up your data by taking a snapshot of the point in time of CBS to prevent data loss caused by tampering and accidental deletion, and ensure quick rollback in case of business failure.



Usage Guidelines

Last Updated At: 2025-08-12 16:28:50

1. **System Disk:** Each system disk supports a capacity range of 50GB–1024GB, with increments of 1GB. Elastic cloud disks (data disks) support a capacity range of 10GB–16,000GB, with increments of 10GB.
2. **Mounting Limits:** A single cloud server can mount up to 20 elastic cloud disks (data disks). A maximum of 10 data disks can be created and mounted simultaneously with the cloud server.
3. **Creation Limits:** Up to 50 elastic cloud disks (data disks) can be created in a single operation.
4. **Availability Zone Requirement:** Cloud servers and cloud disks must reside in the same availability zone.
5. **Performance Considerations:** During disk read/write operations, either IOPS or throughput may reach the product’s maximum limit first, depending on the block size used in the I/O operations.

Performance Specifications:

1. High-Performance Cloud Disk:

- **Maximum Random IOPS:** 6,000
 $\text{Achievable IOPS} = \text{MIN}(1,800 + \text{Storage Capacity (GB)} \times 8, 6,000)$
- **Maximum Throughput:** 150 MB/s
 $\text{Achievable Throughput} = \text{MIN}(100 + \text{Storage Capacity (GB)} \times 0.15, 150) \text{ MB/s}$

2. SSD Cloud Disk:

- **Maximum Random IOPS:** 26,000
- **Maximum Throughput:** 260 MB/s
- **Example:**
 - For a 1TB SSD cloud disk:
 - With 4KB/8KB block size I/O, IOPS reaches the max of 26,000 first (throughput does not reach 260 MB/s).
 - With 16KB block size I/O, throughput reaches the max of 260 MB/s first (IOPS does not reach 26,000).

Additional Guidelines:

6. Snapshot Rollback:

- Snapshots can only be rolled back to their original source disk.



- System disk snapshots can be rolled back when the cloud server is shut down.
- Data disk snapshots can be rolled back when the cloud server is shut down or the disk is unmounted.

7. Creating Disks from Snapshots:

- Only data disk snapshots can be used to create new data disks.
- The new disk's capacity must be $\geq$ the original snapshot's disk capacity.

8. Snapshot Content:

- Snapshots only store disk data, not host memory data.

9. Availability Zone Restrictions:

- A cloud disk's availability zone cannot be changed after creation.
- Cross-availability zone failover migration is not supported.

10. Backup Best Practices:

- The OS on cloud servers is not backed up by default.
- For critical data backup, use the cloud disk snapshot feature (requires approval during feasibility study).
- Set a reasonable backup frequency.
- The ratio of cloud disks (system + data disks) to snapshots per user should not exceed 1:4.
- A single scheduled snapshot policy can bind $\leq$ 500 disks.
- Avoid mass snapshot creation at the same time to prevent failures.
- Regularly validate backup data integrity.



Beginner's Guide

Beginner's Guide

Last Updated At: 2025-08-02 14:56:50

This document will guide you through the basic usage of CCB Cloud Disk by demonstrating how to:

1. Create an empty elastic cloud disk named `cbs-test` in the Wuhan region(for details on elastic cloud disks, refer to Cloud Disk Types ; for creating a disk with existing data, see Creating a Disk from a Snapshot).
2. Attach the disk to a cloud server instance.
3. Create a file system and write a file named `yunccb.txt` to it.

Follow these steps to get started with CCB Cloud Disk.



Prerequisites

Last Updated At: 2025-08-02 14:56:50

Before proceeding with the cloud disk operations described in this document, you must have a running cloud server instance in the same region and availability zone (Wuhan Zone 1) where the cloud disk will be created.

If you don't have a suitable instance, refer to [\[Apply for and Launch a Cloud Server Instance\]](#)for setup instructions.



Creating a New Cloud Disk

Last Updated At: 2025-08-02 14:56:50

This example demonstrates how to create a standard elastic cloud disk via the console. For more details, refer to [Creating a Cloud Disk].

Steps:

1. Log in to the Cloud Disk Console and click "Create New".
2. Configure the Disk:
 - Region/AZ: Select Wuhan Zone 1.
 - Disk Type: Choose Cloud Disk (standard).
 - Capacity: Set to 20GB.
 - Click "Confirm" to proceed.
3. Complete Creation:
 - On the next page, click "Finish" to complete the process.
 - The new disk will appear in the disk list with a default unnamed status and "To Be Mounted" state.
4. Rename the Disk:
 - Click the edit icon next to the disk and name it `cbs-test` .



Attaching the Cloud Disk to a Cloud Server Instance

Attaching the Cloud Disk to a Cloud Server Instance

Last Updated At: 2025-08-02 14:56:50

1. Log in to the CCB Cloud Console.
2. Navigate to:
Cloud Server → Cloud Disk tab.
3. Locate the target disk (cbs-test) in the disk list, then click:
More → Attach to Cloud Server.
4. In the pop-up window:
 - Select the target cloud server (ensure it meets the [prerequisites](#)).
 - Click Confirm to attach the disk.

Formatting, Creating a File System, and Writing Files (Linux Instance)

Step 1: Identify the Disk Device
5. Run the following command to list all attached disks:
bash fdisk -l
 - Locate your 20GB cloud disk (e.g., `/dev/vdb`).

Step 2: Format the Disk (EXT4 File System)

2. Format the disk using:

bash mkfs.ext4 /dev/vdb (Wait for the formatting to complete.) Step 3: Mount the Disk 3. Create a mount point (e.g., `/data`) and mount the disk:

bash mkdir /data mount /dev/vdb /data Step 4: Write a Test File

4. Navigate to the mount point and create a file:

bash cd /data vi yunccb.txt – Press `i` to enter edit mode, then add content (e.g., `"This is my first test"`).

- Press `ESC` , then type `:wq` to save and exit.

5. Verify the file exists:

ls



Connecting to Windows-type Instances

Last Updated At: 2025-08-02 14:56:50

Navigate to Disk Management via a path similar to ([Start] –) [Server Manager] – [Storage/Computer Management] – [Disk Management]. Right-click on the created empty elastic cloud disk and select [Online]. Right-click on the disk, choose [New Simple Volume] from the pop-up shortcut menu, and follow the wizard prompts. Enter the size of the entire disk, click [Next], select the file system (when the disk is larger than 2TB, be sure to choose the GPT partition format), format the partition, click [Next], and then click the [Finish] button. A message indicating that formatting is in progress will appear. Once formatting is completed, the elastic cloud disk can be used. Access the created disk, create a new file named yuncb.txt, enter the content you need, and click the [Save] button. This completes the file writing operation. For more information about system partitioning, formatting, online operations, and file system creation under the Windows system, please refer to the chapter "Initializing Cloud Disks under Windows System".



Lifecycle of Cloud Disks

Creating a Cloud Disk

Creating a Cloud Disk

Last Updated At: 2025-08-12 16:28:50

You can create a cloud disk and attach it to any CVM instance within the same availability zone. Through block storage device mapping, the cloud disk is recognized and used by the cloud server instance. Additionally, users can launch a new cloud disk based on a previously created snapshot. For more information, refer to [Creating a Cloud Disk from a Snapshot](#). A cloud disk can reach its maximum performance immediately after creation without the need for preheating. The method of creating a cloud disk varies depending on its type. For more details about cloud disk types, please refer to [Classification of Cloud Disks](#). [Creating Elastic Cloud Disks\[菜单建错级别\]](#) Creating Elastic Cloud Disks via Console Log in to the Cloud Disk Console and click to start the application. In the pop-up window, select the region/availability zone, capacity, and quantity of disks to apply for, then click the [OK] button. Click [Confirm] on the page to complete the application. You can view the applied cloud disks on the Cloud Disk List page. The elastic cloud disk you just applied for is unnamed by default and displayed in the "To Be Mounted" status. [Creating Elastic Cloud Disks Using Snapshots](#) You can choose this method if you need to retain the data from a data disk snapshot in the newly created disk. Log in to the Snapshot Console, find the snapshot you want to use, and click [Create New Cloud Disk] to start the application.

<http://imgxxfb.yun.ccb.com/raw/336381ffa6c46431353ea082e65a890d.png>

In the pop-up window, select the region/availability zone, capacity, and number of disks to apply for, then click the [OK] button. Click [Confirm] on the page to complete the application. You can view the applied cloud disks on the Cloud Disk List page. The elastic cloud disk you just applied for is named "From snap-xxxxxxx" by default and displayed in the "To Be Mounted" status.



XX

Attaching Cloud Disks

Attaching Cloud Disks

Last Updated At: 2025-08-02 15:07:35

Non-elastic cloud disks (i.e., cloud disks whose lifecycle is tied to the CVM instance) are automatically attached to the CVM instance upon creation and cannot be reattached to another instance. Meanwhile, users can manually attach elastic cloud disks to any instance within the same availability zone. During the attachment process, you can check the remaining number of cloud disks that each instance can still accommodate.



Attaching Elastic Cloud Disks to Instances via Console

Last Updated At: 2025-08-02 15:07:35

Currently, only ordinary elastic cloud disks used as data disks can be attached. System disks cannot be attached.

Log in to the CCB Cloud Console.

Navigate to the [Cloud Server] – [Cloud Disks] tab.

Single Disk Attachment: On the Cloud Disks list page, click [More] – [Attach to Cloud Host] for the disk with a status of "To Be Attached" and support for mounting/unmounting.

Batch Attachment: Select multiple disks with the status of "To Be Attached" and support for mounting/unmounting, then click [Attach] at the top of the page.

In the pop-up window, select the target cloud server and click [OK]. Wait for the attachment to complete, then log in to the cloud server to check the attachment status.

After attaching a cloud disk, it cannot be used immediately. You need to perform operations such as partitioning and formatting. For specific instructions, refer to:

Initializing Cloud Disks in Linux Systems

Initializing Cloud Disks in Windows Systems



Solutions for Elastic Cloud Disks Not Recognized by Some Created

Last Updated At: 2025-08-02 15:07:35

All currently provided images support the connection and unmounting operations of elastic cloud disks. Please note that before unplugging (unmounting) the disk, you must first perform the umount operation (for Linux) or offline operation (for Windows). Otherwise, the disk may not be recognized when remounted.

Automatically Mount Data Disks When Launching New Instances Using Custom Images and Data Disk Snapshots (Linux) When launching a new cloud server instance, if the user specifies a custom image and a data disk snapshot, China Construction Bank Cloud's cloud disks can support automatic mounting after the cloud server instance is started (that is, there is no need to perform a series of operations such as adding, partitioning, formatting, etc., and the data disk can be read and written directly). Users need to perform some operations on the original instance before creating the custom image and data disk snapshot, which will be described in detail below. If users want the cloud disks created from the specified data disk snapshots to be automatically mounted to new cloud server instances under the Linux system, the specified custom image and data disk snapshot must meet the following requirements: The data disk must be formatted before the snapshot is created, that is, it has been successfully mounted on the original cloud server. Before creating the custom image for the system disk, you need to add the following commands to the `/etc/rc.local` file to write the data disk mount point into the file: `mkdir -p mount`

Wherein: `should be replaced` with the mount point of the file system, such as `/mydata`. `should be replaced` with the actual file partition location of the user, such as `/dev/vdb` (device name with a file system but no partition) and `/dev/vdb1` (device name with both a partition and a file system). Only when the above two conditions are met simultaneously can the data disk of the newly launched Linux cloud server instance be automatically recognized and mounted.

When Launching New Instances Using Custom Images and Data Disk S

Last Updated At: 2025-08-12 16:28:50

When launching a new cloud server instance, if you specify a custom image and a data disk snapshot, China Construction Bank Cloud's cloud disks can support automatic mounting after the instance is started. This means you can directly read from and write to the data disk without performing a series of operations such as adding, partitioning, or formatting. You need to perform some operations on the original instance before creating the custom image and data disk snapshot, which will be described in detail below. In the Windows system, if you want the cloud disks created from the specified data disk snapshots to be automatically mounted to new cloud server instances, the specified custom image and data disk snapshot must meet the following requirements: The SAN policy in the custom image must be set to: onlineAll. The Windows public images currently provided by China Construction Cloud have been configured with this setting by default, but it is still recommended that users check this configuration before creating a custom image. The checking method is as follows:

```
PS C:\Users\Administrator>
PS C:\Users\Administrator> diskpart

Microsoft DiskPart 版本 6.1.7601
Copyright (C) 1999-2008 Microsoft Corporation.
在计算机上: [REDACTED]

DISKPART> san

SAN 策略: 使共享磁盘脱机

DISKPART> san policy=onlineall

DiskPart 已成功更改用于当前操作系统的 SAN 策略。

DISKPART> san

SAN 策略: 全部联机
```

The data disk must be formatted as NTFS or FAT32 before creating the snapshot. Only when both of the above conditions are met can the data disk of the newly launched Windows cloud server instance be



automatically recognized and brought online.



Initialize Cloud Disks

Initializing Cloud Disks on Windows Systems

Initializing Cloud Disks on Windows Systems

Last Updated At: 2025-08-02 15:34:52

Instances can recognize connected cloud disks and treat them as regular hard disks. Users can format, partition, and create file systems for cloud block storage devices using any file system. After that, any data written to the file system will be written to the cloud disk, and this process is transparent to the applications using the device. The data disk is offline by default and cannot be used if it has not been partitioned and formatted. This tutorial will guide you through partitioning and formatting in the Windows system.

Although different Windows versions have different paths to access the "Disk Management" interface, the operations for partitioning and formatting disks after entering the Disk Management interface are basically the same.

This article will guide users through data disk mounting, partitioning, and formatting using the Windows 2012 system.



Prerequisites

Last Updated At: 2025-08-12 16:28:50

Please ensure that you have attached the cloud disk to the CVM instance and logged in to the Windows instance.

After formatting, all data in the data disk will be completely erased. Before formatting, make sure there is no data in the data disk or that important data has been backed up. To avoid service abnormalities, please ensure that the cloud server has stopped providing external services before formatting.

When a user has applied for multiple cloud disks, it is recommended to set a custom name for the elastic cloud disk that stores important data.


Users can quickly find cloud disks in the cloud disk console based on the custom name or the internal IP of the associated cloud host.

Partitioning and Formatting Disks in Windows 2012


Partitioning and Formatting Disks in Windows 2012

Last Updated At: 2025-08-12 17:46:36

Bringing a Disk Online The path to access Disk Management in Windows 2012 is: **【Start】** – **【Server Manager】** – **【Tools】** – **【Computer Management】** – **【Disk Management】** . Click the **【Start】** button:




Click **【Server Manager】** :




Click 【Tools】 – 【Computer Management】 :



Click 【Disk Management】 :




As shown in the figure below: Right-click on Disk 1 and select 【Online】 :




If there is already data on the disk (i.e., it is not a blank disk), users can ignore the following operations. Reformatting or partitioning a non-empty disk will erase all existing data.

(Optional) Format the disk 1) Right-click and select 【Initialize Disk】 :




2) Depending on the partitioning method, select **【GPT】** or **【MBR】** , and click the **【OK】** button:




When the disk size exceeds 2TB, only the GPT partition format is supported. If you are unsure whether the subsequent expansion of the disk will exceed this value, it is recommended that you choose the GPT partition; if you are certain that the disk size will not exceed this value, it is recommended that you choose the MBR partition for better compatibility. **(Optional) Partition the disk** 1) Right-click on the unallocated


space and select 【New Simple Volume】 :




2) In the pop-up window, click 【Next】 :




3) Enter the required disk size for the partition and click **【Next】** :




4) Enter the drive letter and click 【Next】 :




5) Select the file system, format the partition, and click **【Next】** :



6) Complete the creation of the new simple volume and click **【Finish】** :



7)Check the new partition:



When Launching a New Instance Using a Custom Image and Data Disk

Last Updated At: 2025-08-12 17:46:36

When launching a new cloud server instance, if the user specifies a custom image and a data disk snapshot, CCB Cloud's cloud disks can support automatic mounting after the cloud server instance is started (that is, the data disk can be directly read and written without a series of operations such as adding, partitioning, and formatting). Users need to perform some operations on the original instance before creating the custom image and data disk snapshot, which will be described in detail below. In the Windows system, if users want the cloud disk generated from the specified data disk snapshot to be automatically mounted to the new cloud server instance, the specified custom image and data disk snapshot must meet the following requirements: The SAN policy in the custom image is: onlineAll. The Windows public images currently provided by CCB Cloud have relevant settings by default, but it is still recommended that users check this configuration before creating a custom image. The checking method

```
PS C:\Users\Administrator>
PS C:\Users\Administrator> diskpart

Microsoft DiskPart 版本 6.1.7601
Copyright (C) 1999-2008 Microsoft Corporation.
在计算机上:

DISKPART> san

SAN 策略: 使共享磁盘脱机

DISKPART> san policy=onlineall

DiskPart 已成功更改用于当前操作系统的 SAN 策略。

DISKPART> san

SAN 策略: 全部联机
```

is as follows:

The data disk must have been formatted as ntfs or fat32 before the snapshot is created Only when both of the above conditions are met can the data disk of the newly launched Windows cloud server instance be automatically recognized and brought online.



Expanding Cloud Disks

Expanding Data Disks of the Cloud Disk Type

Last Updated At: 2025-08-12 17:46:36

Expanding CBS Data Disks via the Cloud Disk Console

1. Open the Cloud Virtual Machine (CVM) console.
 2. Click **【Cloud Disks】** in the navigation pane.
 3. Only disks that are in the **【Unmounted】** state and support **【Mount/Unmount】** can be expanded. Click **【More】** – **【Expand】** at the end, and select the required new size (which must be greater than or equal to the current size) to complete the expansion of the physical disk. For elastic cloud disks that have been attached to an instance, please perform the operation of unmounting the elastic cloud disk first.
- Expanding CBS Data Disks via the Cloud Server Console** Open the Cloud Virtual Machine (CVM) console. Click **【Cloud Hosts】** in the navigation pane. Only instances that are in the **【Shut down】** state and have both system disks and data disks as cloud disks can be expanded. Click **【More】** – **【Cloud Host Settings】** – **【Adjust Disks】** at the end, and select the required new size (which must be greater than or equal to the current size) to complete the expansion of the physical disk. For running instances where both the system disk and data disk are cloud disks, please shut down the instance first if you need to expand the disks.



Expanding System Disks of the Cloud Disk Type

Last Updated At: 2025-08-12 17:46:36

When the system disk is of the cloud disk type, expanding the system disk is supported, but it can only be done by reinstalling the operating system of the cloud server.



Expanding Cloud Disks

Last Updated At: 2025-08-12 17:46:36

A cloud disk is an expandable storage device in the cloud. After creating a cloud disk, users can expand its size at any time to increase storage space without losing the original data on the cloud disk. After the expansion of the cloud disk is completed, it is necessary to expand the partition and file system. You can either allocate the expanded capacity to an existing partition or format the expanded capacity into an independent new partition. If your hard disk partition is in MBR format, the MBR partition format does not support expansion beyond 2TB. It is recommended that you create a new data disk, use the GPT partition, and then copy the data to the new disk by means of replication.



Expanding the File System

Expanding the Linux File System

Last Updated At: 2025-08-12 20:13:57

A cloud disk is an expandable storage device in the cloud. After creating a cloud disk, users can expand its size at any time to increase storage space without losing the original data on the cloud disk. To achieve the goal of expanding and using the expanded space, after expanding the size of the physical cloud disk, users also need to expand the file system on it to recognize the newly added available space. You can follow the steps below: 1)Expand the size of the physical cloud disk. After performing the operation to expand the size of the physical cloud disk, users can check whether the instance has recognized the larger cloud disk space by checking the file system size. On Linux, you can use the `df -h` command to check the file system size. If you do not see the cloud disk size change to the expanded value, you need to expand this file system so that the instance can use the new space. 2) Expand the partition. Use the following command to confirm the partition table format used by the cloud disk before expansion:

<http://imgxxfb.yun.ccb.com/raw/9e3d145822f1b698d236d1e311716dd4.png>


<http://imgxxfb.yun.ccb.com/raw/371c67e403da314bba0705ea6ab5fb52.png>


Expanding Windows File Systems


Last Updated At: 2025-08-12 20:13:57



There are two scenarios for expanding Windows cloud hard disks: For the newly added storage space, create an independent new partition while keeping the old partition unchanged. Expand the old partition to include the newly added storage space, while ensuring that the data in the old partition is not lost. For both of the above scenarios, after the successful upgrade of your Windows cloud hard disk (when you see the change in the cloud hard disk capacity), you can complete the partition expansion through the Windows partition expansion tool – Partition Assistant, and ensure that the original data will not be lost.



Prerequisites Users need to download Partition Assistant first. Then complete the operation of expanding the physical cloud hard disk. If the cloud hard disk has not been formatted or had a file system created, and the capacity has been increased directly on the basis of the original empty cloud hard disk, users can directly refer to the relevant operations in the chapter "Initializing Cloud Hard Disks Under Windows Systems". Format the new space into an independent partition Open Partition Assistant, and you can see the newly expanded and unused disk space:





















2) In the pop-up box, as shown in the figure, drag the small arrow to the right to adjust the space size required for the partition, and then







Open "My Computer" to view the changes in the expanded partition (in this example, the D drive is expanded from 60G to 109G):



Unmounting an Elastic Cloud Disk


Unmounting an Elastic Cloud Disk

Last Updated At: 2025-08-12 20:13:57

When you need to mount an elastic cloud disk to another instance, you can 主动 disconnect the disk from the current instance and connect it to another. Unmounting the disk does not erase the data stored on it.

Unmounting via the Console Currently, only 普通 elastic cloud disks used as data disks can be unmounted. System disks cannot be unmounted. Before unmounting a data disk, ensure you are aware of the following:

Windows Systems: To ensure data integrity, it is recommended to 暂停 all read/write operations on the disk's file systems. Failure to do so may result in data loss for incomplete operations. Before detaching the disk, set it to an "offline" state in Disk Management. Failure to do so may prevent remounting the disk without restarting the cloud host.



- **Under the Linux operating system**, you need to log in to the instance first and perform an `umount` operation on the elastic cloud disk that needs to be unmounted. After the command is executed successfully, you can go to the console to unmount the disk. If the disk is forcibly unmounted without performing the `umount` operation first, the following problems may occur during shutdown and startup:

```
Checking filesystems
/dev/vda1: clean, 35630/524288 files, 335690/2096474 blocks
fsck.ext3: Unable to resolve 'UUID=dabe8ee8-221b-44c7-9074-4d3f8fc4ae44'
fsck.ext3: No such file or directory while trying to open /dev/disk/by-id/virtio-disk-ezy5q5l6-part5
/dev/disk/by-id/virtio-disk-ezy5q5l6-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):
```

```
Checking filesystems
/dev/vda1: clean, 35630/524288 files, 335690/2096474 blocks
fsck.ext3: Unable to resolve 'UUID=dabe8ee8-221b-44c7-9074-4d3f8fc4ae44'
fsck.ext3: No such file or directory while trying to open /dev/disk/by-id/virtio-disk-ezy5q5l6-part5
/dev/disk/by-id/virtio-disk-ezy5q5l6-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):
```

1. Log in to the CCB Cloud console.
2. Navigate to the **【Cloud Servers】** – **【Cloud Disks】** tab.
3. On the cloud disk list page:
 - For unmounting a single disk: Find the cloud disk that is in the "Mounted" state and supports mounting/unmounting, click **【More】** – **【Unmount】** after it.
 - For batch mounting: Check the cloud disks that are in the "Mounted" state and support mounting/unmounting, then click the **【Mount】** button at the top.
4. In the pop-up dialog box, confirm the warning items and click the **【Confirm】** button. **Deleting an Elastic Cloud Disk** The lifecycle of an elastic cloud disk is independent of cloud server instances and



can be destroyed independently of CVM instances.



Snapshots

What is a Snapshot

Last Updated At: 2025-08-12 20:13:57

A snapshot is a data backup method provided by CCB Cloud. It creates a fully usable copy of a specified cloud disk, making the backup independent of the cloud disk's lifecycle. A snapshot includes an image of the disk at the time the copy process started. Snapshots do not consume the user's storage space, further ensuring the reliability of backups. Snapshots are incremental backups, meaning they only store data that has changed since the last snapshot. This minimizes the time required to create a snapshot and reduces storage costs. You can create a new cloud disk based on a snapshot. The new cloud disk will initially contain all the data from the snapshot, making it an exact copy of the original cloud disk. Snapshots are region-specific, so you can only use a snapshot to create a cloud disk within the same region. For more information, see "Creating a Cloud Disk from a Snapshot."

Functions of Snapshots

- Real-time Copies of Online Data:** Snapshots are fully usable copies of cloud disks. In the event of a disk issue, you can quickly restore the disk to its pre-problem state. Taking a snapshot before major changes allows you to roll back if the changes fail.
- Persistent Backups for Critical Milestones:** Snapshots can be used to create persistent backups of business data, preserving milestone states of your business data.
- Rapid Creation of New Disks:** When a snapshot is created for Disk A, you can use this snapshot to quickly clone multiple disks, enabling rapid server deployment.

Business Scenarios for Using Snapshots

Snapshots are a convenient and efficient data protection service recommended for the following business scenarios:

- Daily Data Backup:** Regularly back up system disks and data disks to protect critical business data from accidental deletion, attacks, or viruses.
- Rapid Data Recovery:** Create one or more data snapshots before major operations such as OS upgrades, application software updates, or business data migrations. If any issues arise during the upgrade or migration, you can use the snapshots to restore your business to its normal state quickly.
- Multi-copy Applications of Production Data:** Create snapshots of production data to provide near-real-time, authentic production data for data mining, report queries, development, and testing.

Notes

- Database Operations:** For database services, perform a "Flush & Lock Table" operation before creating a snapshot.
- File System Operations:** For file systems, execute a "Sync" operation before creating a snapshot.

Creating a Snapshot

After data is written to a cloud disk, you can create regular snapshots for data backup. CCB Cloud uses an incremental approach to create snapshots, meaning only the data that has changed since the last snapshot is stored. This allows snapshots to be created quickly when the amount of changed data is small. Although snapshots are stored incrementally, deleting a snapshot will not affect your ability to use other snapshot data. Any undeleted snapshot can restore the cloud disk to the state it was in when the snapshot was taken. You can create a snapshot at any time, regardless of the cloud disk's state. However, a snapshot only captures data that has been successfully written to the cloud disk at the time of creation. Data being written by applications or other processes during snapshot



creation may not be included. For a complete snapshot, it is recommended to pause all write operations temporarily. If this is not possible, unmount the cloud disk from the instance, create the snapshot, and then remount it.

Creating a Snapshot via the Console Open the Cloud Server CVM console. Click Cloud Disks in the navigation pane. Locate the cloud disk for which you want to create a snapshot and click Create Snapshot. Wait for the snapshot creation to complete.

Rolling Back a Snapshot Rolling back a cloud disk using a snapshot restores the disk's data to the state it was in when the snapshot was created. This is particularly useful when changes result in data errors or loss. A snapshot can only be rolled back to the original cloud disk from which it was created. If you need to use snapshot data on another cloud disk, use the "Create Cloud Disk from Snapshot" feature instead. Note: When rolling back a snapshot, if the cloud disk is currently mounted on an instance, the instance must be shut down. If the cloud disk is unmounted, the rollback can be performed directly.

Rolling Back a Snapshot via the Console Open the Cloud Server CVM console. Click Snapshots in the navigation pane. Select the snapshot you want to roll back and click Roll Back.

Creating a Cloud Disk from a Snapshot A cloud disk created from a snapshot will initially contain all the data from the snapshot. Therefore, there is no need to partition, format, or create a file system on the new disk, as formatting would erase all data. After creating a cloud disk from a snapshot and attaching it to a cloud server instance, you can read and write data on the snapshot immediately. Snapshots are thus an important method for data sharing and migration.

Creating a Cloud Disk from a Snapshot via the Console Open the Cloud Server CVM console. Click Snapshots in the navigation pane. Select the data disk snapshot you want to use and click Create Disk from Snapshot. In the pop-up window, select the region/availability zone, disk capacity, and quantity, then click OK. Notes: The default capacity of the new disk is equal to the snapshot size. You can increase the capacity but cannot decrease it. You can create up to 50 elastic cloud disks at a time.

Deleting a Snapshot You can delete snapshots of data states that are no longer needed. Deleting a snapshot only removes the data unique to that snapshot and does not affect the original cloud disk. Deleting an earlier snapshot of a cloud disk also does not affect your ability to use later snapshots. Each snapshot provided by CCB Cloud can independently restore the cloud disk to the state it was in when the snapshot was taken.

Deleting a Snapshot via the Console Open the Cloud Server CVM console. Click Snapshots in the navigation pane. Locate the snapshot you want to delete and click Delete. Or, for batch deletion: Open the Cloud Server CVM console. Click Snapshots in the navigation pane. Select all the snapshots you want to delete (ensure they are not in use) and click Delete.

Scheduled Snapshots Description of Scheduled Snapshot Policies Scope: Applies to all CBS cloud disks (including system disks and data disks). Execution Policy: Scheduled snapshots can be configured to run hourly or daily. Once set, the policy remains effective until modified. Changes to the policy take effect immediately.

Automatic Expiration (Important): Scheduled snapshots can be set to automatically expire after a specified period (default is 7 days) or retained indefinitely.

Batch Operations: Create a scheduled snapshot policy and apply it to multiple cloud disks simultaneously.

Naming Convention: Automatic snapshots are named in the format snap_yyyyMMdd_HH, where yyyyMMdd is the date and HH is the hour. For example, snap_20220418_11 represents a snapshot created on April 18, 2022, at 11:00. You can modify the snapshot name manually.

Lifecycle (Important): Manually Created



定期快照策略

武汉

新建

删除

支持快照策略ID、名称搜索

☐ ID/名称	关联磁盘数量	策略详情	定期快照	创建时间	操作
☐ 策略ID: 00000000000000000000000000000000	0个	每周日，03:00，自动创建快照，...	☒	2020-08-07 14:51:28	修改策略 关联云硬盘 删除
☐ 策略ID: 00000000000000000000000000000001	0个	每周日，04:00，自动创建快照，...	☒	2020-08-07 11:34:22	修改策略 关联云硬盘 删除
☐ 策略ID: 00000000000000000000000000000002	0个	每周一、三、五，14:00，自动创...	☒	2020-08-07 11:07:42	修改策略 关联云硬盘 删除



新建快照策略



名称

所属地域

武汉

备份日期

☐ 每周日

☐ 每周一

☐ 每周二

☐ 每周三

☐ 每周四

☐ 每周五

☐ 每周六

备份时间点

☐ 00:00

☐ 01:00

☐ 02:00

☐ 03:00

☐ 04:00

☐ 05:00

☐ 06:00

☐ 07:00

☐ 08:00

☐ 09:00

☐ 10:00

☐ 11:00

☐ 12:00

☐ 13:00

☐ 14:00

☐ 15:00

☐ 16:00

☐ 17:00

☐ 18:00

☐ 19:00

☐ 20:00

☐ 21:00

☐ 22:00

☐ 23:00

快照保留时间

☒ 保留

—

7

+

天后自动删除

☐ 永久保留

首次备份时间

确定

取消

Associating Cloud Disks You can add cloud disks to a scheduled snapshot policy to enable automatic snapshot capabilities.



关联云硬盘



选择云硬盘

搜索云硬盘名称

×

Q

ID/名称	磁盘属性	容量
☐ disk-0y57fszm tesgs	系统盘	50G
☐ disk-19ewnla8 数据盘	系统盘	20G
☐ disk-qum2m3gm ssssssss	数据盘	210G
☒ disk-8hvjbv52	数据盘	40G
该磁盘未挂载到服务器上，定期快照策略将暂停对该磁盘制作快照		
☐ zhijiangliu_resize_t	数据盘	110G
☐ disk-g7kjiyw aaaaa	数据盘	170G

已选择(1)

ID/名称	磁盘属性	容量	
disk-8hvjbv52 123dsf	数据盘	40G	×

支持按住Shift进行多选

确定

取消



快照信息

名称	111
快照ID	snap-4kua54r1
地域	武汉
大小	20GB
关联磁盘	disk-3i8clzx0 (未命名)
快照类型	数据盘
保留时间	永久保留
创建时间	2022-08-23 15:01:10



☐	snap-8uic4urb snap-8uic4urb	正常	50G	数据盘	disk-backbone disk-backbone	-	2022-10-19 16:00:08	0/50	永久保留	回滚 删除 更多 ▾
☐	snap-9uiddhne snap-9uiddhne	正常	50G	数据盘	disk-backbone disk-backbone	-	2022-09-14 16:00:09	0/50	永久保留	回滚 删除 更多 ▾
☐	snap-4luc54d4 snap-4luc54d4	正常	20G	数据盘	disk-010abuo disk-010abuo	-	2022-08-23 15:01:10	0/50	永久保留	回滚 删除 更多 ▾
☐	snap-ma0u7nuu snap-ma0u7nuu	正常	50G	数据盘	disk-backbone disk-backbone	-	2022-08-20 16:00:47	0/50	永久保留	回滚 删除 更多 ▾
☐	snap-4b375u74 snap-4b375u74	正常	50G	数据盘	disk-backbone disk-backbone	-	2022-08-04 16:00:09	0/50	永久保留	回滚 删除 更多 ▾
☐	snap-5u4u4u4u snap-5u4u4u4u	正常	50G	数据盘	disk-backbone disk-backbone	-	2022-08-03 17:09:31	0/50	永久保留	回滚 删除 更多 ▾

Shared Snapshot

Operation Scenario

Shared snapshot refers to sharing a created data disk snapshot with other users. Users can share snapshots with each other through simple operations, and quickly create cloud disks using the obtained snapshots and mount them for use.

Precautions

- Each snapshot can be shared with a maximum of 50 users.
- Snapshots created from system disks do not support sharing.
- Snapshots that have been shared with other users must be unshared before they can be deleted.
You can click the target snapshot ID to enter the snapshot details page to cancel the sharing.
- Snapshots can only be shared within the same region of the other party’s account.
- It is not allowed to re-share the obtained shared snapshots with other users.

Operation Steps

1. Obtain the unique ID of the opposite account
CCB Cloud shared snapshots are identified by the unique ID of the opposite primary account. That is, you first need to obtain the unique ID of the primary account of the CCB Cloud account used by the other party. The ways to obtain it are as follows:
(1) Log in to the cloud server console.
(2) Click the account name in the upper right corner and select "Account Information".
(3) On the "Account Information" management page, view and record the "Account ID" of the

primary account.

(4) Notify the other party to send the obtained account ID to you.

2. Share the snapshot

You can share a snapshot with other users through the console:

(1) Enter the snapshot list page.

(2) Select "More > Share Snapshot" on the right side of the row where the snapshot to be shared is located. As shown in the following figure:



(3) In the pop-up "Share Snapshot" window, enter the unique ID of the counterpart's account.

(4) Click "Share" to share the snapshot.

(5) Notify the other party to go to the Snapshot List page and select the Shared Snapshots tab to view the shared snapshots.

If you need to share with multiple users, please repeat the above steps.



Monitor Cloud Disks

Last Updated At: 2025-08-12 20:13:57

Providing a good monitoring environment for cloud disks is an important part of maintaining high data reliability. Users can use cloud monitoring to monitor cloud disks that have been mounted on instances. Therefore, when needing to collect cloud disk statistical data, please first perform the operation of connecting the cloud disk to a CVM instance. Through cloud monitoring, users can view the metric data of cloud disks, as well as analyze and set alarms related to cloud disks. Currently, cloud monitoring provides the following monitoring metrics for cloud disks:

The name of the metric	The Chinese name of the indicator	Calculation:	Metric implications	unit	Statistical granularity (period)
disk_read_iops	Hard disk read IOPS	Blocks are stored as the average of read IOPS over a statistical period	The number of IOs read from block storage to memory per second	frequency	5s、10s、60s、300s
disk_read_traffic	Hard disk read traffic	Blocks are stored as the average of the read throughput over a statistical period	The rate at which data is read from block storage into memory	MB/s	5s、10s、60s、300s
disk_write_iops	Write IOPS to the hard disk	Blocks are stored as the average of write IOPS during the statistical period	The number of IOs written from memory to block storage per second	frequency	5s、10s、60s、300s



The name of the metric	The Chinese name of the indicator	Calculation:	Metric implications	unit	Statistical granularity (period)
disk_write_traffic	Hard disk write traffic	Blocks are stored as the average of the write throughput in a statistical period	The rate at which data is written from memory to block storage	MB/s	5s、10s、60s、300s
disk_await	The I/O wait time of the disk	Blocks are stored as the average of ioawait over a statistical period	What percentage of the time during the sampling period the CPU is idle and there are still outstanding I/O requests	ms	5s、10s、60s、300s
disk_svctm	The I/O service time of the disk	Blocks are stored as the average of the SVCTMs over a statistical period	IO Service Time	ms	5s、10s、60s、300s
disk_util	Disk I/O busy ratio	The block is stored as the average of io_util over a statistical period	The ratio of the time that a hard disk has IO operations (i.e., non-spatial time).	%	5s、10s、60s、300s

For specific descriptions of monitoring metrics, please refer to the Cloud Monitoring product documentation. Cloud Monitoring collects raw data of hard disks from running cloud server instances and displays the data in the form of easy-to-read icons. Statistical data is stored by default for one month, allowing users to observe the corresponding status of the disk over the past month and better understand data such as usage, read and write operations. Users can obtain data through the Cloud Monitoring console entrance or Cloud Monitoring API. Meanwhile, the console also provides visual icons for corresponding metrics. For more information, please refer to "Obtaining Monitoring Data of Specific Metrics" and "Viewing Monitoring Charts".



Best Practices

Best Practices

Last Updated At: 2025-08-12 20:13:57

How to Measure the Performance of Cloud Disks Measurement Indicators Due to the different workloads of various applications, the maximum performance of a cloud disk may not be achieved if insufficient I/O requests are provided to fully utilize it. The following indicators are generally used to measure the performance of cloud disks: IOPS: The number of read/write operations per second, measured in counts. The underlying driver type of the storage device determines the different IOPS. Throughput: The amount of data read and written per second, measured in MB/s. Latency: The time elapsed from the sending of an I/O operation to the receipt of confirmation, measured in seconds. Testing Tools FIO is a tool for testing disk performance, used for stress testing and verification of hardware. This article takes FIO as an example. When using FIO, it is recommended to conduct tests with the I/O engine of libaio. Please install FIO and libaio by yourself. Warning: Do not perform FIO tests on the system disk to avoid damaging important system files. To prevent data corruption caused by damage to the underlying file system metadata, do not conduct tests on business data disks. Ensure that there is no mount configuration of the tested hard disk in the /etc/fstab file, otherwise it will cause the cloud server to fail to start.

Recommendations for Test Objects It is recommended to perform FIO tests on idle hard disks that do not store important data, and re-create the file system of the tested hard disk after the test. When testing hard disk performance, it is advisable to directly test the raw data disk (such as /dev/vdb). When testing file system performance, it is recommended to specify a specific file for testing (such as /data/file). Test Examples The testing formulas for different scenarios are basically the same, with differences only in three parameters: rw, iodepth, and bs (block size). For example, the optimal iodepth suitable for each workload varies, depending on the sensitivity of your specific application to IOPS and latency. Common use cases are as follows: – bs = 4k, iodepth = 1: Random read/write test, which can reflect the latency performance of the hard disk. Execute the following command to test the random read latency of the hard disk: `fio -bs=4k -ioengine=libaio -iodepth=1 -direct=1 -rw=randread -time_based -runtime=600 -refill_buffers -norandommap -randrepeat=0 -group_reporting -name=fio-randread-lat --size=10G -filename=/dev/vdb` Execute the following command to test the random write latency of the hard disk: `fio -bs=4k -ioengine=libaio -iodepth=1 -direct=1 -rw=randwrite -time_based -runtime=600 -refill_buffers -norandommap -randrepeat=0 -group_reporting -name=fio-randwrite-lat --size=10G -filename=/dev/vdb` Execute the following command to test the random mixed read/write latency performance of SSD cloud disks: `fio --bs=4k --ioengine=libaio --iodepth=1 --direct=1 --rw=randrw --time_based --runtime=100 --refill_buffers --norandommap --randrepeat=0 --group_reporting --`

name=fio-read --size=1G --filename=/dev/vdb The test results are shown in the following figure:

```
fio-read: (g=0): rw=randrw, bs=(R) 4096B-4096B, (W) 4096B-4096B, (T) 4096B-4096B, ioengine=libaio, iodepth=1
fio-3.1
Starting 1 process
Jobs: 1 (f=1): [m(1)][100.0%][r=3411KiB/s,w=3603KiB/s][r=852,w=900 IOPS][eta 00m:00s]
fio-read: (groupid=0, jobs=1): err= 0: pid=2377: Thu Jun 13 18:23:47 2019
  read: IOPS=880, BW=3523KiB/s (3607kB/s) (344MiB/100001msec)
    slat (nsec): min=2905, max=62479, avg=5254.61, stdev=2075.46
    clat (usec): min=205, max=6921, avg=463.65, stdev=259.48
      lat (usec): min=209, max=6925, avg=469.13, stdev=259.56
    clat percentiles (usec):
      | 1.00th=[ 245], 5.00th=[ 269], 10.00th=[ 293], 20.00th=[ 375],
      | 30.00th=[ 400], 40.00th=[ 416], 50.00th=[ 437], 60.00th=[ 457],
      | 70.00th=[ 478], 80.00th=[ 498], 90.00th=[ 545], 95.00th=[ 619],
      | 99.00th=[ 2057], 99.50th=[ 2376], 99.90th=[ 3294], 99.95th=[ 4015],
      | 99.99th=[ 6259]
    bw ( KiB/s): min= 2168, max= 4024, per=100.00%, avg=3522.64, stdev=310.95, samples=200
    iops        : min=  542, max= 1006, avg=880.66, stdev=77.74, samples=200
  write: IOPS=877, BW=3511KiB/s (3595kB/s) (343MiB/100001msec)
    slat (nsec): min=2981, max=58808, avg=5377.71, stdev=2079.36
    clat (usec): min=421, max=10492, avg=659.10, stdev=219.96
      lat (usec): min=428, max=10496, avg=664.70, stdev=220.05
    clat percentiles (usec):
      | 1.00th=[ 490], 5.00th=[ 523], 10.00th=[ 545], 20.00th=[ 562],
      | 30.00th=[ 578], 40.00th=[ 594], 50.00th=[ 611], 60.00th=[ 635],
      | 70.00th=[ 660], 80.00th=[ 693], 90.00th=[ 783], 95.00th=[ 914],
      | 99.00th=[ 1516], 99.50th=[ 1926], 99.90th=[ 3261], 99.95th=[ 3982],
      | 99.99th=[ 5342]
    bw ( KiB/s): min= 2296, max= 4008, per=100.00%, avg=3510.84, stdev=305.46, samples=200
    iops        : min=  574, max= 1002, avg=877.71, stdev=76.36, samples=200
  lat (usec)    : 250=0.76%, 500=40.51%, 750=51.04%, 1000=5.03%
  lat (msec)    : 2=1.90%, 4=0.71%, 10=0.05%, 20=0.01%
  cpu           : usr=0.50%, sys=1.52%, ctx=175841, majf=0, minf=29
  IO depths     : 1=100.0%, 2=0.0%, 4=0.0%, 8=0.0%, 16=0.0%, 32=0.0%, >=64=0.0%
    submit      : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
    complete    : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
```

- bs = 128k, iodepth = 32: Sequential read/write test, which can reflect the throughput performance of the hard disk. Execute the following command to test the sequential read throughput bandwidth of the hard disk: `fio -bs=128k -ioengine=libaio -iodepth=32 -direct=1 -rw=read -time_based -runtime=600 -refill_buffers -norandommap -randrepeat=0 -group_reporting -name=fio-read-throughput --size=10G -filename=/dev/vdb` Execute the following command to test the sequential write throughput bandwidth of the hard disk: `fio -bs=128k -ioengine=libaio -iodepth=32 -direct=1 -rw=write -time_based -runtime=600 -refill_buffers -norandommap -randrepeat=0 -group_reporting -name=fio-write-throughput --size=10G -filename=/dev/vdb` Execute the following command to test the sequential read throughput performance of SSD cloud disks: `fio --bs=128k --ioengine=libaio --iodepth=32 --direct=1 --rw=write --time_based --runtime=100 --refill_buffers --norandommap --randrepeat=0 --group_reporting -rw --size=1G --filename=/dev/vdb` The test results are shown in the following figure: **Note:** There was a typo in the original text (iodpth instead of iodepth and filname instead of filename), which has been

corrected in the translation.

```
fio-rw: (g=0): rw=write, bs=(R) 128KiB-128KiB, (W) 128KiB-128KiB, (T) 128KiB-128KiB, ioengine=libaio, iodepth=32
fio-3.1
Starting 1 process
Jobs: 1 (f=1): [W(1)][100.0%][r=0KiB/s,w=260MiB/s][r=0,w=2082 IOPS][eta 00m:00s]
fio-rw: (groupid=0, jobs=1): err= 0: pid=2679: Thu Jun 13 18:27:32 2019
  write: IOPS=2081, BW=260MiB/s (273MB/s) (25.4GiB/100045msec)
    slat (nsec): min=2847, max=72524, avg=7739.21, stdev=3233.07
    clat (usec): min=1033, max=250494, avg=15341.09, stdev=28854.07
    lat (usec): min=1041, max=250503, avg=15349.03, stdev=28853.95
    clat percentiles (usec):
      | 1.00th=[ 1565],  5.00th=[ 1860], 10.00th=[ 2057], 20.00th=[ 2311],
      | 30.00th=[ 2540], 40.00th=[ 2769], 50.00th=[ 2999], 60.00th=[ 3326],
      | 70.00th=[ 3818], 80.00th=[ 5014], 90.00th=[82314], 95.00th=[84411],
      | 99.00th=[86508], 99.50th=[86508], 99.90th=[87557], 99.95th=[88605],
      | 99.99th=[90702]
    bw ( KiB/s): min=265708, max=319488, per=100.00%, avg=266498.36, stdev=3766.74, samples=200
    iops        : min= 2075, max= 2496, avg=2082.01, stdev=29.43, samples=200
    lat (msec)  : 2=8.46%, 4=64.09%, 10=11.90%, 20=0.18%, 50=0.01%
    lat (msec)  : 100=15.37%, 500=0.01%
    cpu         : usr=5.34%, sys=1.90%, ctx=63555, majf=0, minf=28
    IO depths   : 1=0.1%, 2=0.1%, 4=0.1%, 8=0.1%, 16=0.1%, 32=100.0%, >=64=0.0%
      submit    : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
      complete  : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.1%, 64=0.0%, >=64=0.0%
      issued rw: total=0,208238,0, short=0,0,0, dropped=0,0,0
      latency   : target=0, window=0, percentile=100.00%, depth=32

Run status group 0 (all jobs):
  WRITE: bw=260MiB/s (273MB/s), 260MiB/s-260MiB/s (273MB/s-273MB/s), io=25.4GiB (27.3GB), run=100045-100045msec

Disk stats (read/write):
  vdb: ios=42/207998, merge=0/0, ticks=21/3173469, in queue=3174831, util=99.95%
```

- bs = 4k, iodepth = 32: Random read/write test, which can reflect the IOPS performance of the hard disk. Execute the following command to test the random read IOPS of the hard disk: `fio -bs=4k -ioengine=libaio -iodepth=32 -direct=1 -rw=randread -time_based -runtime=600 -refill_buffers -norandommap -randrepeat=0 -group_reporting -name=fio-randread-iops --size=10G -filename=/dev/vdb` Execute the following command to test the random write IOPS of the hard disk: `fio -bs=4k -ioengine=libaio -iodepth=32 -direct=1 -rw=randwrite -time_based -runtime=600 -refill_buffers -norandommap -randrepeat=0 -group_reporting -name=fio-randwrite-iops --size=10G -filename=/dev/vdb` Test the random read IOPS performance of the SSD cloud disk. As

shown in the following figure:

```
[root@VM_16_21_centos ~]# fio -bs=4k -ioengine=libaio -iodepth=32 -direct=1 -rw=randread -time_based
-runtime=300 -refill_buffers -norandommap -randrepeat=0 -group_reporting -name=fio-randread --siz
e=100G -filename=/dev/vdc
fio-randread: (g=0): rw=randread, bs=(R) 4096B-4096B, (W) 4096B-4096B, (T) 4096B-4096B, ioengine=lib
aio, iodepth=32
fio-3.1
Starting 1 process
Jobs: 1 (f=1): [r(1)][100.0%][r=18.8MiB/s,w=0KiB/s][r=4804,w=0 IOPS][eta 00m:00s]
fio-randread: (groupid=0, jobs=1): err= 0: pid=2689: Tue Jul 16 13:39:33 2019
  read: IOPS=4800, BW=18.8MiB/s (19.7MB/s) (5626MiB/300017msec)
    slat (usec): min=2, max=3183, avg= 7.55, stdev=14.34
    clat (usec): min=209, max=777668, avg=6656.04, stdev=45385.49
    lat (usec): min=371, max=777673, avg=6664.08, stdev=45385.46
    clat percentiles (usec):
      | 1.00th=[ 635], 5.00th=[ 832], 10.00th=[ 938], 20.00th=[ 1090],
      | 30.00th=[ 1237], 40.00th=[ 1352], 50.00th=[ 1467], 60.00th=[ 1582],
      | 70.00th=[ 1713], 80.00th=[ 1991], 90.00th=[ 9372], 95.00th=[ 19006],
      | 99.00th=[ 38536], 99.50th=[341836], 99.90th=[734004], 99.95th=[750781],
      | 99.99th=[767558]
    bw ( KiB/s): min= 256, max=38424, per=100.00%, avg=19202.16, stdev=13626.10, samples=600
    iops        : min= 64, max= 9606, avg=4800.52, stdev=3406.52, samples=600
    lat (usec)   : 250=0.01%, 500=0.12%, 750=2.64%, 1000=10.86%
    lat (msec)   : 2=66.45%, 4=5.87%, 10=4.59%, 20=6.49%, 50=2.32%
    lat (msec)   : 100=0.01%, 500=0.36%, 750=0.26%, 1000=0.05%
    cpu          : usr=1.38%, sys=5.00%, ctx=233328, majf=0, minf=63
    IO depths    : 1=0.1%, 2=0.1%, 4=0.1%, 8=0.1%, 16=0.1%, 32=100.0%, >=64=0.0%
    submit      : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
    complete    : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.1%, 64=0.0%, >=64=0.0%
    issued rwts: total=1440148,0,0, short=0,0,0, dropped=0,0,0
    latency     : target=0, window=0, percentile=100.00%, depth=32

Run status group 0 (all jobs):
  READ: bw=18.8MiB/s (19.7MB/s), 18.8MiB/s-18.8MiB/s (19.7MB/s-19.7MB/s), io=5626MiB (5899MB), run=
300017-300017msec

Disk stats (read/write):
  vdc: ios=1440052/0, merge=0/0, ticks=9478008/0, in_queue=9485217, util=99.98%
```

Building a RAID Array with Multiple Elastic Cloud Disks

1. RAID Overview

RAID (Redundant Array of Independent Disks) can combine multiple disks into a disk array, improving data read/write performance and reliability. The operating system will recognize the disk array as a single hard drive. Currently, there are multiple RAID levels. Depending on the selected level, compared to a single large hard drive with equivalent capacity, the disk array offers advantages such as enhanced data integration, improved fault tolerance, increased throughput, or expanded capacity.

Notes:

- Please renew elastic cloud disks that are about to expire in a timely manner to avoid impacts on the RAID array due to forced isolation of expired elastic cloud disks by the system.
- It is recommended to use partitions of the same size when creating RAID 1, RAID 01, or RAID 10 to reduce disk space waste.

Building a RAID Note: This article takes the example of building a RAID 0 using four elastic cloud



disks in a CentOS cloud server. Operations may vary for different operating systems and RAID levels. This article is for reference only. For specific operation steps and differences, please refer to the product documentation of the corresponding operating system or relevant RAID documents. The Linux kernel provides an md module for managing RAID devices, and the mdadm tool can be directly used to call the md module to create RAID 0.

```
[root@VM_63_126_centos ~]# fdisk -l | grep /dev/vd | grep Linux | grep -v vda
/dev/vdc1          1      20805      10485688+  83  Linux
/dev/vdc2          20806      27046      3145464    83  Linux
/dev/vdd1          1      20805      10485688+  83  Linux
/dev/vde1          1      20805      10485688+  83  Linux
/dev/vdf1          1      20805      10485688+  83  Linux
```

(1) Log in to the cloud server as the root user. (2) Execute the following command to install mdadm: `yum install mdadm -y` The result of a successful installation is shown in the figure below:

```
[root@VM_63_126_centos ~]# yum install mdadm -y
Loaded plugins: fastestmirror, security
Setting up Install Process
Loading mirror speeds from cached hostfile
Resolving Dependencies
--> Running transaction check
--> Package mdadm.x86_64 0:3.3.2-5.el6 will be installed
--> Finished Dependency Resolution

Dependencies Resolved

=====
Package                Arch                Version              Repository            Size
=====
Installing:
mdadm                  x86_64              3.3.2-5.el6          os                    345 k
Transaction Summary
=====
Install      1 Package(s)

Total download size: 345 k
Installed size: 800 k
Downloading Packages:
mdadm-3.3.2-5.el6.x86_64.rpm                                | 345 kB    00:00
Running rpm_check_debug
Running Transaction Test
Transaction Test Succeeded
Running Transaction
  Installing : mdadm-3.3.2-5.el6.x86_64                    1/1
  Verifying  : mdadm-3.3.2-5.el6.x86_64                    1/1

Installed:
mdadm.x86_64 0:3.3.2-5.el6

Complete!
```

(3) Execute the following command to create a RAID 0 using mdadm: `mdadm --create /dev/md0 --level=0 --raid-devices=4 /dev/vd[cdef]1` The result of a successful creation is shown in the figure below:

```
[root@VM_63_126_centos ~]# mdadm --create /dev/md0 --level=0 --raid-devices=4 /dev/vd[cdef]1
mdadm: Defaulting to version 1.2 metadata
mdadm: array /dev/md0 started.
```

(4) Execute the following command to create a file system using mkfs (taking the EXT3 file system as an example): `mkfs.ext3 /dev/md0`

The result of a successful creation is shown in the figure below:

```
[root@VM_63_126_centos ~]# mkfs.ext3 /dev/md0
mke2fs 1.41.12 (17-May-2010)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=128 blocks, Stripe width=512 blocks
2621440 inodes, 10477056 blocks
523852 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=4294967296
320 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000, 7962624

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

This filesystem will be automatically checked every 24 mounts or
180 days, whichever comes first.  Use tune2fs -c or -i to override.
```

(5) Execute the following commands in sequence to mount the file system. `mkdir md0/` `mount /dev/md0 md0/` `tree md0` The result of a successful mount is shown in the figure below:

```
[root@VM_63_126_centos ~]# mount /dev/md0 md0/
[root@VM_63_126_centos ~]# tree md0
md0
|-- lost+found

1 directory, 0 files
```

(6) Execute the following command to view the details of the file system and record its UUID (take `c5d8a204:c28853ba:3882e9f8:62d078de` as an example). As shown in the figure below: `mdadm --detail -`



–scan

```
[root@VM_163_144_centos ~]# mdadm --detail --scan  
ARRAY /dev/md0 metadata=1.2 name=VM_163_144_centos:0 UUID=c5d8a204:c28853ba:3882e9f8:62d078de
```

(7) Execute the following command to edit the mdadm configuration file: vi /etc/mdadm.conf

(8) In edit mode, enter the relevant content. For elastic cloud disks, it is recommended to enter the following configuration: DEVICE /dev/disk/by-id/virtio-弹性云盘1ID-part1 DEVICE /dev/disk/by-id/virtio-弹性云盘2ID-part1 DEVICE /dev/disk/by-id/virtio-弹性云盘3ID-part1 DEVICE /dev/disk/by-id/virtio-弹性云盘4ID-part1 ARRAY logical device path metadata= UUID=

In this example, the logical device path is /dev/md0 and the metadata is 1.2, so enter: DEVICE /dev/disk/by-id/virtio-弹性云盘1ID-part1 DEVICE /dev/disk/by-id/virtio-弹性云盘2ID-part1 DEVICE /dev/disk/by-id/virtio-弹性云盘3ID-part1 DEVICE /dev/disk/by-id/virtio-弹性云盘4ID-part1 ARRAY /dev/md0 metadata=1.2 UUID=3c2adec2:14cf1fa7:999c29c5:7d739349

(9) Press Esc, then enter :wq to save and exit. **Building an LVM Logical Volume with Multiple Elastic Cloud Disks**

1. Introduction to LVM Logical Volume Manager (LVM) establishes a logical layer on top of hard disks and partitions. It divides disks or partitions into Physical Extents (PE) of the same size. Different disks or partitions can be assigned to the same Volume Group (VG). On a VG, you can create Logical Volumes (LV), and on an LV, you can create a file system. Compared with directly using disk partitions, LVM has the advantage of flexibly adjusting the capacity of the file system: The file system is no longer limited by the size of a physical disk and can be distributed across multiple disks. For example, you can apply for 3 elastic cloud disks of 4TB each and use LVM to create a large file system of nearly 12TB. The size of the logical volume can be adjusted dynamically without re-partitioning the disk. When the space of the LVM volume group cannot meet your needs, you can apply for an elastic cloud disk separately, mount it to the corresponding cloud server, and then add it to the LVM volume group for expansion.
2. Building LVM Note: This article takes the example of creating a dynamically resizable file system using LVM with 3 elastic cloud disks. As shown in the figure below:

```
[root@VM_63_126_centos ~]# fdisk -l | grep vd | grep -v vda | grep -v vdb  
Disk /dev/vdc: 10.7 GB, 10737418240 bytes  
Disk /dev/vdd: 10.7 GB, 10737418240 bytes  
Disk /dev/vde: 10.7 GB, 10737418240 bytes
```

Step 1 Create Physical Volumes (PV) (1) Log in to the Linux cloud server as the root user. (2) Execute the following command to create Physical Volumes (PV). pvcreate <disk path 1> ... In this article, taking



/dev/vdc, /dev/vdd, and /dev/vde as examples, execute: `pvcreeate /dev/vdc /dev/vdd /dev/vde` If the creation is successful, it will be as shown in the figure below:

```
[root@VM_63_126_centos ~]# pvcreeate /dev/vdc /dev/vdd /dev/vde
Physical volume "/dev/vdc" successfully created
Physical volume "/dev/vdd" successfully created
Physical volume "/dev/vde" successfully created
```

```
[root@VM_63_126_centos ~]# lvm diskscan | grep LVM
/dev/vdc [ 10.00 GiB] LVM physical volume
/dev/vdd [ 10.00 GiB] LVM physical volume
/dev/vde [ 10.00 GiB] LVM physical volume
3 LVM physical volume whole disks
0 LVM physical volumes
```

```
[root@VM_63_126_centos ~]# vgcreate lvm_demo0 /dev/vdc /dev/vdd
Volume group "lvm_demo0" successfully created
```

```
[root@VM_63_126_centos ~]# vgextend lvm_demo0 /dev/vdf
Volume group "lvm_demo0" successfully extended
```

```
[root@VM_63_126_centos ~]# vgs
VG      #PV #LV #SN Attr   VSize  VFree
lvm_demo 3   0   0 wz--n- 29.99g 29.99g
```

Create a Logical Volume (LV) Execute the following command to create an LV: `lvcreate -L` In this example, to create an 8GB logical volume named "lv_0", execute: `lvcreate -L 8G -n lv_0 lvm_demo0` If the creation is successful, it will be as shown in the figure below:

```
[root@VM_63_126_centos ~]# lvcreate -L 8G -n lv_0 lvm_demo
Logical volume "lv_0" created
```

```
[root@VM_63_126_centos ~]# pvs
PV      VG      Fmt Attr PSize PFree
/dev/vdc lvm_demo lvm2 a-- 10.00g 2.00g
/dev/vdd lvm_demo lvm2 a-- 10.00g 10.00g
/dev/vde lvm_demo lvm2 a-- 10.00g 10.00g
```



```
mkfs.ext3 /dev/lvm_demo0/lv_0 mount /dev/lvm_demo0/lv_0 vg0/
```

```
[root@VM_63_126_centos ~]# mount /dev/lvm_demo/lv_0 vg0/
[root@VM_63_126_centos ~]# mount | grep lvm
/dev/mapper/lvm_demo-lv_0 on /root/vg0 type ext3 (rw)
```

```
lvextend [-L +/- ] lvextend -L +4G /dev/lvm_demo0/lv_0
```

```
[root@VM_63_126_centos vg0]# lvextend -L +4G /dev/lvm_demo/lv_0
Size of logical volume lvm_demo/lv_0 changed from 8.00 GiB (2048 extents) to 12.00 GiB (3072 extents).
Logical volume lv_0 successfully resized
```

```
[root@VM_63_126_centos vg0]# pvs
PV          VG          Fmt Attr PSize  PFree
/dev/vdc    lvm_demo  lvm2 a--  10.00g    0
/dev/vdd    lvm_demo  lvm2 a--  10.00g  7.99g
/dev/vde    lvm_demo  lvm2 a--  10.00g 10.00g
```

```
resize2fs /dev/lvm_demo0/lv_0
```

```
[root@VM_63_126_centos vg0]# resize2fs /dev/lvm_demo/lv_0
resize2fs 1.41.12 (17-May-2010)
Filesystem at /dev/lvm_demo/lv_0 is mounted on /root/vg0; on-line resizing required
old desc_blocks = 1, new_desc_blocks = 1
Performing an on-line resize of /dev/lvm_demo/lv_0 to 3145728 (4k) blocks.
The filesystem on /dev/lvm_demo/lv_0 is now 3145728 blocks long.

[root@VM_63_126_centos vg0]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/vda1       7.9G 1019M  6.5G  14% /
/dev/mapper/lvm_demo-lv_0
                12G  549M   11G   5% /root/vg0
```

```
df -h
```



Frequently Asked Questions

Last Updated At: 2025-08-12 20:13:57

General Questions Q: What happens to the data when a CVM is terminated? A: Since the lifecycle of a local disk is completely tied to the CVM instance, the data stored in the local disk will be destroyed along with the termination of the CVM instance. Data stored on Cloud Block Storage (CBS) is independent of the instance's lifecycle, and you can freely choose whether to retain the CBS volume beyond the instance's lifecycle. Therefore, we recommend using CBS to store data that needs to be kept for a long time. Q: Which CBS type should I choose? A: Before choosing a disk type, please first confirm the type of work and data storage you need it for: If your disk is used to store data for regular scenarios such as small and medium-sized databases, Web/App applications, transactional workloads, etc., it is recommended to use high-performance cloud disks to get the best cost-effectiveness. If your disk is used to store data for scenarios with high workloads and high performance requirements, such as large core databases, OLTP services, NoSQL databases, etc., it is recommended to use SSD cloud disks to get the highest performance. Q: After successfully applying for a cloud disk, can I change its type? A: The current version supports upgrading the elastic cloud disk type from high-performance cloud disk to SSD cloud disk. You can back up the data with a snapshot, and then upgrade the high-performance cloud disk to SSD cloud disk through [Adjust Cloud Disk Type]. Q: After successfully applying for a cloud disk, can I adjust its capacity? A: Cloud disks support capacity adjustment, which can be expanded but not reduced. Cloud Disk Usage and Lifecycle Questions Q: Does it support multiple instances accessing the same cloud disk? A: Although you can mount up to 20 cloud disks to one instance, currently multiple instances are not supported to access the same cloud disk. Q: Can multiple disks be logically merged into one to get better performance? A: Yes. You can stripe multiple cloud disks mounted to a CVM instance to achieve higher performance than a single disk. For details, please refer to "Building LVM Logical Volumes with Multiple Elastic Cloud Disks". Q: Will data be lost when unmounting a cloud disk? A: The data in the cloud disk will not change due to mounting or unmounting. To maintain data consistency, we strongly recommend: On Windows operating systems, to ensure data integrity, it is recommended to suspend all read and write operations to the file systems of the disk; otherwise, uncompleted read and write data will be lost. On Linux operating systems, you need to log in to the cloud server instance and perform the umount command operation on the disk. After the command is executed successfully, enter the console to unmount the disk. Q: Will the cloud disk lose data when unmounting it? A: We strongly recommend: On Windows operating systems, to ensure data integrity, it is recommended to suspend all read and write operations to the file systems of the disk; otherwise, uncompleted read and write data will be lost. On Linux operating systems, you need to log in to the cloud server instance and perform the umount command operation on the disk. After the command is executed successfully, enter the console to unmount the disk. Q: Why can't I find the cloud server I want to mount when mounting a cloud disk? A: Please ensure that your cloud server instance and the disk are in the same availability zone of the same



region. Cloud disks can only be mounted within the same availability zone of the same region. Also, ensure that your instance has not been released. Q: Several cloud disks of the same size and type are mounted on the same sub-machine. How to distinguish the correspondence between these disks and the actually applied CBS? A: For Linux operating systems, you can view the correspondence between elastic cloud disks and device names by executing the command `ls -l /dev/disk/by-id/`:

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

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

Q: What are the precautions for using cloud disks? A:

- For independently applied cloud disks, when configuring static file system information using `fstab`, the file system identifier should use the UUID or label of the file system. This prevents changes in the kernel name of the cloud disk on the instance due to multiple mount/unmount operations of multiple cloud disks on the same instance.
- Please pay special attention that if a cloud disk expires before the cloud server, it will be throttled, unmounted, and reclaimed within a period after expiration. To avoid impact on your business, please renew it in a timely manner.
- If removing the cloud disk from the cloud server will not have a serious impact on your critical business, you can consider using the `nofail` option when configuring `fstab`. This prevents system restart errors caused by the removal of the cloud disk from the cloud server.
- It is recommended that you execute the `san policy=OnlineAll` operation in `diskpart` before using the cloud disk on the Windows operating system.
- When unmounting a cloud disk from the Windows operating system, it is recommended that you first interrupt all read and write operations to the disk and perform an offline operation.

Common Questions about Snapshots Q: Do I need to unmount the disk or interrupt all read/write operations to take a snapshot? A: No. You can take snapshots in real-time while the disk is connected and in use, without affecting your normal business. However, snapshots can only capture data that has been written to the CBS volume, not data cached in the memory by applications or the operating system.



To ensure that all application data is captured in the snapshot, we recommend that you completely pause IO operations on the disk before taking the snapshot. For CBS volumes used as system disks, we recommend shutting down the instance first to take a complete snapshot.

Q: Can I read previous snapshots for recovery? A: Each snapshot has a unique global identifier. Users can roll back or create a new CBS volume based on any existing snapshot.

Q: Is it necessary to shut down the instance to roll back a snapshot? A: For disks already mounted on the instance, the instance needs to be shut down during rollback; for unmounted disks, the rollback operation can be performed directly.

Q: Will taking a snapshot affect disk performance? A: Taking a snapshot will have a slight impact on the read/write performance of the disk depending on the amount of data written to the disk. Therefore, we recommend taking snapshots during off-peak business hours.

Q: How long does it take for a snapshot to be available after it is created? A: The time required to create a snapshot varies depending on factors such as the amount of data written to the disk and the underlying read/write conditions, making it difficult to predict. However, you don't need to worry, as taking a snapshot does not affect your normal use of the disk.

Q: Why can some snapshots not be used to create images? A: Both system disks and data disks can be used to create snapshots, but only system disk snapshots can be used to create custom images, which can then be used to create one or more instances with the same environment.

Q: When a cloud disk is destroyed, will its associated snapshots be deleted? A: No. Destroying a cloud disk only makes data disk snapshots unable to be rolled back, but the associated snapshots will not be deleted. If you need to delete related snapshots, please do so through the console or API.

Common Questions about Disk Performance Q: What are random IO and sequential IO? A: IO is either random or sequential. Random IO refers to access to non-contiguous addresses, which are randomly distributed in the addressing space of the disk LUN. Businesses that mainly generate random IO include: OLTP services, SQL, instant messaging services, etc. Sequential IO refers to read/write operations that access data continuously from adjacent addresses, block by block logically. In sequential IO access, the disk seek time is significantly reduced because the read/write head basically does not need to move to access the next block. Businesses such as data backup and log writing mostly generate sequential IO.

Q: What is IOPS? A: IOPS (Input/Output Per Second) refers to the amount of input/output (or the number of read/write operations) per second, and is one of the important indicators to measure disk performance. IOPS refers to the number of I/O requests that the system can process per unit time, generally measured by the number of I/O requests processed per second. I/O requests usually include read or write data operations. Traditional disks are essentially mechanical devices, such as FC, SAS, and SATA disks, with rotation speeds usually ranging from 5400/7200/10000/15000 rpm. The key factor affecting disks is the disk service time, which is the time it takes for the disk to complete an I/O request. It consists of three



parts: seek time, rotation delay, and data transmission time. Typically, a 7200-rpm mechanical hard disk can provide 75–150 IOPS, and a 15000-rpm mechanical hard disk can provide 175–210 IOPS (the specific value depends on access patterns such as sequential/random access and IO size).

Q: Does the IO size of application read/write operations affect the IOPS performance I get? A: Yes. For given resources, the IOPS you get depends on the IO size of the application's read and write operations. Generally, when the IO size is small (such as 256K), the IOPS performance of the disk can be fully utilized.

Q: Does the IO size of application read/write operations affect the throughput performance I get? A: Yes. For given resources, the throughput you get depends on the IO size of the application's read and write operations. Generally, when the IO size is large (such as 1M), the throughput performance of the disk can be fully utilized. Glossary

Elastic Cloud Disk

A cloud disk with an independent lifecycle that can be freely mounted and unmounted between different cloud server instances in the same region and same availability zone (currently, simultaneous mounting on multiple cloud server instances is not supported), also known as an elastic cloud drive.

Non-elastic Cloud Disk

A non-elastic cloud disk refers to a cloud disk created together with a cloud server, whose lifecycle is tied to the cloud server and does not support elastic mounting.

Rollback

The act of restoring a program or data to the last correct state when an error occurs in program or data processing. Rollback includes types such as program rollback and data rollback.

IOPS

IOPS (Input/Output Per Second) refers to the amount of input/output (or the number of read/write operations) per second, and is one of the important indicators for measuring hard disk performance. IOPS refers to the number of I/O requests that a system can process within a unit of time, usually measured by the number of I/O requests processed per second. I/O requests are typically read or write data operation requests. Traditional disks are essentially mechanical devices, such as FC, SAS, and SATA disks, with rotation speeds usually ranging from 5400/7200/10000/15000 rpm. A key factor affecting disks is the disk service time, which is the time it takes for the disk to complete an I/O request. It consists of three parts: seek time, rotational delay, and data transmission time. Generally, a single 7200rpm mechanical hard disk can provide 75 – 150 IOPS, and a single 15000rpm mechanical hard disk can provide 175 – 210 IOPS. The specific values depend on factors such as access patterns (e.g., sequential, random) and I/O size.

Full Snapshot

The first snapshot created for a disk, which saves the full amount of data on the disk.



Throughput

The amount of data successfully transmitted per unit time for a network, device, port, virtual circuit, or other facilities. Cloud Disk Snapshot

A copy of a cloud disk at a certain point in time, which can be used to restore the cloud disk to the state at the time the snapshot was created. Incremental Snapshot

After a full snapshot or the previous incremental snapshot, each subsequent snapshot only needs to back up the files that have been added or modified compared to the previous one, which is called an incremental snapshot.