Cloud Block Storage

Product Documentation

Contents

Cloud Block Storage 4
- Product Introduction 4
  - CBS Overview 4
  - Strengths 6
  - Use Cases 7
  - Cloud Disk Instances 10
    - Cloud Disk Types 10
    - Cloud Disk Status 12
    - Cloud Disk Encryption 13
    - Enhanced SSD Performance Descriptions 17
  - Cloud Disk Snapshots 19
    - Snapshot Overview 19
    - Snapshot Principle 22
  - CAM 27
    - CAM Overview 27
    - Authorizable Resource Types 29
    - Authorization Policy Syntax 31
  - Related Products 34
  - Previous Generation Products 35
- Purchase Guide 36
  - Billing Overview 36
  - Purchase Guide 38
  - Viewing Bills 39
- Getting Started 40
  - Step 1: Creating Cloud Disks 40
  - Step 2: Attaching Cloud Disks 41
  - Step 3: Initializing Cloud Disks 42
- Operation Guide 45
  - Operation Overview 45
  - Use Limits 47
  - Managing Cloud Disks 49
    - Creating Cloud Disks 49
    - Attaching Cloud Disks 51
    - Initializing Cloud Disks 56
      - Initialization Scenarios 56
      - Initializing Cloud Disks (< 2 TB) 58
      - Initializing Cloud Disks (≥ 2 TB) 71
    - Expanding Cloud Disks 84
      - Expansion Scenarios 84
      - Expanding Cloud Disks 86
      - Expanding System Disk Partitions and File Systems Online 91
      - Expanding Partitions and File Systems (Windows) 95
      - Expanding the Partition and the File System (Linux) 97
    - Adjusting the Cloud Disk Type 116
    - Detaching Cloud Disks 118
    - Terminating Cloud Disks 120
    - Adjusting Cloud Disk Performance 122
  - Managing Snapshots 124
    - Creating Snapshots 124
    - Rolling Back Data from Snapshots 126
    - Creating Cloud Disks from Snapshots 127
    - Scheduled Snapshot 129
    - Cross-Region Snapshot Replication 134
    - Deleting Snapshots 135
    - Sharing Snapshots 136
  - Managing Snapshot Groups 138
    - Creating Snapshot Groups 138
    - Rolling Back Data from Snapshot Groups 140
    - Creating Cloud Disks from Snapshot Groups 141
    - Creating Custom Images from Snapshot Groups 143
    - Deleting Snapshot Groups 144
  - Monitoring and Alarming 145
  - Tag 147
    - Using Tags to Manage Cloud Disks 147
    - Editing Tags 150
- Best Practices 151
  - How to Measure the Cloud Disk Performance 151
  - Building LVM Logic Volumes with Multiple Elastic Cloud Disks 157
- Troubleshooting 162
  - Disk Not Mounting upon Linux CVM Restart 162
- FAQs 167
  - Performance-related Issues 167
  - Snapshot-related Issues 168
  - Usage-related Issues 171
- Glossary 176
- API Documentation 178
  - CBS(cbs) 178
    - version(2017-03-12) 178
      - API overview 178
      - CallingMethod 181
        - signaturev1 181
        - signaturev3 189
        - request 198
        - result 199
        - PublicParameter 202
      - Cloud Disk APIs 204
        - AttachDisks 204
        - CreateDisks 205
        - DescribeBlockStorages 206

- DescribeDiskAssociatedSnapshots 207
- DescribeDiskConfigQuota 209
- DescribeDiskSupportFeatures 211
- DescribeDisks 213
- DescribeDisksDeniedActions 214
- DescribeInstancesDiskNum 216
- DescribeUserDiskResources 218
- DescribeZoneDiskConfigInfos 219
- DetachDisks 221
- InitializeDisks 222
- InquirePriceModifyDiskExtraPerformance 223
- InquiryPriceCreateDisks 225
- InquiryPriceModifyDiskAttributes 226
- InquiryPriceResizeDisk 228
- ModifyDiskAttributes 229
- ModifyDiskExtraPerformance 230
- ResizeDisk 231
- SwitchParameterCreateDisks 232
- SwitchParameterResizeDisk 233
- TerminateDisks 234

• Other APIs 236
- DescribeDedicatedClusterDiskStatistics 236
- ModifyDisksSynchronizeFlag 238

• Snapshot APIs 240
- ApplySnapshot 240
- ApplySnapshotGroup 242
- BindAutoSnapshotPolicy 244
- CopyAutoSnapshotPolicyCrossAccount 246
- CopySnapshotCrossAccount 248
- CopySnapshotCrossRegions 250
- CreateAutoSnapshotPolicy 251
- CreateSnapshot 252
- CreateSnapshotGroup 253
- DeleteAutoSnapshotPolicies 255
- DeleteSnapshotGroup 257
- DeleteSnapshots 258
- DescribeAutoSnapshotPolicies 259
- DescribeCopyAutoSnapshotPolicy 260
- DescribeDiskAssociatedAutoSnapshotPolicy 261
- DescribeSnapshotGroups 263
- DescribeSnapshotGroupsDeniedActions 265
- DescribeSnapshotSharePermission 267
- DescribeSnapshots 269
- DescribeSnapshotsDeniedActions 271
- GetSnapOverview 273
- ModifyAutoSnapshotPolicyAttribute 275
- ModifySnapshotAttribute 276
- ModifySnapshotsSharePermission 278
- UnbindAutoSnapshotPolicy 279

• Data structure 281

• ErrorCode 287

Product Introduction
CBS Overview

What Is CBS

Cloud Block Storage (CBS) is a block storage device with high availability, high reliability, high performance, low cost and customization, which can be used as an independent and scalable hard disk of Cloud Virtual Machine (CVM) to provide efficient and reliable storage for CVM instances. CBS provides persistent storage at the data block level. It is usually used as the main storage device for data (such as file system and database) that needs frequent and fine-grained updates, featuring high availability, high reliability and high performance.

CBS supports triple replication and EC distributed mechanism, backing up your data on different physical machines to avoid data loss caused by single point of failure, thereby improving data reliability. You can easily purchase, adjust, and manage your cloud disk devices through the console and create storage space larger than a single cloud disk capacity by building a file system.

According to different types of lifecycle, cloud disks can be divided into:

- Non-elastic cloud disk: Its lifecycle completely follows that of Cloud Virtual Machine. And it is purchased together with the Cloud Virtual Machine and used as a system disk, while it cannot be mounted and unmounted.
- Elastic cloud disk: Its lifecycle is independent of the Cloud Virtual Machine. It can be purchased separately and manually mounted to the Cloud Virtual Machine, or purchased together with the Cloud Virtual Machine and automatically mounted to the Cloud Virtual Machine as a data disk. Elastic cloud disk can be mounted or unmounted on the Cloud Virtual Machine in the same availability zone at any time. You can mount multiple elastic cloud disks to the same Cloud Virtual Machine, or you can unmount the elastic cloud disk from Cloud Virtual Machine A and then mount it on Cloud Virtual Machine.

Typical Use Cases

- If the hard disk space of a CVM is insufficient during use, you can purchase one or more cloud disks and mount to the CVM to meet the storage capacity requirements.
- No additional storage space is required when you purchase a CVM. When there is a need for storage, you can purchase a cloud disk to scale out the storage capacity of the CVM.
- When there is a demand for data exchange between multiple CVMs, you can unmount the cloud disk (data disk) and remount it to other CVMs.

- You can purchase multiple cloud disks and configure logical volume of LVM (Logical Volume Manager) to reach the upper limit of storage capacity of a single cloud disk.
- You can purchase multiple cloud disks and configure the RAID (Redundant Array of Independent Disks) policy to reach the upper limit of I/O capacity of a single cloud disk.

Product Features

The cloud platform provides diversified persistent storage devices. Users can flexibly select the type of hard disk, and store files and build databases on the hard disk.

- Hard Drive Selection: Premium Cloud Disk and SSD.
- Elastic mounting and unmounting: All types of elastic cloud disk support elastic mounting and unmounting. Multiple cloud disks can be mounted on the Cloud Virtual Machine to build a large-capacity file system.
- Elastic scale-out: You can scale out the cloud disk at any time, and a single disk supports a maximum of 32TB.
- Snapshot backup: It supports snapshot creation and rollback to back up key data. It also supports using snapshots to create hard disks for rapid service deployment.

Strengths

Reliability

CBS adopts a distributed mechanism with triple replicas. The system only returns a response confirming successful writing after ensuring that the data has been written to all three copies. The backend data replication mechanism can quickly create a copy through data migration or other methods if any one of the copies fails, always ensuring that three copies are available. This provides you with a secure and reliable data storage service. Data is stored across racks, achieving a reliability of 99.9999999%.

Elasticity

You can freely configure storage capacity and expand on demand without interrupting your business.
The maximum capacity of a cloud disk is 32 TB, and a single CVM can mount up to 20 elastic cloud disks as data disks, efficiently handling big data processing scenarios with TB/PB-level data.

High Performance

SSD adopts the NVMe standard high-performance SSD and provides up to 26,000 random IOPS and a maximum throughput of 260 MB/s per disk, achieving a combination of superior performance and ultra-high reliability.

Premium Cloud Disk helps you easily support high-throughput DB access on the business side.

Usability

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

You can create snapshots for cloud disks at any time to back up data, or quickly create cloud disks by using snapshot files to rapidly deploy business operations.

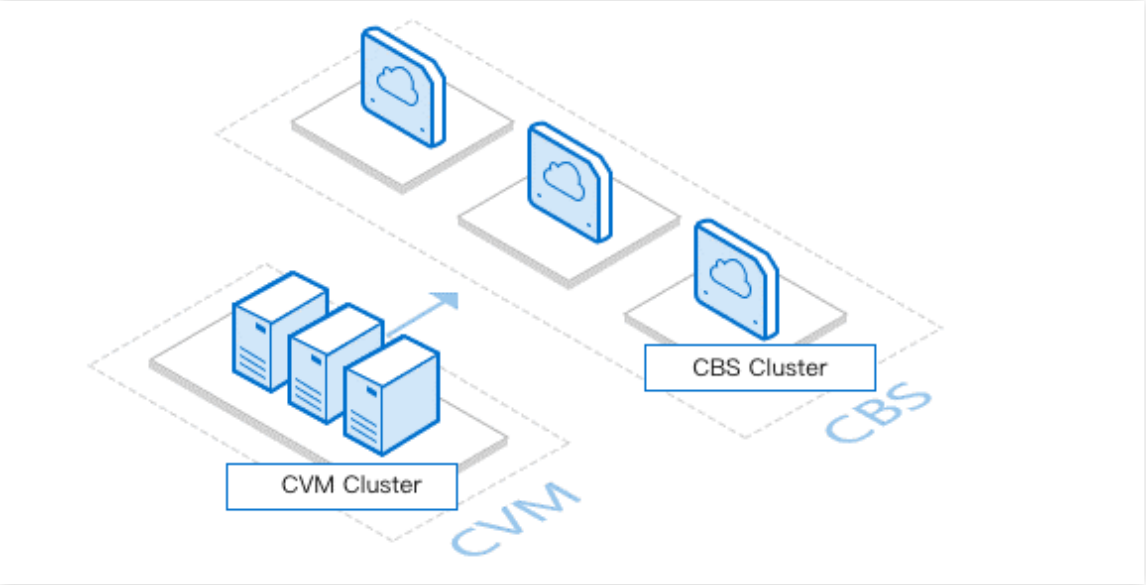
Use Cases

Typical Use Cases for CBS

De-localization

- High-Performance and High-Reliability Data Storage
  CBS can efficiently support hot migration of CVM and avoid business interruption caused by physical failures in advance. CBS provides three copies of data redundancy, with perfect data backup, snapshot and second-level data recovery capabilities. It is suitable for high-load and core critical business systems.

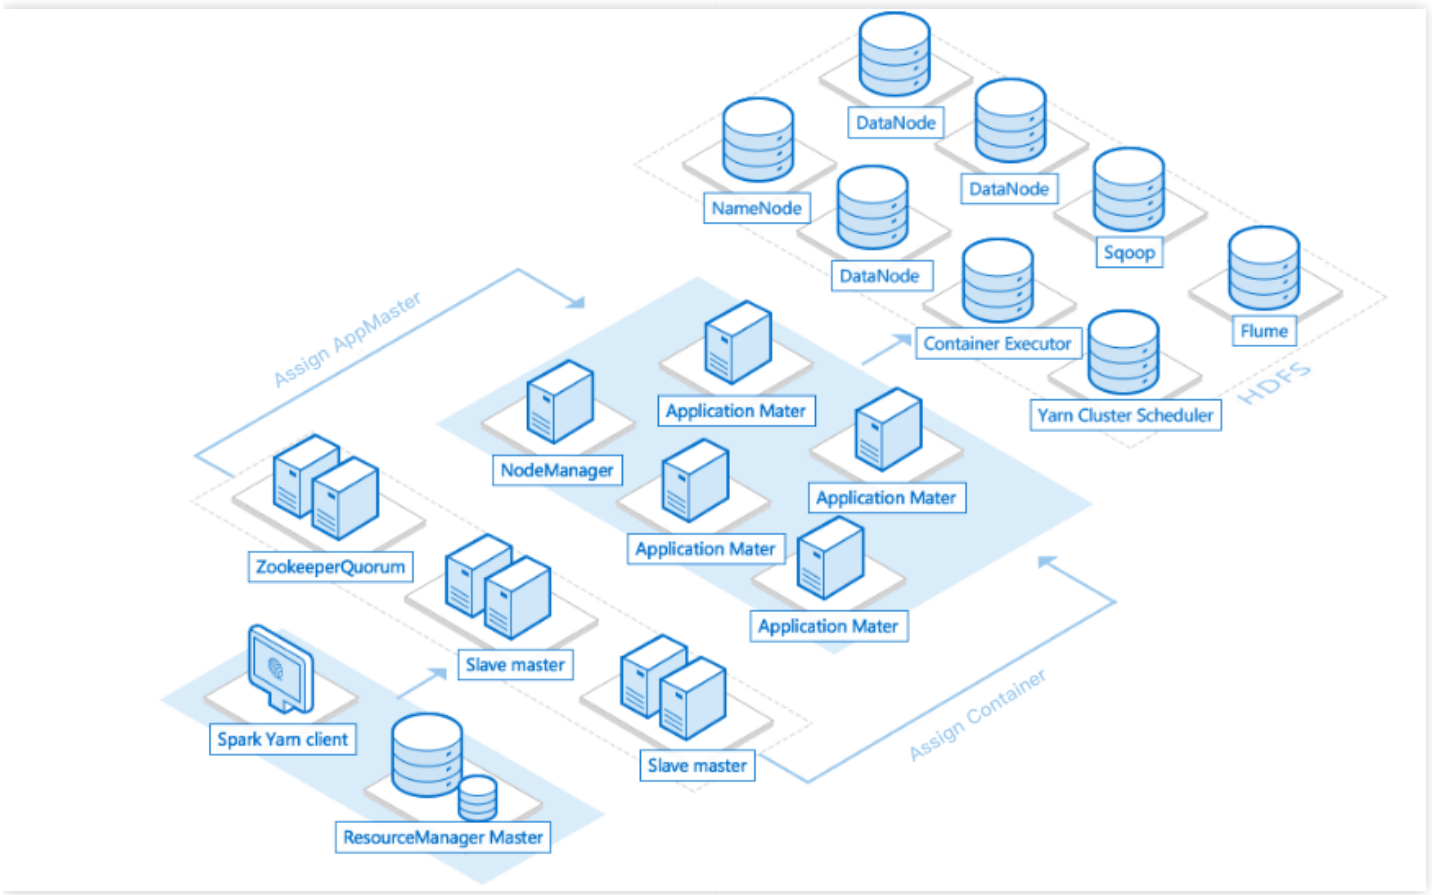
- Elastic Scale-out
  Cloud disks can be freely attached and detached within the same availability zone without shutting down or restarting the CVM. The capacity of cloud disks can be elastically configured and expanded on demand.



Massive Data Analysis

In the typical Spark-HDFS offline data analysis framework, RDD read/write and shuffle write are sequential I/O for disk read/write. Only shuffle read I/O is random I/O, and the proportion of sequential I/O is as high as 95%. CBS has excellent multi-thread concurrent throughput performance and efficiently supports offline processing data of Hadoop-Mapreduce, HDFS, Spark, and TB/PB levels.

Multiple disks are concurrent, and the throughput performance of a single HDFS cluster can reach 1GB/s.



Deployment Environment

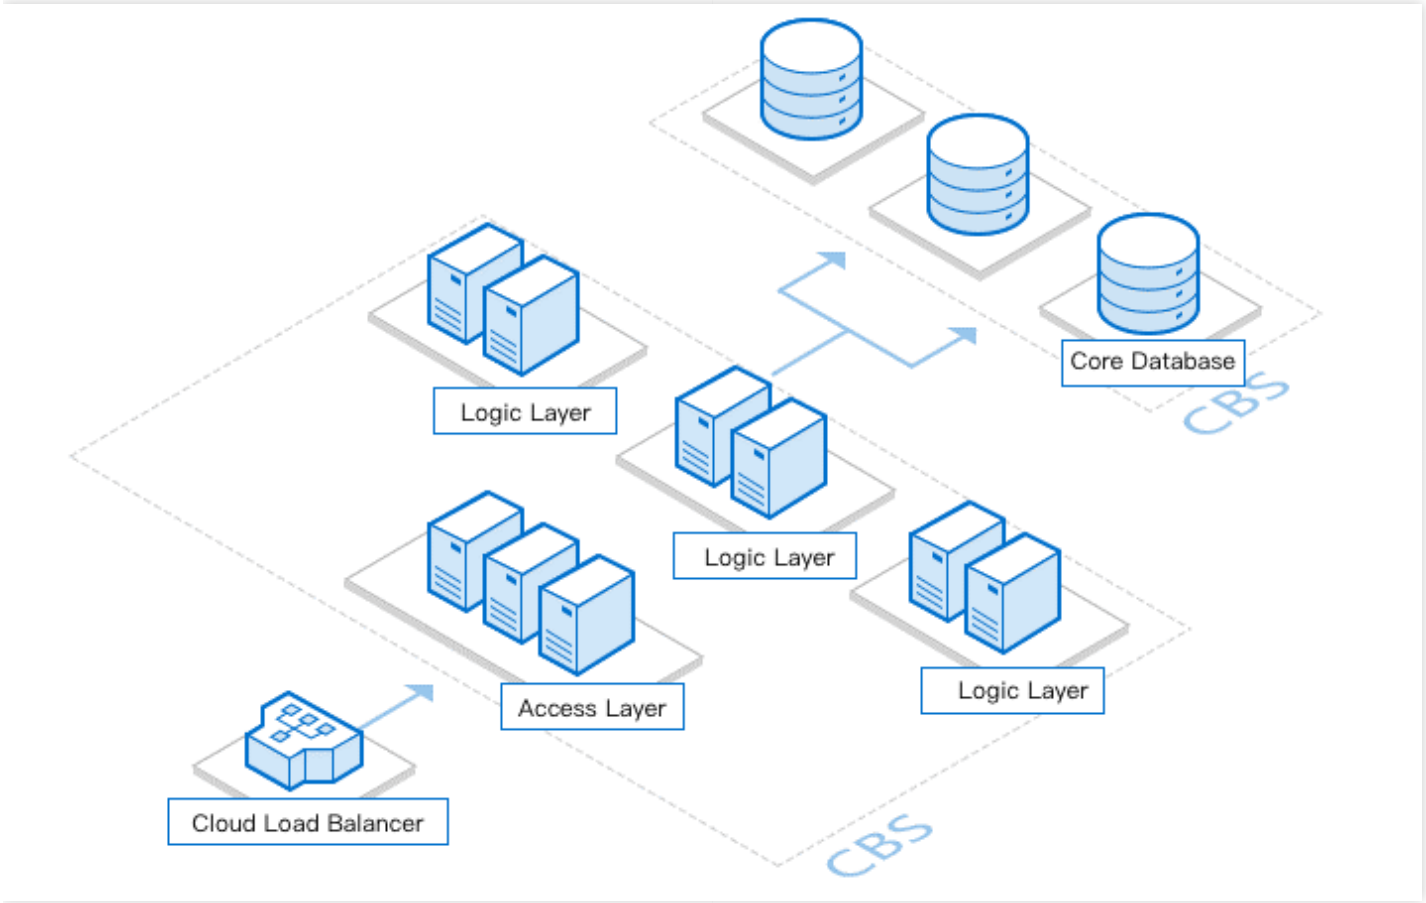
Five CVM instances with 12 cores and 40 GB RAM, simulating the offline data to analyze 1.5TB data volume.

Test Performance

- Each CVM has one attached HDD cloud disk of 1 TB. Five HDD cloud disks provide a read speed of 500 MB/s, and the data is read into memory in 50 minutes.
- Each CVM has one attached SSD of 1 TB, and the data is read into memory in 25 minutes.

Core Database

A SSD is suitable for scenarios requiring high I/O performance and high data reliability. It is particularly suitable for medium-to-large relational database applications such as PostgreSQL, MySQL, Oracle, and SQL Server, I/O-intensive core workload systems with high data reliability requirements, and medium-to-large development and test environments with high data reliability requirements.



Deployment Environment

Four CVM instances with 4 cores and 8 GB RAM, each attached with one 800G SSD and deployed with MySQL version 5.5.42.

Test Performance

Simulate the OLTP performance test with sysbench. The test set has 10 million records, TPS up to 1616 and QPS up to 29000. A single disk is enough to support tens of thousands of online simultaneous transactions per second.

Cloud Disk Instances
Cloud Disk Types

A cloud disk is a highly available, reliable, low-cost, customizable network block device that can be used as an independent and scalable hard disk for CVM. It provides block-level data storage and uses a distributed mechanism with triple replicas and EC to ensure data reliability for the CVM. Different hard disk types vary in performance, features, and pricing. You can choose based on the requirements of your deployed applications.

- Premium Cloud Disk: Premium Cloud Disk is a hybrid storage type that provides high-performance storage capabilities close to solid-state storage through a cache mechanism. It uses a distributed mechanism with three replicas to ensure data reliability. Premium Cloud Disk is suitable for small-to-medium-sized application scenarios with high data reliability requirements and moderate performance requirements, such as Web/App servers, business logic processing, and small-to-medium website building.
- SSD: SSD is based on full NVMe SSD storage medium, uses a distributed mechanism with three replicas, and provides low latency, high random IOPS, high throughput I/O capacity, and data security up to 99.9999999%. Cloud SSD is suitable for scenarios with high I/O performance requirements.

The performance metrics of different types of cloud disk are as follows:

Performance Metrics	SSD	Premium Cloud Disk
Maximum capacity per disk (GiB)	32000	32000
Maximum IOPS of a single disk	26000	6000
Calculation formula for random IOPS performance.	Maximum random IOPS = 1800 + Storage capacity (GiB) x 30 The maximum random IOPS shall not exceed 26000.	Maximum random IOPS = 1800 + Storage capacity (GiB) × 8 The maximum random IOPS shall not exceed 6000.
Maximum throughput of a single disk (MB/s)	260MB/s	150MB/s
Calculation formula for throughput performance (MB/s)	Maximum throughput = 120 + Storage capacity (GiB) x 0.2 The maximum throughput shall not exceed 260 MB/s.	Maximum throughput = 100 + Storage capacity (GiB) x 0.15 The maximum throughput shall not exceed 150 MB/s.
Single-channel random read/write latency	0.5 - 3ms	0.8 - 4ms

The main difference between SSD and Premium Cloud Disk lies in I/O performance.

SSD is more suitable for scenarios with medium-to-large loads, for example:

- Medium and large database: It can support medium and large relational database applications such as MySQL.
- Image process: It can support data analysis and storage services such as image process.

Premium Cloud Disk mainly applies to the following data scenarios:

- It is suitable for small and medium databases, Web/App servers, etc., providing long-term stable I/O performance output.
- It is suitable for scenarios that require a balance between storage capacity and performance, such as enterprise office services.
- Meet the I/O requirements of core business testing and development and joint debugging environment.

Cloud Disk Status

A cloud disk has the following states:

Status	Attributes	Description
To be mounted	Steady State	The state of a cloud disk after it is successfully created and before it is attached to a CVM.
Mounting	Intermediate State	The state of a cloud disk when it is being attached and before it is in use.
In use	Steady State	The state of a cloud disk after it has been attached to a CVM in the same availability zone.
Unmounting	Intermediate State	The state of a cloud disk when it is being detached and before it enters the state of to be attached.
Expanding	Intermediate State	During the scaling out of cloud disk.

Cloud Disk Encryption

When your business requires encryption protection for data stored on cloud disks due to security or compliance requirements, you can enable the cloud disk encryption feature provided by Key Management Service (KMS) to effectively protect data privacy.

Key Management

The industry-standard AES-256 algorithm is used to encrypt your cloud disk data with a data key. When you use an encrypted cloud disk for the first time, the system automatically creates a customer master key (CMK) in KMS for the corresponding region, dedicated to cloud disk encryption. This automatically created key is unique and stored in the KMS protected by strict physical and logical security controls.

Encrypted cloud disks in each region use a unique 256-bit data key (DK) for encryption. Snapshots created from encrypted cloud disks and encrypted cloud disks created from these snapshots are all associated with this key. The key is protected by KMS, effectively preventing unauthorized access. The DK of cloud disk is only used in the memory of the instance's host and is never stored in plaintext on any persistent medium (including the cloud disk itself).

How It Works

When you set a cloud disk as an encrypted cloud disk, KMS encrypts the data and decrypts it automatically when reading it. The encryption and decryption process runs on the host where the CVM instance is located, with almost no impact on the read-write performance of the cloud disk. For how to test the performance of a cloud disk, see How to Measure the Performance of a Cloud Disk.

After an encrypted cloud disk is created and attached to an instance, the system encrypts the following data:

- Static data in the cloud disk.
- Data transmitted between the cloud disk and the instance (data in the instance's operating system is not encrypted).
- All snapshots created from the encrypted cloud disk.

Use Limits

The encryption feature for cloud disks has the following limitations:

Limit Type	Description
Limitations on cloud disks	- Cloud disk encryption currently supports all cloud disk types and instance types. - Only the cloud disk can be encrypted, but not the local disk. - Only the data disk can be encrypted, but not the system disk. - The existing non-encrypted disk cannot be directly converted into an encrypted disk. - The encrypted cloud disk cannot be converted into a non-encrypted cloud disk. - After an encrypted cloud disk is expanded, it should be detached from the CVM and then attached again, so that the new capacity can be recognized. - Encrypted cloud disks cannot be attached to instances with local storage.
Limits on snapshot and image	- The snapshot generated by existing non-encrypted disk cannot be directly converted into an encrypted snapshot. - The encrypted snapshot cannot be converted into a non-encrypted snapshot. - The image with an encrypted snapshot cannot be shared. - The snapshot generated by existing non-encrypted disk cannot be directly converted into an encrypted snapshot. - Encrypted snapshots and images created by encrypted snapshots cannot be copied across regions.
Other limits	- The cloud disk encryption feature depends on the KMS in the same region. If you have no other operation requests, no additional operations are required in the KMS console. - When using the cloud disk encryption feature for the first time, you need to authorize and activate KMS according to the page prompts. Otherwise, you will not be able to purchase encrypted cloud disks. - The CMK specifically created by the system for cloud disk encryption can be queried through the KMS console, but it cannot be specified, deleted, or modified by users.

Operation Fees

The CBS encryption feature and the CMK used do not incur additional fees, nor do read and write operations on the cloud disk data. However, when managing the encrypted cloud disk, whether through the console or API, KMS will be used as an API. This will count towards your KMS invocation count in that region. KMS is itself billed based on the number of invocations.

Management operations for encrypted cloud disks include:

- Create an encrypted cloud disk.
- Attach a cloud disk.
- Detach a cloud disk.

- Create a snapshot.
- Roll back snapshot.

> Note:
>
> Ensure that your account balance is sufficient, otherwise the operation will fail.

Creating an Encrypted Cloud Disk

You can create an encrypted cloud disk using the following three methods:

Creating an Encrypted Cloud Disk via the Console

1. Log in to the CBS console, select the corresponding region, and then click Create.
2. In the "Purchase Data Disk" dialog box, check the Encrypt option.

   > Note:
   >
   > If this is your first time using encrypted cloud disks in this region, you need to authorize KMS first.

3. Select the CBS configuration as needed and click OK.
4. After the purchase is complete, you can view the created encrypted cloud disk on the CBS list page. The newly created encrypted cloud disk is in the To be mounted status. To mount the cloud disk to the Cloud Virtual Machine in the same availability zone, see Mounting Cloud Disks.

Creating an Encrypted Cloud Disk via a Snapshot

As described in Creating a Cloud Disk via a Snapshot, select an encrypted snapshot to create an encrypted cloud disk that already contains relevant data.

Creating an Encrypted Cloud Disk via API

You can use the CreateDisks API to create an encrypted cloud disk in the following two ways:

- Specify the encryption option `Encrypt` to `true`.
- Specify `SnapshotId` of the encrypted snapshot.

Convert Data Encryption Status

If you need to convert existing data on the cloud disk from non-encrypted to encrypted, it is recommended that you use the rsync command in Linux or the robocopy command in Windows to copy data from the non-encrypted cloud disk to the new encrypted cloud disk.

If you need to convert existing data on the cloud disk from encrypted to non-encrypted, it is recommended that you use the same command to copy data from the encrypted disk to a new non-encrypted disk.

Enhanced SSD Performance Descriptions

The single-disk performance of the Enhanced SSD is composed of the baseline performance and the additional performance, where:

- Baseline performance is linearly related to the cloud disk capacity and reaches the maximum baseline performance at the critical point.
- If you have requirements that exceed the benchmark performance, you can enable the extra performance to obtain the higher single-disk performance.

This document will introduce the benchmark performance and extra performance separately.

Total Performance of an Enhanced SSD

The performance of each enhanced SSD consists of benchmark performance and extra performance. Regardless of how it is allocated, the performance of the current enhanced SSD follows the following rules:

Performance of Enhanced SSD	Maximum Value
Random IOPS	100000
Maximum throughput (MB/s)	1000

Baseline Performance of Enhanced SSD

For enhanced SSD, the benchmark performance varies linearly with capacity and cannot be adjusted independently. The calculation formula is as follows:

Baseline Performance of Enhanced SSD	Calculation Formula for Benchmark Performance	Maximum Value of Benchmark Performance
Benchmark Performance - Random IOPS	min {1800 + Capacity (GiB) × 50, 50000}	50000
Benchmark Performance - Maximum Throughput (MB/s)	min {120 + Capacity (GiB) × 0.5, 350}	350

It can be calculated from the above formula:

- When the capacity reaches 460 GiB, the upper limit of throughput is reached first. At this point, the baseline performance values are: 24,800 random IOPS and 350 MB/s maximum throughput.
- When capacity reaches 964 GiB, the random IOPS reaches its upper limit. At this point, the baseline performance values are: 50,000 random IOPS and 350 MB/s maximum throughput.

Additional Performance of Enhanced SSD

If you have requirements that exceed the maximum value of benchmark performance, you can enable the extra performance to obtain the higher performance.

Prerequisites

The following conditions must be met to enable the extra performance:

- Currently, this feature is only available to Enhanced SSD, other types of cloud disk do not support extra performance.
- Only when any item of baseline performance reaches its maximum value can additional performance be configured. According to the baseline performance calculation formula, it is supported only when the capacity is higher than 460 GiB.

Extra Performance Formula

Extra performance is configurable by the user, and the calculation formula is as follows:

Additional Performance of Enhanced SSD	Extra Performance Formula	Maximum Value of Extra Performance
Random IOPS	min {configuration value × 128, 50000}	50000
Max throughput (MB/s)	min {configuration value × 1, 650}	650

Example

Example 1: An Enhanced SSD with a capacity requirement of 2,000 GiB and a throughput requirement of 500 MB/s.

According to the baseline performance description, it can be seen that the baseline performance (throughput) has reached its upper limit of this configuration capacity. To meet the requirements, an additional performance of (500 - 350)/1 = 150 needs to be purchased. Therefore, the overall configuration to be purchased is: 2,000 GiB Enhanced SSD with additional performance configured as 150.

Example 2: An Enhanced SSD with a capacity requirement of 1,000 GiB and an IOPS requirement of 50,000.

According to the baseline performance description, it can be seen that the baseline performance (random IOPS) has reached its upper limit of this configuration capacity. However, since the user does not need more IOPS performance, there is no need to configure additional performance. Therefore, the overall configuration to be purchased is: 1,000 GiB Enhanced SSD without additional performance.

Cloud Disk Snapshots
Snapshot Overview

Description

- Real-Time Replica of Online Data
  A snapshot is a fully available copy of a cloud disk. When a cloud disk with created snapshots encounters an issue, you can quickly restore it to the state before the issue occurred via the snapshot. We recommend that you create snapshots for the relevant cloud disks before major business operations changes. Therefore, if the change fails, data can be quickly recovered.

- Persistent Backup of Key Milestones
  Snapshots can serve as persistent backups of business data, preserving milestone status of business data.

- Rapid business deployment
  You can use snapshot files of your business to quickly clone multiple cloud disks for rapid server deployment.

Use Cases

Snapshot is a convenient and efficient data protection service and is recommended for the following business scenarios:

- Daily Data Backup
  You can use snapshots to regularly back up important business data to mitigate the risk of data loss caused by misoperation, attacks, or viruses.

- Quick Data Recovery
  You can create one or more snapshots before performing major operations such as changing the operating system, upgrading application software, or migrating business data. If any problem occurs during the change operation, the business data can be restored in time through the created snapshot.

- Multi-replica Application of Production Data
  You can create production data snapshots to provide near-real-time production data for applications

such as data mining, report query, development and testing.

- Rapid Deployment Environment
  You can create a snapshot of CVM and use this system snapshot to create custom images. You can create one or more instances from the created image to quickly deploy CVMs in the same environment in batches and save time on repeated configuration.

Snapshot Type

- Manual Snapshot
  You can manually create a snapshot of the cloud disk data at a specific point. This snapshot can be used to quickly create more cloud disks with the same data or restore the cloud disk to its state at that point in the future.

- Scheduled Snapshot
  When your business undergoes continuous updates, you can use scheduled snapshots to provide a continuous backup feature. Simply create a backup policy and associate it with cloud disks to achieve continuous backup of cloud disk data during a certain period, significantly improving data security.

> Note:
>
> During the snapshot creation process, some application data may be saved in memory but not persisted. This phenomenon may cause the snapshot to fail to capture the latest and most complete cloud disk data.

Case Review

Case 1: No manual snapshots were created before performing high-risk operations, resulting in data loss

For example, customer A has never created snapshots for cloud disks. On a certain day in May 2019, due to the operator performing an FIO test on the cloud disk, the file system was corrupted, and the data could not be recovered.

Analysis: If customer A takes a snapshot of the CBS before testing, then after data corruption, snapshot rollback can be quickly enabled to restore business on time.

Case 2: No scheduled snapshots were created for important data disks, resulting in data loss

For example, customer B has created snapshots for multiple cloud disks. However, after January 2019, due to cost considerations, no snapshots are created for newly purchased cloud disks. On a certain day in June

2019, data on an unprotected cloud disk was accidentally deleted at the file system level and could not be recovered.

Analysis: If customer B has performed regular snapshot protection on the cloud disk, then if the data is accidentally deleted, it can be restored to the data status at the last snapshot time point to reduce losses. After the incident, customer B took the initiative to create a snapshot for the cloud disk, which fully enhanced data protection.

Case 3: Using the scheduled snapshot to roll back data and restore business after an incorrect operation

For example, customer C has implemented snapshot protection for all cloud disks. On a certain day in May 2019, an abnormal start-up occurred due to misoperation.

Analysis: Customer C promptly restored the data using the regular snapshot from 2 days ago, and the business was not affected.

The above cases are all data loss caused by operational errors. By comparison, we can find that:

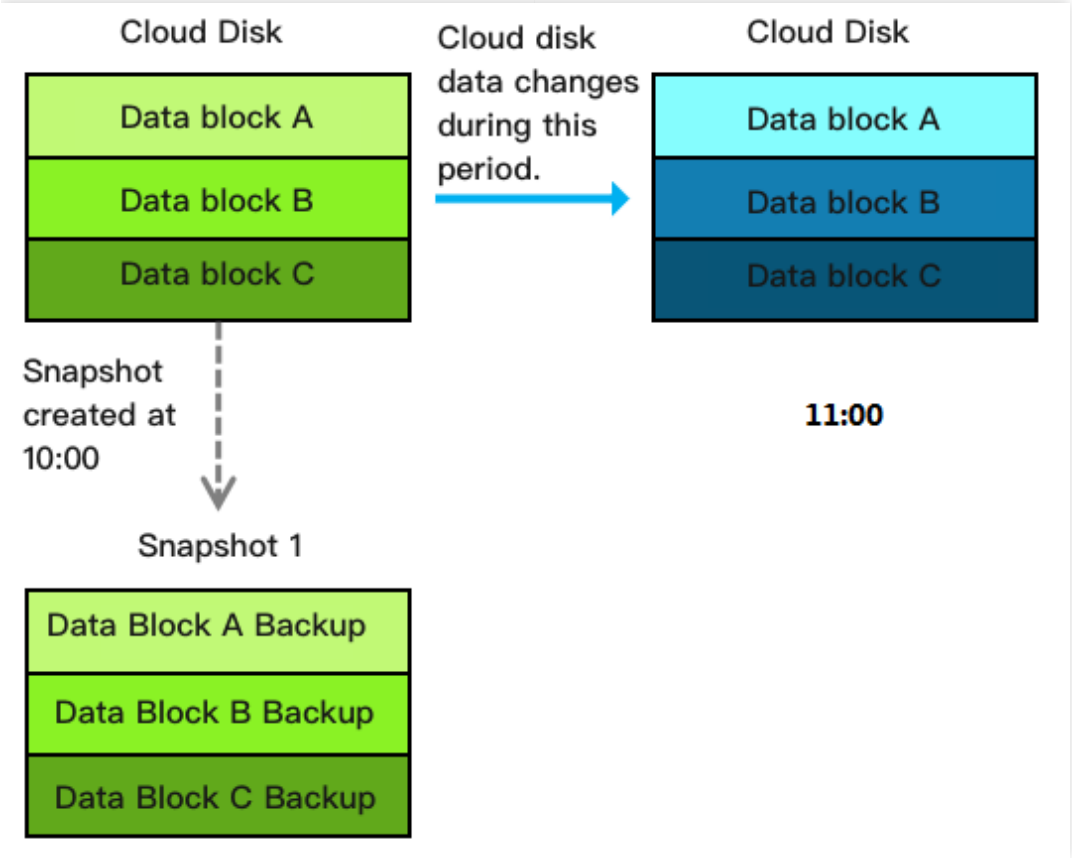
- If no snapshots have been created, data recovery becomes very difficult if servers or cloud disks encounter issues, potentially causing significant losses.
- If snapshots have been created, data can generally be recovered if servers or cloud disks encounter issues, resulting in smaller losses.
  It is recommended that you create snapshots periodically for different types of business to improve data security and achieve low-cost and high disaster recovery for the business.

Snapshot Principle

Snapshot and Source Disk

A snapshot is a data backup of the cloud disk at a certain time. The write-in and modification of data in the cloud disk will not affect the created snapshot. Based on this feature, you can use snapshot to record the data of cloud disk at different times for system recovery, disaster recovery and cloud disk replication.

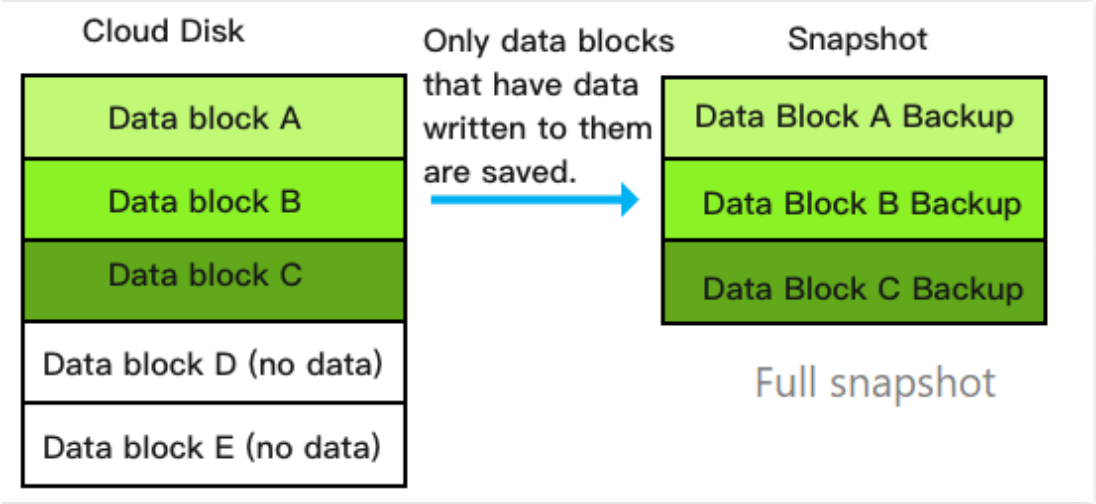
As shown in the following figure, snapshot 1 retains the information of data blocks on the cloud disk at 10:00 (the time point when the snapshot is created), which has nothing to do with subsequent changes.



Snapshot Size and Source Disk

The snapshot only saves the written or modified data blocks in the cloud disk, so generally, the size of the snapshot corresponding to the cloud disk will be smaller than that of the cloud disk.

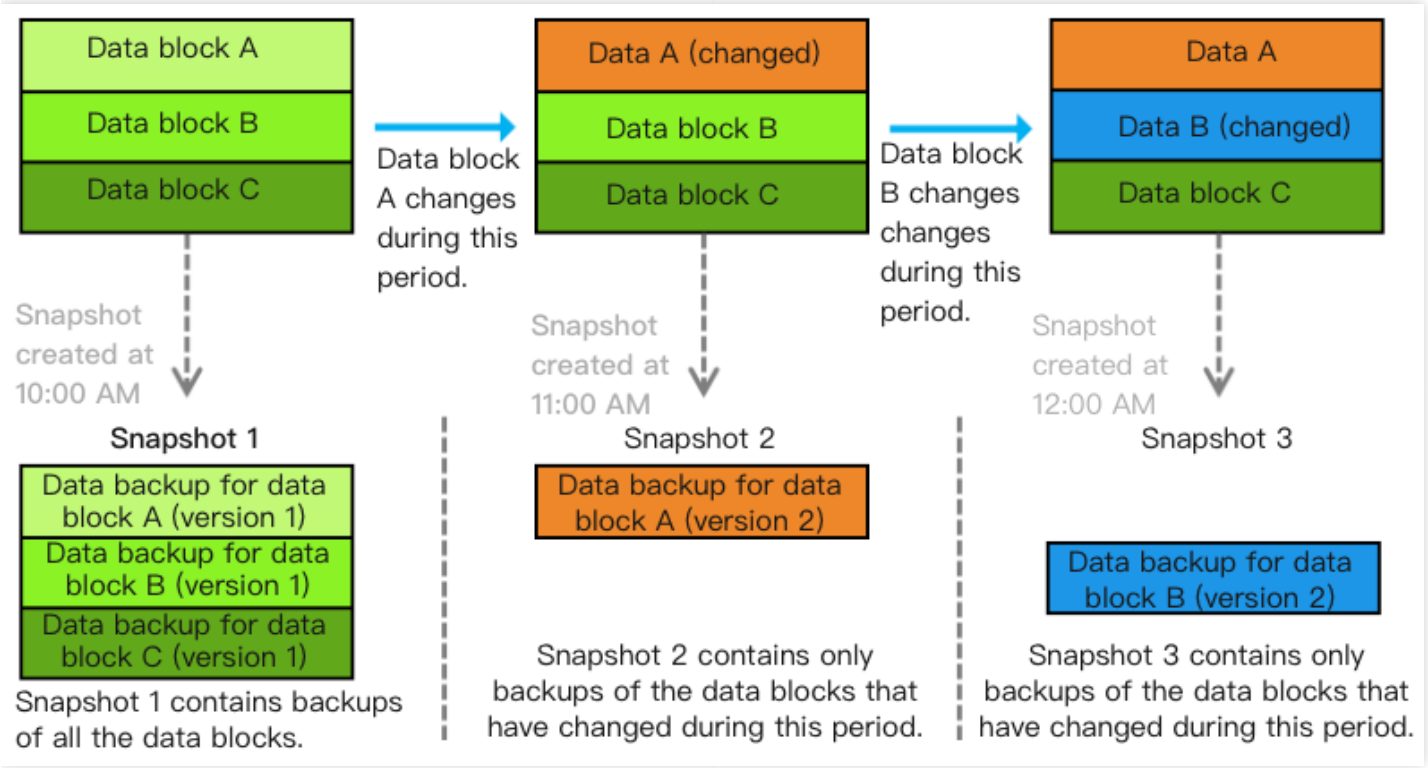
The relationship between snapshot size and source disk is as follows:



Description of Incremental Snapshot Creation Process

Snapshots use an incremental snapshot mechanism. When you continuously create multiple snapshots of the same CBS, only the first snapshot is a full snapshot, and subsequent snapshots only contain data changes relative to the previous snapshot (incremental snapshot). This minimizes the total storage capacity occupied by consecutive snapshots, reducing user overhead.
For example: Assume the CBS has three data blocks, A, B, and C. You create snapshots at 10:00, 11:00, and 12:00 respectively. Changes of data blocks on the CBS during each time interval are shown in the following figure. Each snapshot should contain the following data:

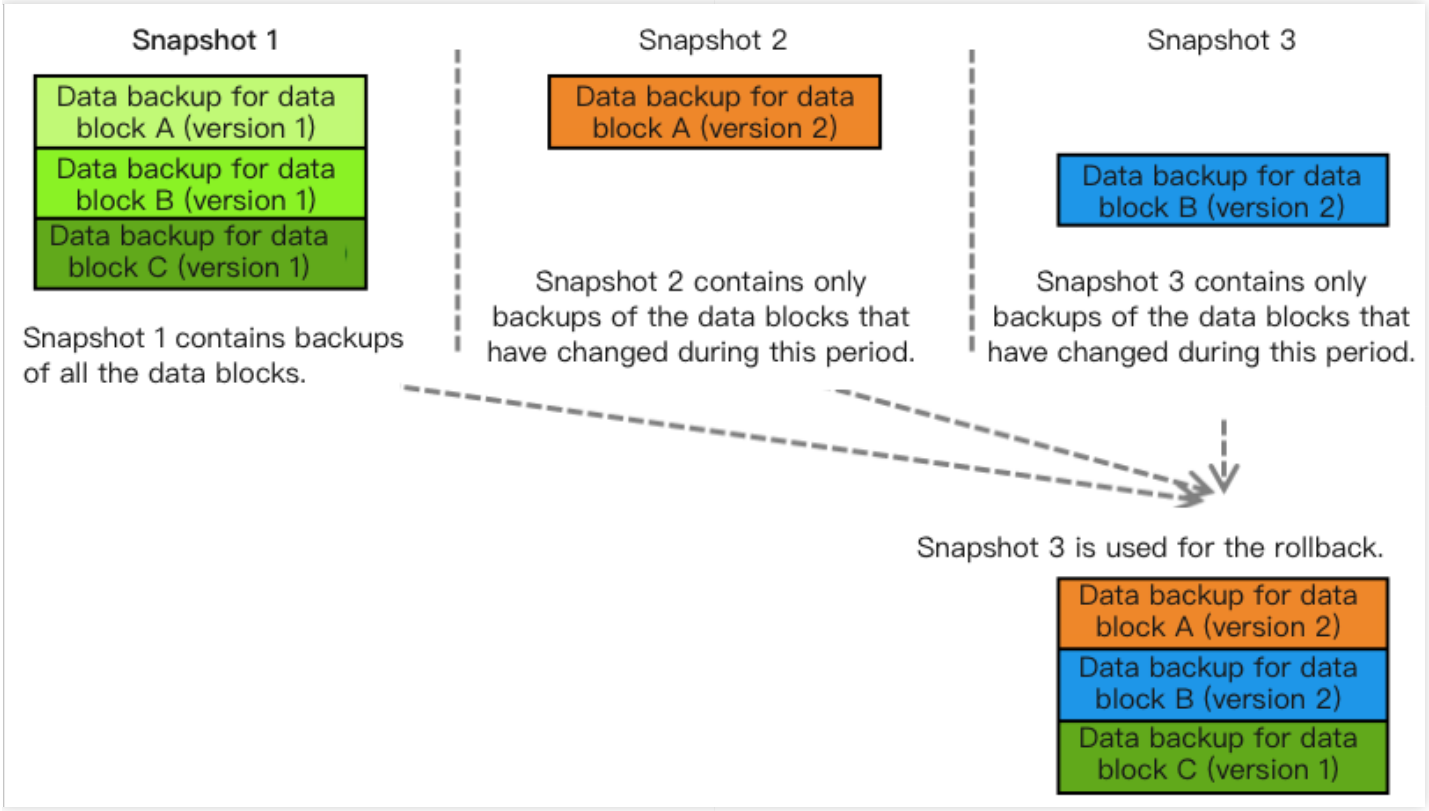
- Snapshot 1 (the first snapshot): It refers to the data backup that contains all data blocks on the cloud disk at that time (generally called full snapshot).
- Snapshot 2: Data block A on the cloud disk changes during the period, and snapshot 2 only contains the latest data backup of data block A (generally called incremental snapshot).
- Snapshot 3: Data block B on the cloud disk changes during the period, and snapshot 3 only contains the latest data backup of data block B (generally called incremental snapshot).



Description of Incremental Snapshot Rollback Process

Based on the previous examples, when you use snapshot 3 for data rollback, the system will merge the data of snapshot 1, snapshot 2 and snapshot 3. If there are data blocks at the same position, the data in the latest snapshot will be taken. During the final rollback, the merged data set is written to the cloud disk to be rolled back.

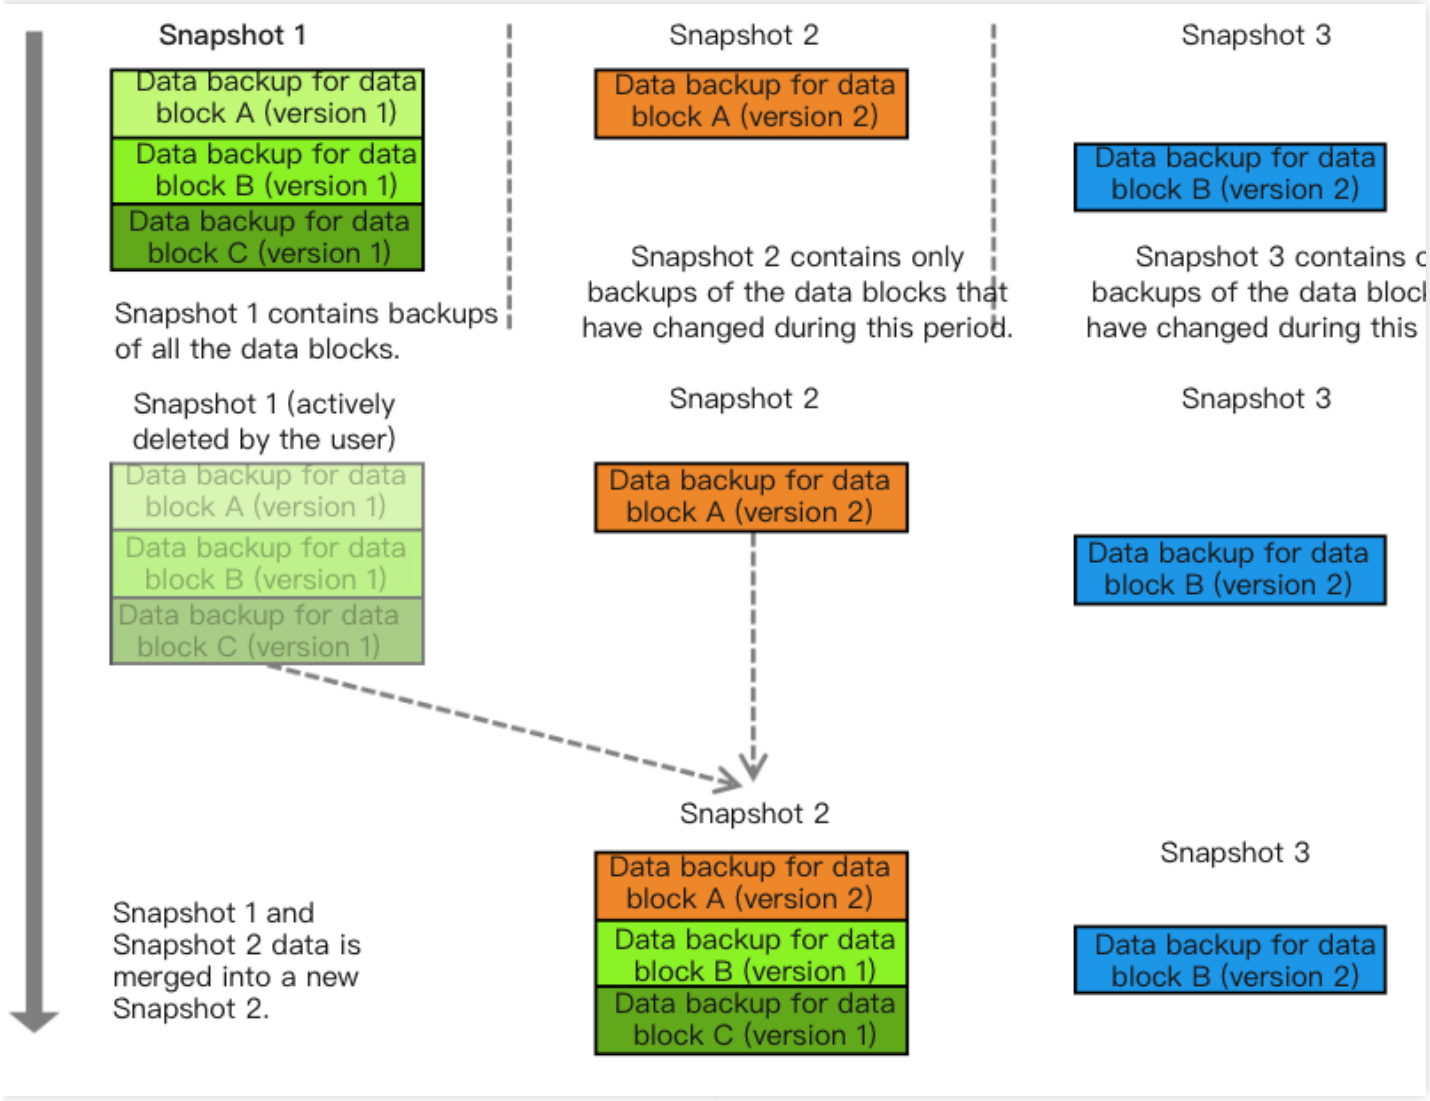
The incremental snapshot rollback process is shown in the following figure:



Description of Incremental Snapshot Deletion and Merging Process

- When the full snapshot (i.e. the first snapshot) is deleted, the system will automatically merge this full snapshot with the next incremental snapshot.
- When the incremental snapshot is deleted, the system will automatically merge this incremental snapshot with the next incremental snapshot. If there is no next incremental snapshot, the system will directly delete this incremental snapshot.
  Based on the previous examples, if you delete snapshot 1, the system will merge snapshot 1 and snapshot 2 and use the data in snapshot 2 to overwrite the data at the same position of snapshot 1. The merged snapshot 2 becomes a new full snapshot.

The incremental snapshot deletion and merging process is shown in the following figure:


Snapshot 1
Data backup for data block A (version 1)
Data backup for data block B (version 1)
Data backup for data block C (version 1)
Snapshot 1 contains backups of all the data blocks.
Snapshot 2
Data backup for data block A (version 2)
Snapshot 2 contains only backups of the data blocks that have changed during this period.
Snapshot 3
Data backup for data block B (version 2)
Snapshot 3 contains c
backups of the data block
have changed during this
Snapshot 1 (actively deleted by the user)
Data backup for data block A (version 1)
Data backup for data block B (version 1)
Data backup for data block C (version 1)
Snapshot 2
Data backup for data block A (version 2)
Snapshot 3
Data backup for data block B (version 2)
Snapshot 1 and Snapshot 2 data is merged into a new Snapshot 2.
Snapshot 2
Data backup for data block A (version 2)
Data backup for data block B (version 1)
Data backup for data block C (version 1)
Snapshot 3
Data backup for data block B (version 2)

CAM
CAM Overview

If you use CVM, CBS, VPC, database, and other services, which are managed by different users but share your cloud account tokens, the following problems will occur:

- The risk of your key being compromised is high since multiple users are sharing it.
- The access permission of other users is not under control. They can introduce security risks caused by misoperations.
  To prevent these problems, you can use sub-accounts to allow different users to manage different services. By default, a sub-account has no permission to use CVM or CVM-related resources. Therefore, we need to create policies to allow sub-accounts to use the resources or permissions they need.

Cloud Access Management (CAM) is a set of web services mainly used to help users securely manage access permissions to resources under their cloud platform accounts. With CAM, you can create, manage, and terminate users (groups), and control who can use which resources through identity and policy management.

You can use CAM to bind a user or user group to a policy that allows or denies them access to specified resources to complete specified tasks.

If you do not need to manage the cloud access to CLB-related resources for your sub-account, you can skip this section. Skipping these sections does not affect your understanding and use of the rest of the documentation.

Getting Started

The CAM policy must authorize or deny one or more CBS-related operations. You must also specify the resources that can be used for the operation (it can be all resources, or some operations can be partial resources). The policy can also contain the conditions set by the operation resources.

Task	Link
Understanding the Basic Policy Structure	Policy Syntax
Defining Operations in a Policy	CBS-related Operations
Defining Resources in a Policy	CBS-related Resource Path
Using Conditions to Restrict Policies	CBS Conditional Key
Resource-Level Permissions Supported by CBS	Resource-Level Permissions Supported by CBS

Console Example	Example of Cloud Access Management

Authorizable Resource Types

Resource-level permissions refer to the ability to specify which resources a user can perform operations on. CBS supports resource-level permissions, meaning that for CBS operations supporting resource-level permissions, it controls when users are allowed to perform operations or which specific resources users can use.
The types of resources that can be authorized in CAM are as follows:

Resource Type	Resource Description Method in Authorization Policies
CBS instance-related	qcs::brc:$region::volume/*

CBS instance-related introduces the CBS API operations that currently support resource-level permissions and the resources and conditional keys supported by each operation. When setting the resource path, you need to modify variable parameters such as $region and $account to your actual parameter information. You can also use the * wildcard in the path.

> Note:
>
> CBS API operations not listed in the table indicate that the API operations do not support resource-level permissions. For API operations that do not support resource-level permissions, you can still grant users permission to use the operation, but the resource element of the policy statement must be specified as * .

CBS Instance-related

API Operation	Resource Path	Conditional Key
Attaching cloud disks (AttachDisks)	qcs::cvm:$region:$account:volume/* qcs::cvm:$region:$account:volume/$diskId	cvm:region cvm:zone cvm:disk_type
Creating cloud disks (CreateDisks)	qcs::cvm:$region:$account:volume/* qcs::cvm:$region:$account:volume/$diskId	cvm:region cvm:zone cvm:disk_type
Querying the list of cloud disk operation logs (DescribeDiskOperationLogs)	qcs::cvm:$region:$account:volume/* qcs::cvm:$region:$account:volume/$diskId	cvm:region cvm:zone cvm:disk_type

API Operation	Resource Path	Conditional Key
Querying the cloud disk list (DescribeDisks)	qcs::cvm:$region:$account:volume/* qcs::cvm:$region:$account:volume/$diskId	cvm:region cvm:zone cvm:disk_type
Detaching cloud disks (DetachDisks)	qcs::cvm:$region:$account:volume/* qcs::cvm:$region:$account:volume/$diskId	cvm:region cvm:zone cvm:disk_type
Modifying cloud disk attributes (ModifyDiskAttributes)	qcs::cvm:$region:$account:volume/* qcs::cvm:$region:$account:volume/$diskId	cvm:region cvm:zone cvm:disk_type
Expanding cloud disks (ResizeDisk)	qcs::cvm:$region:$account:volume/* qcs::cvm:$region:$account:volume/$diskId	cvm:region cvm:zone cvm:disk_type
Returning cloud disks (TerminateDisks)	qcs::cvm:$region:$account:volume/* qcs::cvm:$region:$account:volume/$diskId	cvm:region cvm:zone cvm:disk_type

Authorization Policy Syntax

Policy Syntax

CAM policy:

```
{
      "version":"2.0",
      "statement":
      [
        {
          "effect":"effect",
          "action":["action"],
          "resource":["resource"],
           "condition": {"key":{"value"}}
        }
      ]
}
```

- version: It is required, currently only allowing a value of 2.0 .
- statement: used to describe the detailed information of one or more permissions. This element includes the permissions or permission sets of multiple other elements such as effect, action, resource, condition, etc. A policy has one and only one statement element.
    i. action: used to describe the operations that are allowed or denied. Operations can be APIs (described with a name prefix) or feature sets (a specific set of APIs described with a permid prefix). This element is required.
    ii. resource: describes the specific data of the authorization. Resources are described using six paragraphs. The resource definition details will vary for each product. This element is required.
    iii. condition: describes the constraints under which the policy takes effect. Conditions consist of operators, operation keys, and operation values. Condition values can include information such as time and IP address. Some services allow you to specify additional values in the condition. This element is optional.
    iv. effect: describes whether the statement produces an allow or explicit deny result. There are two situations: allow and deny (explicitly deny). This element is required.

CBS Actions

In CAM policy statements, you can specify any API operation from any services that support CAM. For CBS, use APIs prefixed with name/cvm: . For example: name/cvm:CreateDisks or name/cvm:DescribeDisks .

If you want to specify multiple operations in a single statement, separate them with commas as shown below:

```
"action":["name/cvm:action1","name/cvm:action2"]
```

You can also use wildcard characters to specify multiple actions. For example, you can specify all operations whose names begin with the word Describe as follows:

```
"action":["name/cvm:Describe*"]
```

To specify all operations for the corresponding CVM, use the * wildcard as follows:

```
"action":["name/cvm:*"]
```

CBS Resource Paths

Each CAM policy statement has its own resources that apply to it. The general format of resource paths is as follows:

```
qcs:project_id:service_type:region:account:resource
```

- project_id: Describe the project information. It is only for compatibility with early CAM logic and does not need to be filled in.
- service_type: The product abbreviation, such as CVM.
- region: regional information, such as bj.
- account: the root account information of the resource owner, such as uin/164256472.
- resource: The specific resource details of each product, such as volume/diskid1 or volume/*.
  For example, you can specify it in a statement using the specific CBS resource (disk-abcdefg), as follows:

```
"resource":[ "qcs::cvm:bj:uin/164256472:volume/disk-abcdefg"]
```

You can also use the * wildcard to specify all CBS resources belonging to a specific account, as follows:

```
"resource":[ "qcs::cvm:bj:uin/164256472:volume/*"]
```

If you want to specify all resources, or if a specific API operation does not support resource-level permissions, use the * wildcard in the Resource element, as follows:

```
"resource": ["*"]
```

To specify multiple resources in one instruction, separate them with a comma. The following is an example of specifying two resources:

```
"resource":["resource1","resource2"]
```

CBS Conditional Key

In a policy statement, you can optionally specify conditions that control when a policy takes effect. Each condition consists of one or more key-value pairs. Conditional keys are not case sensitive.

- If you specify multiple conditions or multiple keys in a single condition, we evaluate them using a logical AND operation.
- If you specify a key with multiple values in a single condition, we will evaluate it using a logical OR operation. Permissions will be granted only if all conditions are met.
  The following table describes the service-specific condition keys used by CBS:

Condition Key	Reference Type	Key-Value Pair
cvm:region	String	cvm:region= region Where, region refers to the region (e.g. ap-city)
cvm:disk_type	String	cvm_disk_type= disk_type Where, disk_type refers to the disk type (e.g. CLOUD_PREMIUM)

Related Products

Product Name	Relationship with CBS
CVM	CBS can be used as the system disk and data disk of CVM to provide the elastic scalability for block storage devices.
Cloud Monitor	The performance metrics of CBS can be viewed through Cloud Monitor, including hard disk read IOPS, hard disk read traffic, hard disk write IOPS, hard disk write traffic, hard disk I/O waiting time, hard disk I/O service time, and hard disk I/O busy ratio.
Tag	Tags are used to identify cloud resources, which can be used to classify and search for CBS resources.

Previous Generation Products

HDD Cloud Disk

The HDD cloud disk is the previous generation cloud disk type, suitable for business scenarios with low I/O loads where data is infrequently accessed. It uses magnetic media as the storage medium and employs a distributed mechanism with three replicas to achieve highly reliable data storage.

Performance Metrics

Performance Metrics	HDD Cloud Disk
Random IOPS	Provide specifications from 10 GiB to 16,000 GiB and hundreds of random IOPS.
Throughput (MB/s)	Support 50MB/s I/O throughput performance.
Latency	-

Use Cases

For scenarios with relatively infrequent requests, such as system logs, enterprise work files, data repositories, small blogs, and BBS, we recommend using an HDD cloud disk to reduce costs.

Purchase Guide
Billing Overview

CBS Billing Overview

The billing mode of CBS is described as follows:

Billing mode	Pay-as-you-go
Payment Method	Settled every hour
Billing Unit	USD/GB * hour
Minimum Service Time	Settle by hour, purchase at any time and release at any time.
Application Scenario	Applicable to scenarios where the demand for devices fluctuates greatly in an instant, such as e-commerce flash sales.

> Note:
>
> The billing feature is not supported by cloud disks in Cloud Dedicated Cluster (CDC).

Billing Overview of Snapshot

- Snapshot charging is based on the total number of snapshots for users, and pushed to billing according to the actual total space occupied by snapshots per hour.
- The snapshot charging method is a post-paid method. It needs to be opened (it will be judged at the first push, and if it is not opened, the billing opening API will be called for opening). This is not perceived by the user), but there is no need to freeze and unfreeze.
- The lifecycle of snapshot billing connection billing management. It provides 3 APIs: service suspension, correction and destruction.
- Regular snapshot. The created snapshot will have an expiration deletion time, and after expiry, it will be automatically deleted -> billing is required.
- Snapshot created when a full machine image is created -> billing is required.

Measurement Method

- Take User+Region+Total Snapshot Capacity as the measurement unit.
- Total capacity: Count the total capacity of all snapshots of a user in a certain region based on the

actual size of snapshots.

- Total regional capacity monitoring: The total snapshot capacity of each user region every 10 minutes needs to be reported to the monitoring side.

Billing Mode

- Billing interval: deducted by hour
- Display: resource bills and cost details

Arrears Policy

- Arrears Definition: Account remaining amount <= 0.
- Freezing Policy:
    - Feature Freezing

      Once the user account is in arrears, all user snapshot operation features (create, rollback, cross-regional replication, create cloud disk, create image, periodic snapshot creation/modification/associated cloud disk/availability, and no longer create snapshot when regular snapshot arrives at the point) are stopped except for Terminate.

Purchase Guide

Purchase Channels for Cloud Disks

Two methods are provided for purchasing cloud disks: Purchase via the console or API.

Direct Purchase via Console

1. Log in to the CBS console, and click Create.
2. Set the cloud disk type and capacity.
3. Confirm to enable.

Purchase by Snapshot from Console

If you want the new disk to contain the data from a snapshot of a data disk by default, you can create a cloud disk from a snapshot on the Snapshot List page, or specify the target snapshot by configuring the parameter Snapshot when purchasing the disk separately through the console.

1. Log in to the Snapshot List page.
2. In the row of the target snapshot, select More > Create Cloud Disk.
3. Set the cloud disk type and capacity.

   > Note:
   >
   > When you create a cloud disk using a snapshot, the capacity cannot be smaller than the snapshot size. If you do not specify the capacity of a cloud disk, the capacity defaults to the same as the snapshot size.

4. Confirm to enable.

Purchasing Along with a CVM

- You can purchase a non-elastic cloud disk by setting the System Disk as a cloud disk when purchasing a CVM.
- You can purchase an elastic cloud disk by setting the Data Disk parameter when purchasing a CVM.

Purchase via API

You can use the CreateDisks API to create a cloud disk. For specific operations, see CreateDisks in the API documentation.

Viewing Bills

Operation Steps

1. Log in to the console and choose Cloud Products > Management & Audit > Billing Management to enter the Billing Management page.
2. In the left sidebar, choose Bill Management > Bill Details > Bill by Instance.

Getting Started
Step 1: Creating Cloud Disks

This document takes creating an empty cloud disk cbs-test in City1 as an example to help you understand how to quickly create the cloud disk.

Prerequisites

- Before creating a cloud disk, you need to register a cloud platform account and complete the qualification review.
- Make sure that you have an available CVM in the region and availability zone where you created the cloud disk (City1 in this example).

Operation Steps

1. Log in to the CBS console.
2. Select City1, and click Create.
3. Set the following parameters.
    - AZ: Select City1.
    - Cloud Disk Type: Select SSD Cloud Disk.
    - Capacity: Set to 20 GiB.
    - Snapshot: Whether to create cloud disk from a snapshot.
    - Disk Name: Enter cbs-test .
    - Purchase Quantity: Select 1.
4. Click OK.

Subsequent Operations

After the cloud disk is created, you need to attach it on a CVM in the same availability zone to be used as a data disk.

Step 2: Attaching Cloud Disks

This document will take an example of attaching an empty elastic cloud disk cbs-test to a CVM in the City1 region to help you quickly understand how to attach a cloud disk. For more information about attaching cloud disks, please see Attaching Cloud Disks.

> Note:
>
> Cloud disks can only be attached to the CVM in the same availability zone.

Prerequisites

- Successfully created the cbs-test cloud disk.
- Make sure that you have an available CVM in the region and availability zone where you created the cloud disk (City1 in this example). For information on how to purchase and start a CVM, please see Create an Instance.

Connecting to a CVM Instance

1. Log in to the CBS console.

2. On the Cloud Block Storage page, locate the cbs-test cloud disk, and select More > Attach.

3. In the pop-up window, select a CVM instance to which the cloud disk will be attached, and click OK.

   > Note:
   >
   > You can check Release on Instance Termination according to the actual situation.

   Return to the CBS list page, and at this time, the cloud disk status will be Mounting, indicating that cloud disk is in the process of being attached to the CVM. When the CBS status is In use, it indicates that the attachment to the CVM was successful.

Subsequent Operations

After the cloud disk is attached, it is used as the data disk of the CVM and is offline by default. You need to perform initialization operations such as formatting, partitioning, and creating a file system on the data disk. For specific operations, please see Initializing Cloud Disks.

Step 3: Initializing Cloud Disks

Scenarios

This document takes initializing a newly attached empty cloud disk to the CVM, creating a file system, and writing a file named test.txt to it as an example to help you understand how to simply use cloud disks. For more information on initializing cloud disks, see Initialization Scenarios.

Prerequisites

The cloud disk has been successfully attached; that is, the cloud disk is in the state of In use.

> Note:
>
> The Windows and Linux CVM operating systems used in this document are as follows:
>
> - Windows Server 2008
> - CentOS 7.5

Formatting, Creating File System and Writing Files (Windows)

1. Log in to the Windows instance as an administrator.

2. Choose CVM Management > Storage > Disk Management to enter the Disk Management interface.

3. (Optional) Right-click the target empty disk and select Online.
   When the disk's state changes to Not Initialized, it means that the online process is complete.

4. (Optional) Right-click the online disk and select Initialize Disk. In the pop-up window, select Master Boot Record and click OK.

   > Note:
   >
   > The maximum disk capacity supported by the MBR (Main Boot Record) format partition is 2TB, and that supported by the GPT (GUID Partition Table) partition table is 18EB. If you need to use a disk capacity greater than 2TB, please use GPT partitioning method.
   >
   > After the disk is put into use, if you switch the disk partition form, the original data on the disk

> will be cleared. Therefore, please select the disk partition form carefully when you initialize the disk.

5. Right-click the target disk and select New Simple Volume.

6. Input the simple volume size and click Next.

7. Select a drive letter or drive path and click Next. This document takes assigning drive letter E as an example.

8. Select the file system, format the partition, and click Next.

9. Click Complete.
   The target disk shows as formatting. Wait a moment for the system to complete the initialization. When the volume status is Healthy, it means the disk has been successfully initialized. After successful initialization, you can see the new partition's data disk in the PC interface.

10. Enter the newly partitioned data disk, create a new file test.txt , input the content you need, and choose File > Save.

Formatting, Creating File System and Writing Files (Linux)

> Note:
>
> - This document uses the EXT4 file system as an example.
> - After a Linux CVM is restarted or started up, the data disk will not be automatically attached. For details, see Formatting and Attaching Data Disks.

1. Run the following command to view the disk name connected to the instance.

   ```
   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.
   This document uses /dev/vdb as an example for the disk name connected to the instance. As shown below:

```
[root@VM_0_81_centos ~]# fdisk -l

Disk /dev/vda: 53.7 GB, 53687091200 bytes, 104857600 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x0009ac89

   Device Boot      Start         End      Blocks   Id  System
/dev/vda1   *        2048   104857566    52427759+  83  Linux

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

2. Run the following command to format this disk.

   mkfs.ext4 /dev/vdb

3. Run the following command to mount the disk to the /data mount point.

   mount /dev/vdb /data

4. Run the following commands in sequence to enter this disk and create a new file test.txt .

   cd /data
   vi test.txt

5. Press i to enter the editing and input This is my first test. Press ESC to exit the editing and input :wq to save and exit the file.

6. Run the ls command, and you can see that the test.txt file has been written to the disk.

Operation Guide
Operation Overview

When using CBS, you may encounter issues in managing cloud disks, managing snapshots, and monitoring and alarms. This document describes common operations for you to use CBS.

Cloud Disk

- Creating Cloud Disks
- Attaching Cloud Disks
- Initializing Cloud Disks
    - Initialization Scenarios
    - Initializing Cloud Disks (< 2 TB)
    - Initializing Cloud Disks (≥ 2 TB)
- Expanding Cloud Disks
    - Expansion Scenarios
    - Expanding Cloud Disks
    - Expanding Partitions and File Systems (Windows)
    - Expanding Partitions and File Systems (Linux)
- Detaching Cloud Disks
- Terminating Cloud Disks
- Adjusting cloud disk performance

Snapshot

- Creating Snapshots
- Rolling Back Data from Snapshots
- Creating Cloud Disks Using Snapshots
- Scheduled Snapshot
- Cross-Region Snapshot Replication
- Deleting Snapshots
- Sharing Snapshots

Snapshot Group

- Creating Snapshot Groups
- Rolling Back Data from Snapshot Groups
- Creating Cloud Disks Using Snapshot Groups
- Creating Custom Images Using Snapshot Groups
- Deleting Snapshot Groups

Monitoring and Alarms

Monitoring and Alarming

Use Limits

Use Limits on Cloud Disk

Limit Type	Restrictions
Elastic cloud disk capability	The data disks purchased with the CVM are all elastic cloud disks, which can be unmounted from the CVM and remounted.
Cloud disk performance	I/O performance takes effect at the same time. For example, for a 1 TB SSD, the maximum random IOPS can reach 26,000, which means that both read and write IOPS can reach this value. At the same time, due to several performance limitations, I/O with a block size of 4 KB/8 KB can reach the maximum value of IOPS in this example, but I/O with a block size of 16 KB cannot reach the maximum value of IOPS (the throughput has reached the limit of 260 MB/s).
Max attachable elastic cloud disks per CVM	Up to 32 disks.
Max cloud disks that can be created at a time	Up to 50 disks.
Restrictions for cloud disk mounting on CVM	The CVM and cloud disk must be in the same availability zone.

Use Limits on Snapshot

Limit Type	Restrictions
Limits on the creation types of cloud disks by snapshot	Only the data disk snapshot can be used to create an elastic cloud disk, and the system disk snapshot can only be used to create a custom image.
Limits on the creation sizes of cloud disks by snapshot	The capacity of the cloud disk created using a snapshot should be greater than or equal to that of the cloud disk from which the snapshot originated.

Limit Type	Restrictions
Limits on snapshot rollback	A snapshot can only be rolled back to the source cloud disk from which the snapshot was created. If you want to create a cloud disk using an existing snapshot, see Creating Cloud Disks Using Snapshots.
Total snapshot capacity	No limit.
Snapshot quota per cloud disk	1,000 snapshots.

Use Limits on Scheduled Snapshot Policies

Limit Type	Restrictions
Limits on cloud disks that can be bound to a scheduled snapshot policy	Only cloud disks in the same region as the scheduled snapshot policy can be bound. A maximum of 200 cloud disks can be bound.

Managing Cloud Disks
Creating Cloud Disks

Scenarios

You can create a cloud disk and connect it to any CVM in the same availability zone. Through block storage device mapping, the cloud disk is recognized and used by the CVM. The cloud disk can achieve its maximum performance without preheating after being created.
You can choose to create different types of cloud disks based on your business needs. For more details about cloud disk types, see Cloud Disk Type.

Operation Steps

Creating Cloud Disks Using the Console

1. Log in to the CBS console.
2. Select a region and click Create.
3. In the Purchase Data Disk dialog box, set the following parameters:

Parameter	Parameter Description
AZ	This parameter is required. The availability zone where a cloud disk is located cannot be changed after a cloud disk is created.
Cloud Disk Type	This parameter is required. Cloud disk types include: - Premium Cloud Storage - SSD Cloud Disk
Capacity	This parameter is required. When you create a cloud disk using a snapshot, the capacity cannot be smaller than the snapshot size. If you do not specify the capacity of a cloud disk, the capacity defaults to the same as the snapshot size.
Snapshot	This parameter is optional. When a cloud disk is created with a snapshot, check Create Cloud Disk with Snapshot and select the snapshot to be used. - The capacity of a cloud disk created with a snapshot is equal to the snapshot size by default, and you can increase the capacity to a value larger than the default value. - When a cloud disk is created with a snapshot, the disk type is consistent with the type of the source cloud disk by default. You can adjust the cloud disk type.

Parameter	Parameter Description
Disk Name	This parameter is optional. The name of a cloud disk supports up to 64 characters. After the cloud disk is created, it can be modified. - Create a single cloud disk: The disk name is the name of the cloud disk. - Batch create cloud disk: When multiple cloud disks is created at a time, the disk name will be the prefix of the cloud disk name, and the final cloud disk name is Disk Name_Digit: from Hard Disk Name_0 to Hard Disk Name_49.
Scheduled Snapshots	This parameter is optional. When you create a cloud disk, you can associate it with the regular snapshot policy to regularly manage the snapshot of the cloud disk.
Project	This parameter is optional. When you create a cloud disk, you can set the project to which it belongs.
Tag	This parameter is optional. You can bind tags to cloud disks when you create them. Tags are used to identify cloud resources, and you can use tags to classify and search for cloud resources.
Encrypt	This parameter is optional. If the KMS service is enabled and authorized for use, this item supports enabling cloud disks encryption. The encryption algorithm is AES256, and the sources of supported tokens are as follows: - Use the token created by KMS for Cloud Services by default - Use the existing custom token
Purchase Quantity	This parameter is optional. Quantity is 1 by default, meaning that only one cloud disk will be created. Currently, a maximum of 50 cloud disks can be created in batches.

4. Click OK.

   You can view the created cloud disk on the Cloud Disk page. The newly created elastic cloud disk is at To be mounted status. To mount the cloud disk to the CVM in the same availability zone, see Attaching Cloud Disks.

Creating Cloud Disks Using Snapshots

If you want to create a new cloud disk containing relevant data directly, you can select Create Cloud Disk from Snapshot.

Creating Cloud Disks Using the API

You can use the CreateDisks API to create a cloud disk

Attaching Cloud Disks

You can attach an elastic cloud disk (used as a data disk of the CVM) to any CVM in the same availability zone. Each CVM supports up to 20 data disks. You can attach a cloud disk by using the following methods:

- When starting a new CVM, specify the corresponding custom image and data disk snapshot.
  After automatic attachment, there is no need to perform initialization operations such as partitioning or formatting, and read-write operations can be directly performed on the data disk.
- For the cloud disks purchased separately, manually mount the elastic cloud disks to an existing CVM instance in the same availability zone through the console or API.
- After manually mounting, you'll need to perform initialization operations such as partitioning and formatting the disk. For details, see Initializing a Cloud Drive (less than 2TB) or Initializing a Cloud Drive (2TB or greater).
- Cloud disks created from snapshots
  If the cloud disk capacity equals the snapshot capacity, no initialization operations such as partitioning or formatting are required after manual attachment, and read-write operations can be directly performed on the data disk.
  If the cloud disk capacity is greater than the snapshot capacity, you need to extend the file system or convert the partition format.

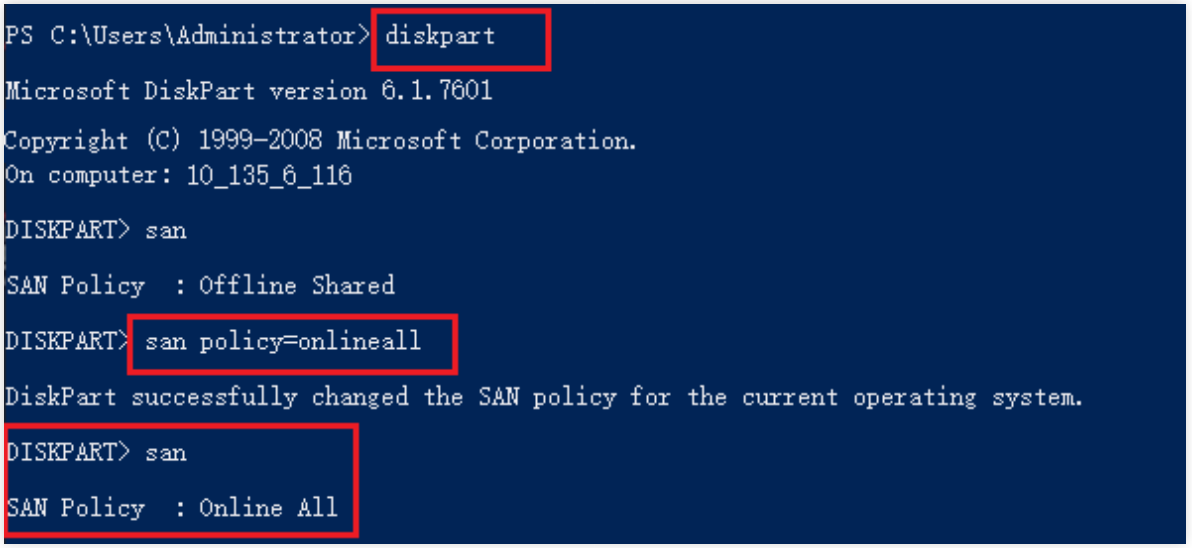
> Note:
>
> Some Linux CVMs may fail to recognize elastic cloud disks. You can first enable the hot-swappable disk feature in the CVM. For details, see Enabling Hot-Swappable Disk Feature.

Automatic Attachment

Attaching Data Disks (Windows)

To realize automatic attachment of a cloud disk created from the corresponding data disk snapshot when starting a Windows CVM instance, the specified custom image and data disk snapshot should meet the following requirements:

- Before you create a snapshot, the data disk must be formatted into the ntfs or fat32 format.
- The SAN policy in the custom image is onlineAll .
  Note:
  The Windows public image currently provided has been set by default, but it is still recommended that you check this configuration before you create a custom image. The checking method is as follows:



Attaching Data Disks (Linux)

To realize automatic attachment of a cloud disk created from the corresponding data disk snapshot when starting a Linux CVM instance, the specified custom image and data disk snapshot should meet the following requirements:

- Before you create a snapshot, the data disk must be formatted, that is, it has been successfully mounted on the source CVM.
- Before you create a custom image for the system disk, you need to add the following commands in the /etc/rc.local file to write the mount point of the data disk into the file.

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

> Note:
>
> - <mount-point> should be set to the mount point of the file system, for example /mydata .

- <device-id> should be set to the actual file partition location. For example, fill in /dev/vdb when there is a file system but no partition, and fill in /dev/vdb1 when there is a partition and a file system.

Manual Mounting

Attaching Cloud Disks Using the Console

1. Log in to the CBS console.
2. On the cloud disk list page, you can attach cloud disks using the following methods:
   a. Click More > Attach in the row of the cloud disk in the state of To be mounted.
   b. Select the cloud disks in the state of To be mounted, and click Attach above the cloud disk list to perform batch attachment.
3. Select the target CVM in the pop-up box, and click OK.
4. Refresh the cloud disk list.
   If the cloud disk's state changes to In use, it indicates that the attachment is successful.
5. Depending on the cloud disk's state, you need to select and perform the corresponding subsequent operations to make it available.

Creation Model	Cloud Disk Capacity	Subsequent Operations
Creating Directly	Cloud disk capacity < 2 TB	Initializing Cloud Disks (<2 TB)
Creating Directly	Cloud disk capacity ≥ 2 TB	Initializing Cloud Disks (≥ 2 TB)
Creating from Snapshots	Cloud disk capacity = Snapshot capacity	No subsequent operation is required, and it can be used directly after it is mounted.
Creating from Snapshots	Snapshot capacity < Cloud disk capacity ≤ 2 TB or 2 TB < Snapshot capacity < Cloud disk capacity	- Mount to Windows CVM: scale out the partition and file system (Windows) - Mount to Linux CVM: scale out the partition and file system (Linux)
Creating from Snapshots	Snapshot capacity ≤ 2 TB < Cloud disk capacity	If MBR partition is used in the snapshot: See Initializing Cloud Disks (≥ 2 TB) for repartition using GPT. This operation will delete existing data. If the GPT partition is used in the snapshot: - Mount to Windows CVM: scale out the partition and file system (Windows) - Mount to Linux CVM: scale out the partition and file system (Linux)

Attaching Cloud Disks Using the API

You can use the AttachDisks API to create a snapshot. For details, see API Documentation > AttachDisks.

Enabling the Disk Hot Swapping Feature

All currently provided images support attaching and detaching elastic cloud disks. Before detaching a

cloud disk, you need to perform umount (Linux) or offline (Windows) operations. Otherwise, the CVM may fail to recognize the elastic cloud disk when it is attached again.
However, if you have purchased a CVM with the following operating systems before and plan to attach an elastic cloud disk to it, it is recommended that you first add the relevant driver to the CVM to enable the hot-plug feature.

1. Log in to the Linux CVM as root user.

2. Execute the following command to add a driver.

   ```
   modprobe acpiphp
   ```

3. (Optional) According to different operating systems, select the corresponding operation method to set the `acpiphp` module as automatically load upon startup:

   - Centos5 series

     a. Execute the following command to create and open the `acpiphp.modules` file.

     ```
     vi /etc/sysconfig/modules/acpiphp.modules
     ```

     b. Add the following to the file and save it.

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

     c. Execute the following command to add the executable permission.

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

   - Debian 6 series, Ubuntu 10.04 series

     a. Execute the following command to modify the file.

     ```
     vi /etc/modules
     ```

     b. Add the following to the file and save it.

acpiphp

- openSUSE 12.3 series

  a. Execute the following command to modify the file.

  vi /etc/sysconfig/kernel

  b. Add the following to the file and save it.

  MODULES_LOADED_ON_BOOT="acpiphp"

Initializing Cloud Disks
Initialization Scenarios

After a cloud disk is manually attached, the cloud disk is used as a data disk for the CVM and is in the offline state by default. You need to perform initialization operations on the data disk, such as formatting, partitioning, and creating a file system. Choose the initialization method based on your actual usage scenario:

- If the whole hard disk is only presented as an independent partition (that is, there are no multiple logical disks, such as disk D /vdb1 and disk E /vdb2), it is recommended that you do not use the partition but build the file system directly on the raw device.
- 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 build a file system on the partition.

Common disk partition forms include the Main Boot Record (MBR) and Global Partition Table (GPT). After the disk is put into use, switch the disk partition form to clear the original data on the disk. Therefore, a reasonable partition form should be selected according to actual needs.

The following table briefly describes the two partition formats.

Partition Form	Maximum Disk Capacity Supported	Quantities of Partitions Supported	Partition Tools
MBR	2TB	- 4 main partitions - 3 main partitions and 1 scale-out partition	Windows operating system: disk management Linux operating system: - fdisk tool - parted tool
GPT	18EB The current maximum cloud disk capacity supported is 16 TB.	Do not limit the partition quantities.	Windows operating system: disk management Linux operating system: parted tool

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

- Disk capacity < 2 TB:
    - Initializing Cloud Disks (Windows)
    - Initializing Cloud Disks (Linux)
- Disk capacity ≥ 2 TB:
    - Initializing Cloud Disks (Windows)
    - Initializing Cloud Disks (Linux)

Initializing Cloud Disks (< 2 TB)

Scenarios

This document takes the capacity of cloud disks less than 2 TB as an example to guide the initialization of cloud disks. For more information about the disk initialization scenario, please see Initialization Scenarios.

Prerequisites

> Note:
>
> - Formatting the data disk will empty 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.

The cloud disk has been attached to a CVM.

Operation Steps

Initializing Cloud Disks (Windows)

> Note:
>
> This document takes the Windows Server 2008 operating system as an example. Different operating systems may have different formatting operations, so this document is for reference only.

1. Log in to the Windows CVM instance.
2. Click Start on the CVM desktop.
3. Right-click Computer in the Start menu, and select Manage.
4. In the left sidebar, choose Storage > Disk Management to enter the disk management page.

   > Note:
   >
   > If the new disk is offline, you need to perform step 5 to bring it online before performing step 6 for initialization. Otherwise, directly perform step 6 for initialization.

5. The disk list appears in the right pane. Right-click the Disk 1 area, and select Online in the menu list to connect. After online, disk 1 changes from offline to Not initialized.

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

7. Display the disk to be initialized in the initialize disk dialog box, select MBR or GPT, and click OK.

   > Note:
   >
   > After a disk is put into use, switching the disk partition scheme will result in the deletion of all existing data on the disk. Therefore, please select the appropriate partition scheme based on your actual needs.

8. Right-click an unallocated region on the disk and select New Simple Volume.

9. The Create simple volume wizard dialog box is displayed. Click Next as prompted on the page.

10. Specify the volume size according to the actual situation. The default is the maximum value, and click Next.

11. Assign a drive number and click Next.

12. Select Format this volume with the following settings, set parameters according to the actual situation, format the new partition, and click Next to complete the partition creation.

13. Click Complete to finish the wizard. You need to wait a moment for the system to complete initialization. When the volume status is Healthy, it means that the disk has been successfully initialized.

14. After successful initialization, enter the computer interface to check the new disk.

Initializing Cloud Disks (Linux)

Select the initialization method according to your actual use cases:

- If the entire hard disk is only presented as an independent partition (i.e., there are no multiple logical disks such as vdb1 and vdb2), it is highly recommended that you do not use a partition and directly Creating file systems on bare devices.
- If the entire hard disk needs to be presented as multiple logical partitions (that is, there are multiple logical disks), you need to partition it first, and then build a file system on the partitions.

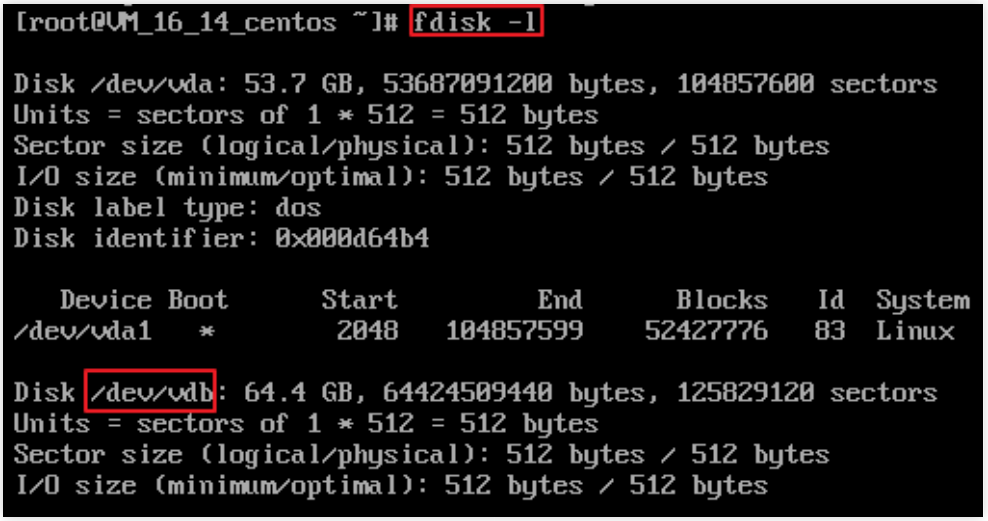
Creating File Systems on Bare Devices

1. Log in to the Linux CVM.

2. Execute the following command as root user to check the disk name.

   fdisk -l

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



3. Execute 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. Take setting the file system to EXT4 as an example:

   mkfs -t ext4 /dev/vdb

> Note:
>
> Formatting takes some time. Observe the system's running status and do not exit.

4. Execute 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. Execute 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. Execute the following command to check the mounting result.

   df -TH

   > Note:
   >
   > If there is no need to set automatic disk mounting at startup, you can skip the subsequent steps.

7. Confirm the mounting mode and obtain corresponding information.
   You can use the soft link of elastic cloud disks, the UUID (universally unique identifier) of the file system, or the device name to mount the disk automatically according to business requirements. The relevant instructions and ways of obtaining information are as follows:

Mounting Mode	Advantages and Disadvantages	Ways to Obtain Information
Using a Soft Link from an Elastic Cloud Disk (Recommended)	Advantages: The soft link of each elastic cloud disk is fixed and unique, remaining unchanged even during operations such as unmounting and remounting, or formatting partitions. Disadvantages: Only elastic cloud disks has soft link. Elastic cloud disks do not recognize or detect partition formatting operations.	Execute the following command to check the soft link of elastic cloud disks. ls -l /dev/disk/by-id
Using the UUID of the File System	Automatic mount settings may be invalidated due to file system UUID changes. For example, the UUID of a file system will change after it is reformatted.	Execute the following command to check the UUID of the file system. blkid /dev/vdb

Mounting Mode	Advantages and Disadvantages	Ways to Obtain Information
Using the Name of Devices	The automatic mount setting may be invalidated due to the change of device name. For example, when the elastic cloud disks on the CVM are unmounted and remounted during data migration, the name may change when the operating system recognizes the file system again.	Execute the following command to check the device name. fdisk -l

8. Execute the following command to back up the /etc/fstab file. Take backup to the /home directory as an example:

   cp -r /etc/fstab /home

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

   vi /etc/fstab

10. Press i to enter the 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> <Installation options of file system> <File system dump frequency> <File system check sequence at startup>

- (Recommended) Take automatic mounting by the soft link from elastic cloud disks as an example. Based on the previous examples, the following is added:

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

- Take automatic mounting by the disk partition's UUID as an example. Based on the previous examples, the following is added:

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

- Take automatic mounting using the name of devices as an example. Based on the previous examples, the following is added:

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

12. Press Esc, enter :wq, and press Enter.
    Save the settings and exit the editor.

13. Execute the following command to check whether the /etc/fstab file is written successfully.

```
mount -a
```

If the operation succeeds, it means that the file is written successfully, and the newly created file system will be automatically mounted when the operating system starts.

Creating a File System on the Partition

> Note:
>
> This document uses the fdisk partition tool in the CentOS 7.5 operating system to configure the /dev/vdc data disk 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.

2. Execute the following command as root user to check the disk name.

   ```
   fdisk -l
   ```

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

```
[root@VM_16_14_centos ~]# fdisk -l

Disk /dev/vda: 53.7 GB, 53687091200 bytes, 104857600 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x000d64b4

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

Disk /dev/vdb: 64.4 GB, 64424509440 bytes, 125829120 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
```

3. Execute the following command to enter the fdisk partition tool and start partitioning the new data disk.

   fdisk <Create data disk>

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

   fdisk /dev/vdb

   The output is similar to the following figure:

```
[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 0xcb7892a8.

Command (m for help):
```

4. Enter n and press Enter to create a new partition.
   The output is similar to the following figure:

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

   Indicates that the disk has two partition types:

- P indicates the main partition.
- E indicates the extended partition.

5. Take the creation of a primary partition as an example. Enter p and press Enter to start creating a primary partition.
   The output is similar to the following figure:

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

   Partition number indicates the main partition number, ranging from 1 to 4.

6. Take the selection of partition number 1 as an example, enter the main partition number 1 and press Enter.
   The output is similar to the following figure:

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

   First sector indicates the initial magnetic column area. You can select 2048 - 20971519, and the default is 2048.

7. Take the selection of the default initial cylinder number 2048 as an example, press Enter.
   The output is similar to the following figure:

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

   Last sector indicates the cut-off magnetic cylinder area. You can select 2048-20971519, and the default is 20971519.

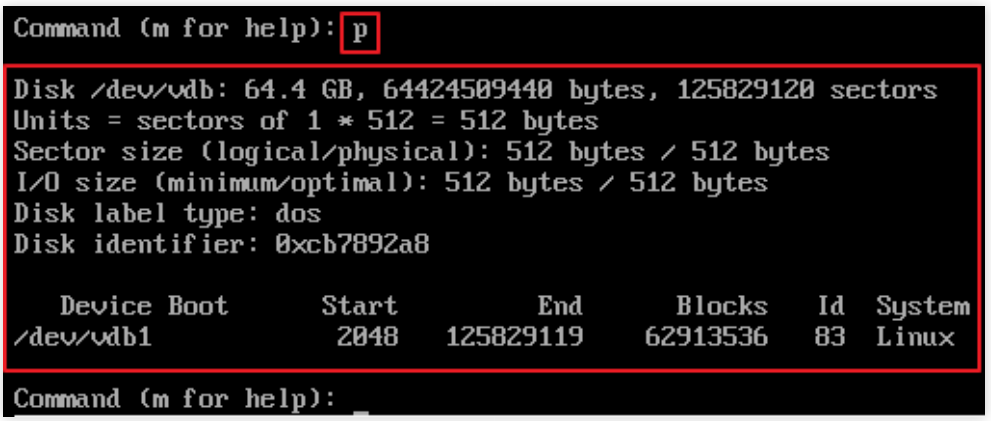
8. Take the selection of default cut-off cylinder number 20971519 as an example, press Enter.
   The output is similar to the following figure:

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

Command (m for help):
```

   Indicates that the partition is completed, i.e. a new partition has been created for the 60 GB data disk.

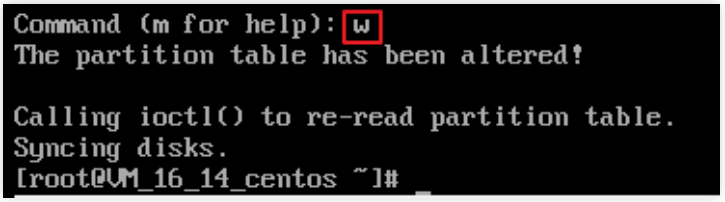
9. Enter p and press Enter to check the details of the new partition.
   The output is similar to the following figure:



It indicates the details of the new partition /dev/vdb1 .

> Note:
>
> If the above partition operation is incorrect, please enter q to exit the fdisk partition tool. The previous partition results will not be retained.

10. Enter w and press Enter to write the partition results into the partition table.
    The output is similar to that shown in the following figure, indicating that the partition has been created.



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

    partprobe

12. Execute the following command to set the file system of the new partition in the format 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. Take setting the file system as EXT4 as an example:

mkfs -t ext4 /dev/vdb1

The output is similar to the following figure:

```
[root@VM_16_14_centos ~]# mkfs -t ext4 /dev/vdb1
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
3932160 inodes, 15728384 blocks
786419 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2164260864
480 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
        32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
        4096000, 7962624, 11239424

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

Formatting takes some time. Observe the system's running status and do not exit.

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

mkdir <Mount point>

Take the new mount point /data/newpart as an example:

mkdir /data/newpart

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

mount /dev/vdb1 <Mount point>

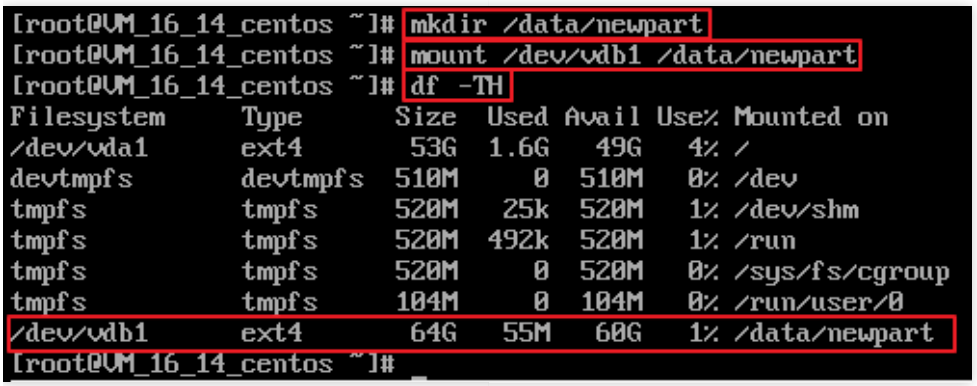
Take the new mount point /data/newpart as an example:

mount /dev/vdb1 /data/newpart

15. Execute the following command to check the mounting result.

df -TH

The output is similar to the following figure:



It indicates that the new partition /dev/vdb1 has been mounted to /data/newpart .

> Note:
>
> If there is no need to set automatic disk mounting at startup, you can skip the subsequent steps.

16. Confirm the mounting mode and obtain corresponding information.
    You can use the soft link of elastic cloud disks, the UUID (universally unique identifier) of the file system, or the device name to mount the disk automatically according to business requirements. The relevant instructions and ways of obtaining information are as follows:

Mounting Mode	Advantages and Disadvantages	Ways to Obtain Information
Using a soft link from an elastic cloud disks (recommended)	Advantages: The soft link of each elastic cloud disk is fixed and unique, remaining unchanged even during operations such as unmounting and remounting, or formatting partitions. Disadvantages: Only elastic cloud disks has soft link. Elastic cloud disks do not recognize or detect partition formatting operations.	Execute the following command to check the soft link of elastic cloud disks. ls -l /dev/disk/by-id
Using the UUID of the File System	Automatic mount settings may be invalidated due to file system UUID changes. For example, the UUID of a file system will change after it is reformatted.	Execute the following command to check the UUID of the file system. blkid /dev/vdb1

Mounting Mode	Advantages and Disadvantages	Ways to Obtain Information
Using the Name of Devices	The automatic mount setting may be invalidated due to the change of device name. For example, when the elastic cloud disks on the CVM are unmounted and remounted during data migration, the name may change when the operating system recognizes the file system again.	Execute the following command to check the device name. fdisk -l

17. Execute the following command to back up the /etc/fstab file . Take backup to the /home directory as an example:

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

18. Execute the following command to open the /etc/fstab file with the VI editor.

```
vi /etc/fstab
```

19. Press i to enter the 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> <Installation options of file system> <File sy
stem dump frequency> <File system check sequence at startup>
```

- (Recommended) Take automatic mounting by the soft link from elastic cloud disks as an example. Based on the previous examples, the following is added:

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

- Take automatic mounting by the disk partition's UUID as an example. Based on the previous examples, the following is added:

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

- Take automatic mounting using the name of devices as an example. Based on the previous examples, the following is added:

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

21. Press Esc, enter :wq, and press Enter.
    Save the settings and exit the editor.

22. Execute the following command to check whether the /etc/fstab file is written successfully.

```
mount -a
```

If the operation succeeds, it means that the file is written successfully, and the newly created file system will be automatically mounted when the operating system starts.

Initializing Cloud Disks (≥ 2 TB)

Scenarios

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 a maximum disk capacity of 2 TB. Therefore, when partitioning a disk with a capacity greater than 2 TB, use GPT as the partition format. For the Linux operating system, when GPT is selected as the partition format, the fdisk partitioning tool cannot be used, and you need to use the parted tool.

Prerequisites

> Note:
>
> - Formatting a data disk will erase all data. Make sure that the disk does not contain data or that important data has been backed up.
> - To prevent service exceptions, ensure that the Cloud Virtual Machine (CVM) has stopped providing external services before formatting.

The cloud disk has been attached to a CVM.

Operation Steps

Initializing Cloud Disks (Windows)

> 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. This document is for reference only.

1. Log in to the Windows CVM instance.

2. Click  on the desktop. Go to the server manager page.

3. In the left sidebar, click File and Storage Services.

4. In the left sidebar, select Volume > Disk.

   > Note:
   >
   > If the newly added disk is offline (as shown in the above figure), perform Step 5 to connect it, and then perform Step 6 for initialization. Otherwise, directly perform Step 6 for initialization.

5. The disk list appears in the right pane. Right-click the line where 1 is located, and select Online in the menu list to connect. After online, 1 changes from Offline status to Online.

6. Right-click the line where 1 is located and select Initialize in the menu list.

7. Click Yes as prompted.

8. After initialization, 1 changes from Unknown partition to GPT. Right-click the line where 1 is located and select New Simple Volume in the menu list.

9. The New Volume Wizard dialog box pops up. Click Next as prompted on the page.

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. Select Format this volume with the following settings, set parameters according to the actual situation, format the new partition, and click Next to complete the partition creation.

14. After the information is confirmed, click Create.

15. Wait for the completion of the volume creation and click Finish.
    Upon successful initialization, you can enter the This PC interface to view the new disk.

Initializing Cloud Disks (Linux)

Select the initialization method according to your actual use cases:

- If the entire hard disk is presented as a single independent partition (that is, there are no multiple logical disks like vdb1 and vdb2), it is highly recommended that you do not partition it and directly build a file system on the raw device.
- If the entire hard disk needs to be presented as multiple logical partitions (that is, there are multiple logical disks), you need to partition it first, and then build a file system on the partitions.

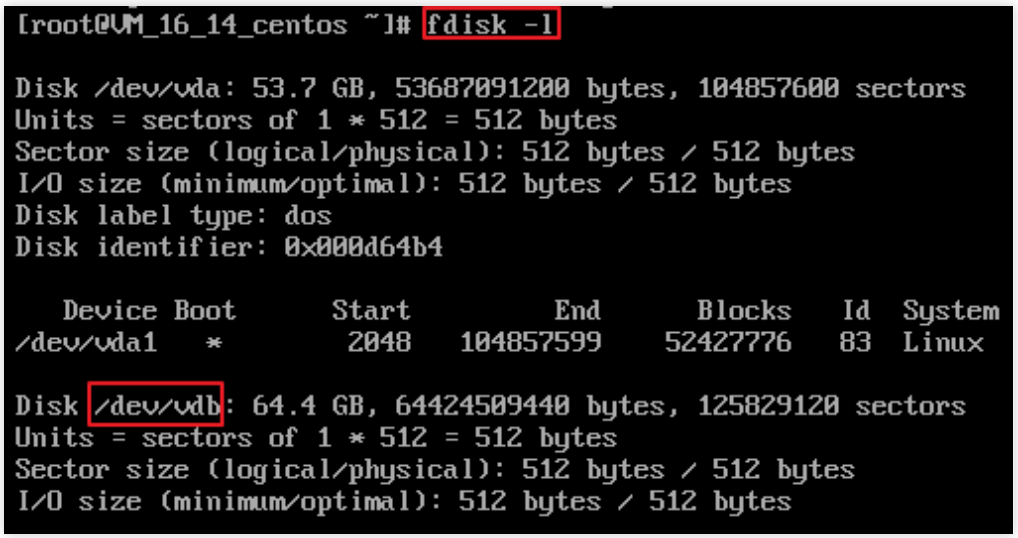
Creating File Systems on Bare Devices

1. Log in to the Linux CVM.

2. Execute the following command as root user to check the disk name.

   fdisk -l

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



3. Execute 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. Take setting the file system to EXT4 as an example:

   mkfs -t ext4 /dev/vdb

> Note:
>
> Formatting takes some time. Observe the system's running status and do not exit.

4. Execute 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. Execute 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. Execute the following command to check the mounting result.

   df -TH

> Note:
>
> If there is no need to set automatic disk mounting at startup, you can skip the subsequent steps.

7. Confirm the mounting mode and obtain corresponding information.
   You can use the soft link of elastic cloud disks, the UUID (universally unique identifier) of the file system, or the device name to mount the disk automatically according to business requirements. The relevant instructions and ways of obtaining information are as follows:

Mounting Mode	Advantages and Disadvantages	Ways to Obtain Information
Using a soft link from an elastic cloud disks (recommended)	Advantages: The soft link of each elastic cloud disk is fixed and unique, remaining unchanged even during operations such as unmounting and remounting, or formatting partitions. Disadvantages: Only elastic cloud disks has soft link. Elastic cloud disks do not recognize or detect partition formatting operations.	Execute the following command to check the soft link of elastic cloud disks. ls -l /dev/disk/by-id

Mounting Mode	Advantages and Disadvantages	Ways to Obtain Information
Using the UUID of the File System	Automatic mount settings may be invalidated due to file system UUID changes. For example, the UUID of a file system will change after it is reformatted.	Execute the following command to check the UUID of the file system. blkid /dev/vdb
Using the Name of Devices	The automatic mount setting may be invalidated due to the change of device name. For example, when the elastic cloud disks on the CVM are unmounted and remounted during data migration, the name may change when the operating system recognizes the file system again.	Execute the following command to check the device name. fdisk -l

8. Execute the following command to back up the /etc/fstab file. Take backup to the /home directory as an example:

   cp -r /etc/fstab /home

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

   vi /etc/fstab

10. Press i to enter the 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> <Installation options of file system> <File system dump frequency> <File system check sequence at startup>

- (Recommended) Take automatic mounting by the soft link from elastic cloud disks as an example. Based on the previous examples, the following is added:

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

- Take automatic mounting by the disk partition's UUID as an example. Based on the previous examples, the following is added:

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

- Take automatic mounting using the name of devices as an example. Based on the previous examples, the following is added:

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

12. Press Esc, enter :wq, and press Enter.
    Save the settings and exit the editor.
13. Execute the following command to check whether the /etc/fstab file is written successfully.

```
mount -a
```

If the operation succeeds, it means that the file is written successfully, and the newly created file system will be automatically mounted when the operating system starts.

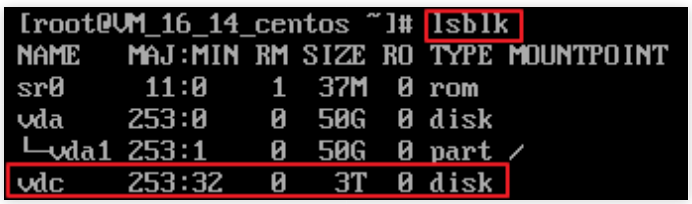
Creating a File System on the Partition

> 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.
2. Execute the following command as root user to check the disk name.

   ```
   lsblk
   ```

   The output is similar to the following figure, indicating that the current CVM has two disks. /dev/vda is the system disk, and /dev/vdc is the newly added data disk.



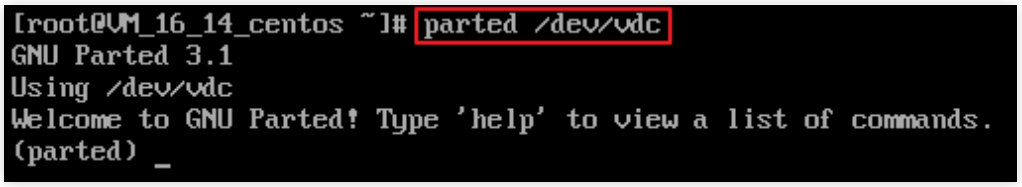
3. Execute the following command to enter the parted tool and start partitioning the new data disk.

   parted <Create data disk>

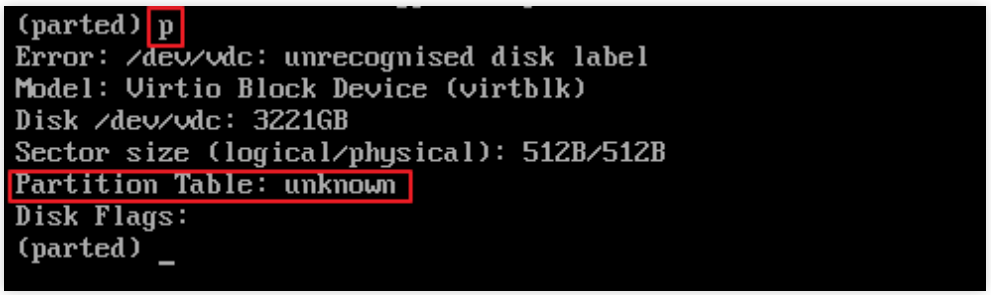
   Take the newly mounted data disk /dev/vdc as an example:

   parted /dev/vdc

   The output is similar to the following figure:



4. Enter p and press Enter to check the current disk partition form.
   The output is similar to the following figure:



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

5. Execute the following command to configure the disk partition form.

   mklabel <Disk partitioning method>

   When the disk capacity is greater than or equal to 2 TB, only GPT partitioning can be used:

   mklabel gpt

6. Enter p and press Enter to check whether the disk partition is set successfully.
   The output is similar to the following figure:

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

Number  Start  End  Size  File system  Name  Flags

(parted)
```

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

7. Enter unit s and press Enter to set the measurement unit of the disk as the cylinder.

8. Taking the creation of a partition for an entire disk as an example, enter mkpart opt 2048s 100% and press Enter.
   2048s indicates the initial capacity of the disk, and 100% indicates the ending capacity of the disk. It is for reference only. You can plan the number and capacity of disk partitions according to business needs.

9. Enter p and press Enter to check the details of the new partition.
   The output is similar to the following figure:

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

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

(parted) _
```

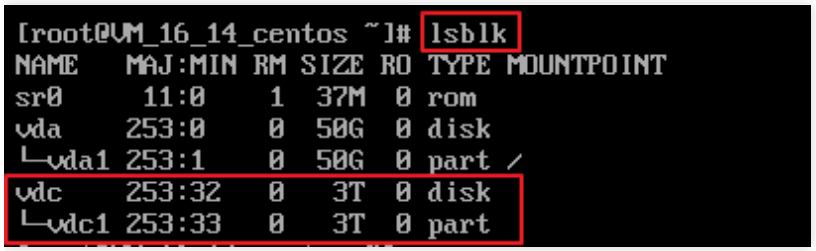
It indicates the details of the new partition /dev/vdc1 .

10. Enter q and press Enter to exit the parted tool.

11. Execute the following command to check the disk name.

lsblk

The output is similar to that shown in the following figure. At this time, you can see the new partition /dev/vdc1.



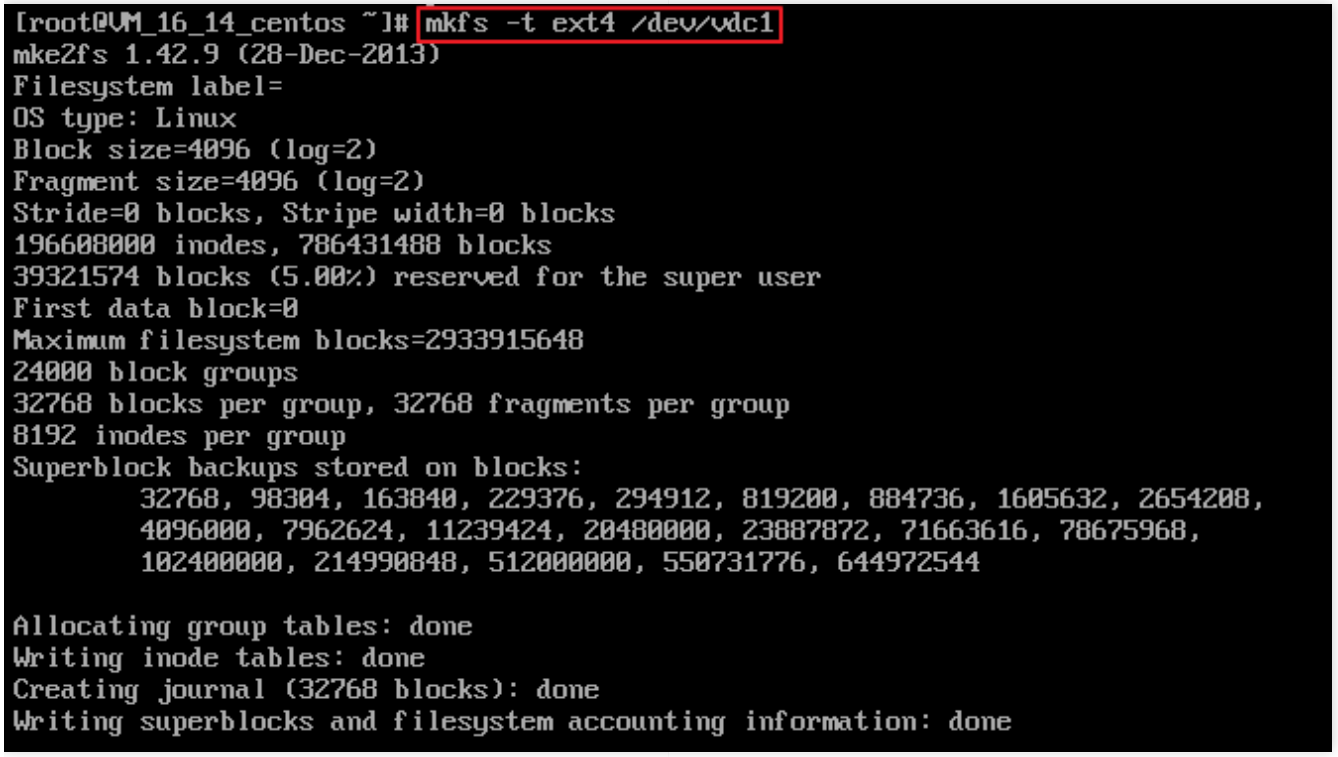
12. Execute the following command to set the file system of the new partition in the format required by the system.

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

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

mkfs -t ext4 /dev/vdc1

The output is similar to the following figure:



Formatting takes some time. Observe the system's running status and do not exit.

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

mkdir <Mount point>

Take the new mount point `/data/newpart2` as an example:

mkdir /data/newpart2

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

mount /dev/vdc1 <Mount point>

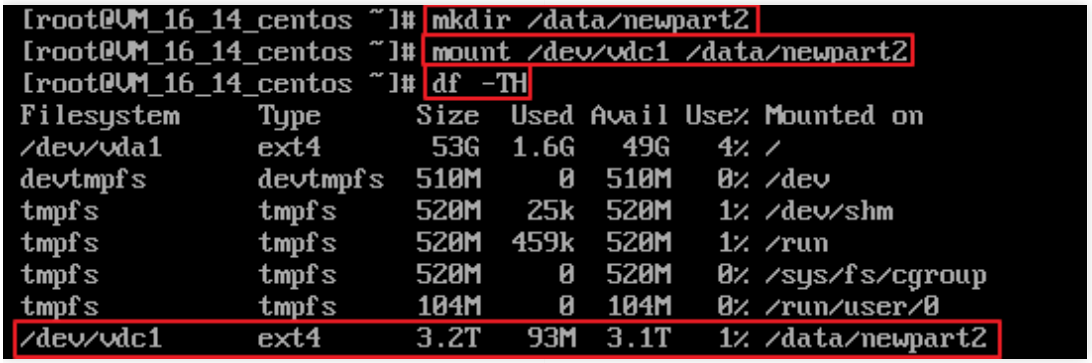
Take the new mount point `/data/newpart2` as an example:

mount /dev/vdc1 /data/newpart2

15. Execute the following command to check the mounting result.

df -TH

The output is similar to the following figure:



It indicates that the new partition /dev/vdc1 has been mounted to /data/newpart2 .

> Note:
>
> If there is no need to set automatic disk mounting at startup, you can skip the subsequent steps.

16. Confirm the mounting mode and obtain corresponding information.
    You can use the soft link of elastic cloud disks, the UUID (universally unique identifier) of the file system, or the device name to mount the disk automatically according to business requirements. The relevant instructions and ways of obtaining information are as follows:

Mounting Mode	Advantages and Disadvantages	Ways to Obtain Information
Using a soft link from an elastic cloud disks (recommended)	Advantages: The soft link of each elastic cloud disk is fixed and unique, remaining unchanged even during operations such as unmounting and remounting, or formatting partitions. Disadvantages: Only elastic cloud disks has soft link. Elastic cloud disks do not recognize or detect partition formatting operations.	Execute the following command to check the soft link of elastic cloud disks. ls -l /dev/disk/by-id
Using the UUID of the File System	Automatic mount settings may be invalidated due to file system UUID changes. For example, the UUID of a file system will change after it is reformatted.	Execute the following command to check the UUID of the file system. blkid /dev/vdc1
Using the Name of Devices	The automatic mount setting may be invalidated due to the change of device name. For example, when the elastic cloud disks on the CVM are unmounted and remounted during data migration, the name may change when the operating system recognizes the file	Execute the following command to check the device name. fdisk -l

Mounting Mode	Advantages and Disadvantages	Ways to Obtain Information
	system again.	

17. Execute the following command to back up the /etc/fstab file. Take backup to the /home directory as an example:

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

18. Execute the following command to open the /etc/fstab file with the VI editor.

```
vi /etc/fstab
```

19. Press i to enter the 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> <Installation options of file system> <File system dump frequency> <File system check sequence at startup>
```

- (Recommended) Take automatic mounting by the soft link from elastic cloud disks as an example. Based on the previous examples, the following is added:

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

- Take automatic mounting by the disk partition's UUID as an example. Based on the previous examples, the following is added:

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

- Take automatic mounting using the name of devices as an example. Based on the previous examples, the following is added:

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

21. Press Esc, enter :wq, and press Enter.
    Save the settings and exit the editor.

22. Execute the following command to check whether the /etc/fstab file is written successfully.

```
mount -a
```

If the operation succeeds, it means that the file is written successfully, and the newly created file system will be automatically mounted when the operating system starts.

Expanding Cloud Disks
Expansion Scenarios

The Expansion Type Is the Cloud Disk of the System Disk

When expanding a cloud disk of the system disk type, you can expand it using the following three methods.

- Expand it via the CVM console.
- Expand it via the CBS console.
- Expand it via API.

The Expansion Type Is the Cloud Disk of the Data Disk

When expanding a cloud disk of the data disk type, you can expand it using the following three methods.

- Expand it via the CVM console.
- Expand it via the CBS console.
- Expand it via API.

Depending on the "detachable" attribute of the CBS data disk, you can select different operation entries to expand the data disk.

- If the current CBS data disk can be unmounted, you can expand the capacity in the CBS console or by using ResizeDisk.
- If the current CBS data disk cannot be unmounted, you can expand the capacity in the CVM console or by using ResizeDisk.

> Note:
>
> If the maximum capacity of a cloud disk still cannot meet your business needs, you can use multiple elastic cloud disks to build an LVM logical volume.

After the data disk is scaled out, you need to perform subsequent operations before identifying and using instances:

Before Expansion	Subsequent Operations After Expansion
File system not created	- Disk capacity less than 2 TB: Initialize cloud disks (< 2 TB). - Disk capacity greater than or equal to 2 TB: Initialize cloud disks (≥ 2 TB).

Before Expansion	Subsequent Operations After Expansion
File system created	- Disk capacity < 2 TB: - The scaled-out disk is the cloud disk of Windows CVM: Scale out the partition and file system (Windows). - The scaled-out disk is the cloud disk of Linux CVM: Scale out the partition and file system (Linux). - Disk capacity ≥ 2 TB: - GPT partition format: Expand partitions and file systems (Windows) or Expand partitions and file systems (Linux). - MBR partition format: Not supported. MBR partition format supports a maximum disk capacity of 2TB. If your disk partition is in MBR format and its capacity needs to be expanded beyond 2TB, we recommend you create and mount a new data disk, and use GPT partition format to copy the data to the new disk.

Expanding Cloud Disks

Scenarios

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 the cloud disk capacity is expanded, you need to extend the partition and file system. You can allocate expanded capacity to an existing partition or format the expanded capacity into a new independent 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.

Operation Limits

1. The cloud disks only support expanding and do not support shrinking.
2. Only cloud disks with normal lifecycle support expanding, and isolated cloud disks do not support expanding.
3. The expanded capacity of the cloud disk should be an integer multiple of 10 GiB.

Expanding Data Disks

> Note:
>
> If multiple cloud disks of the same capacity and type are mounted to your CVM, you can identify them by seeing the operations shown in Distinguishing data disks. Select a data disk and expand it as instructed below.

Expanding via the CVM Console

1. Log in to the CVM console.
2. Select More > CVM Configuration > Adjust Disk in the row of the target CVM.
3. Select the required new capacity size (must be greater than the current capacity).
4. Depending on the operating system type of the target cloud service, you need to execute Expanding partitions and file system (Windows) or Expanding partition and file system (Linux) to allocate the capacity of the expanded part to the existing partition or format the capacity of the expanded part

into a separate new partition.

Expanding via the CBS Console

1. Log in to the CBS console.
2. Select More > Expansion on the target CBS.
3. Select the required new capacity size (must be greater than the current capacity).
4. Depending on the operating system type of the target cloud service, you need to execute Expanding partition and file system (Windows) or Expanding partition and file system (Linux) to allocate the capacity of the expanded part to the existing partition or format the capacity of the expanded part into a new independent partition.

Expanding via API

You can use the ResizeDisks API to expand the specified elastic cloud disks.

Expanding System Disks

Expanding via the CVM Console

1. Log in to the CVM console.
2. Select More > CVM Configuration > Adjust Disk in the row of the target CVM.
3. Select the required new capacity size (must be greater than the current capacity).
4. After the console expansion is completed, log in to the instance to confirm whether the file system has been automatically expanded.

Expanding via the CBS Console

1. Log in to the CBS console.
2. Select More > Expansion on the target CBS.
3. Select the required new capacity size (must be greater than the current capacity).
4. After the console expansion is completed, log in to the instance to confirm whether the file system has been automatically expanded.

Expanding via API

You can use the ResizeInstanceDisks API to expand the capacity of specified non-elastic cloud disks.

Related Operations

Distinguishing Data Disks

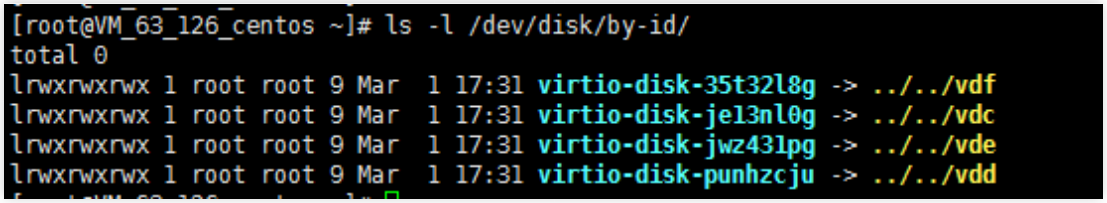
Check cloud disks according to the operating system of the CVM, and select your way:

Linux

1. Log in to a Linux instance.
2. Execute the following command to check the corresponding relationship between cloud disks and device names.

   ls -l /dev/disk/by-id

   The returned result is as shown in the following figure:



   In which, disk-xxxx is the ID of cloud disks. You can check it in the cloud disks console.

Windows

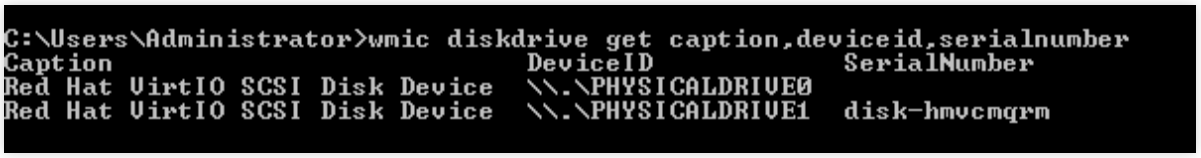
1. Log in to a Windows instance.
2. Right-click , and select Run.
3. Enter cmd in the popup window and press Enter.
4. Run the following command to check the corresponding relationship between cloud disks and device names.

   wmic diskdrive get caption,deviceid,serialnumber

   Or execute the following command:

   wmic path win32_physicalmedia get SerialNumber,Tag

The returned result is as shown in the following figure:



In which, disk-xxxx is the ID of cloud disks. You can check it in the cloud disks console.

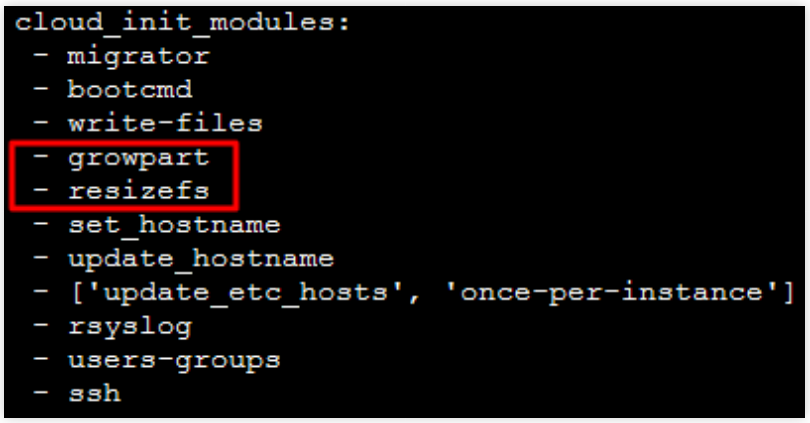
Checking the Cloudinit Configuration

Check cloud disks according to the operating system of the CVM, and select your way:

Checking the Cloudinit Configuration of Linux Instances

After the system disk is expanded, log in to the Linux instance and check whether the /etc/cloud/cloud.cfg file contains the growpart and resizefs configuration items.

- Yes. If the system disk is expanded and online expansion cannot be performed, you need to restart the machine. During system startup, cloud-init will call the growpart/resizefs command to complete expansion. As shown below:



  - growpart: Expand the partition size to the disk size.
  - resizefs: Expand and adjust the size of / partition file system to the partition.

- No. Manually expand the file system and partition according to the operating system type of the target cloud service. You need to execute Expanding partition and file system (Linux) to allocate the capacity of the expanded part to the existing partition or format the capacity of the expanded part as a new independent partition.

Checking the Cloudinit Configuration of Windows Instances

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 .

- Yes. If the cloudbase-init configuration file cloudbase-init.conf contains the ExtendVolumesPlugin configuration item, you need to restart the machine cloudbase-init before automatically expanding the volume and adding the blank space behind partition C to partition C. And there should be no other partition interference between the C partition and the blank space. If there are no other partitions between the C partition and the blank space, and you do not want to restart, or a third-party security software intercepts cloudbase-init so that it cannot complete the extend volume, you need to manually implement the following powershell command.

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

- No. Manually expand the file system and partition according to the operating system type of the target cloud service. You need to execute Expanding partition and file system (Windows) to allocate the capacity of the expanded part to the existing partition or format the capacity of the expanded part as a new independent partition.

Expanding System Disk Partitions and File Systems Online

Scenarios

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.

Prerequisites

- Before following the steps in this document, create a snapshot of the system disk for the corresponding instance to back up data. For details, see Creating a Snapshot.
  In case of data loss due to misoperation, you can roll back the snapshot to recover the data.
- You have expanded the cloud disk in the console. For detailed directions, see Expanding Cloud Disks.
- The kernel version of Linux CVM should be later than 3.6.0. You can use the uname -a command to check the kernel version.

Operating Environment

Linux instance

Resources	Description
Operating System	CentOS 8.0 64-bit
Cloud disk (system disk)	/dev/vda: uses MBR partition and ext4 file system, and has been expanded online from 50 GB to 60 GB via the console.

Windows instance

Resources	Description
Operating System	Windows Server 2012 R2 IDC 64-bit
Cloud disk (system disk)	C drive: uses MBR partition and NTFS file system, and has been expanded online from 50 GB to 100 GB via the console.

Operation Steps

Linux Operating System

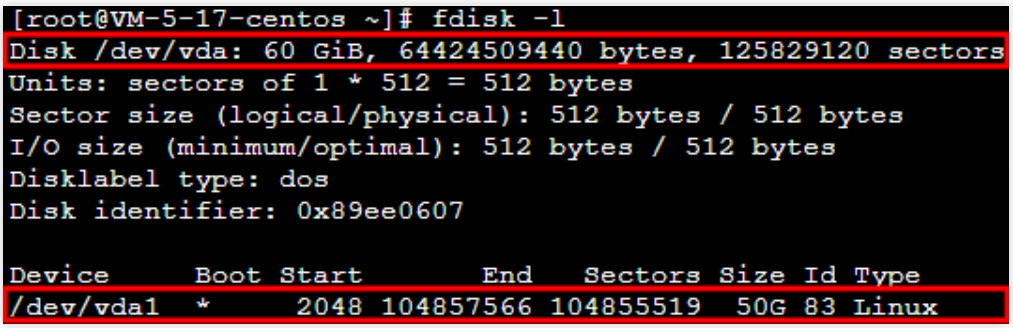
Perform the following steps based on your operating system:

1. Log in to the CVM instance as instructed in Logging In to Linux Instances.

2. Execute the following command to query the partition information of the cloud disk.

   fdisk -l

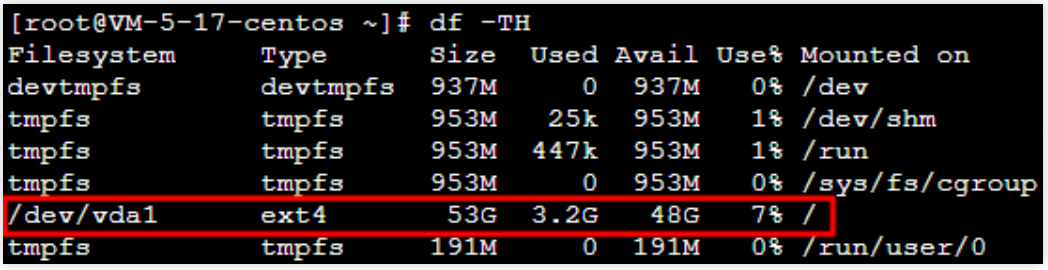
   The returned result is as shown in the figure below. It can be seen that the dev/vda data disk has a capacity of 60 GB, containing an MBR partition /dev/vda1 with a capacity of 50 GB.



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

   df -TH

   The returned result is as shown in the figure below, indicating that /dev/vda1 has a file system type of ext4.



4. Based on the CVM operating system type, execute the following command to install the growpart tool.

   - CentOS

yum install -y cloud-utils-growpart

- Ubuntu or Debian

  apt-get install -y cloud-guest-utils

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

   growpart /dev/vda 1

   The returned result is as shown in the following figure:

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

6. Execute the following command to expand the ext4 file system.

   resize2fs /dev/vda1

   The returned result is as shown in the following figure:

```
[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 scaling-out result.

   df -TH

   The returned result is shown in the following figure, indicating that the expansion succeeds.

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

After the expansion succeeds, check data integrity and observe whether the business in the CVM is running normally.

Windows Operating System

Perform the following steps based on your operating system:

1. Log in to the CVM instance as instructed in Logging In to Windows Instances.
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 the disk at the top of the page.
   After the scan is completed, you can view the new 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 expansion succeeds, check data integrity and observe whether the business in the CVM is running normally.

Expanding Partitions and File Systems (Windows)

Scenarios

After the cloud disk capacity is expanded, you need to allocate the expanded capacity to an existing partition or format the expanded capacity into a new independent partition.

- If you have performed a hard disk expansion operation while the cloud disk is connected to a CVM and the CVM is running normally, you need to first execute the Rescan Disk operation to identify the expanded hard disk space.
- If you scale out a hard disk when the hard disk is ready to mount or mounted but the server has been shut down, the space of the scaled-out hard disk will be automatically identified.

Prerequisites

> Note:
>
> - Inadvertently scaling-out the file system may affect existing data, so it is strongly recommended that you manually create a snapshot to backup data before operation.
> - Scaling out the file system requires restarting instances or rescanning disks, which will cause service interruption for a certain period. You are advised to choose the appropriate time with caution.

- The cloud disk capacity has been expanded.
- The cloud disk has been attached to a Windows CVM, and a file system has been created.
- You have logged in to the Windows CVM of the partition and file system to be scaled out.

Operation Steps

> Note:
>
> - If the CVM to which the disk is attached is running normally during cloud disk expansion, you need to rescan the disk to recognize the expanded cloud disk space before extending the volume.
> - If the disk is in the state of To Be Mounted or the CVM to which the disk is attached is shut down during cloud disk expansion, you can directly extend the volume.

> - This document takes the Windows Server 2008 operating system as an example. Different operating systems may have different formatting operations, so this document is for reference only.

(Optional) Rescan the Disk

1. Open Computer Management.
2. In the left sidebar, choose Storage > Disk Management.
3. Right-click Disk Management, and select Rescan Disks.
4. After the scan is complete, check whether the size of the data disk has changed to the scaled-out size.

Scaling out Volume

1. Right-click any blank area in the disk space and select Expand Volume.
2. Complete volume scaling-out operations according to the guidance of the Scale-out Volume. After that, the new data disk space will be merged into the original volume.

Expanding the Partition and the File System (Linux)

Scenarios

A cloud disk is an expandable cloud storage device. You can expand its capacity at any time after it is created to increase storage space without losing existing data on the cloud disk.

After the cloud disk capacity is expanded, you need to allocate the expanded capacity to an existing partition or format the expanded capacity into a new independent partition.

Prerequisites

> Note:
>
> Inadvertently scaling-out the file system may affect existing data, so it is strongly recommended that you manually create a snapshot to back up data before operation.

- The cloud disk capacity has been expanded.
- The cloud disk has been attached to a Linux CVM, and a file system has been created.
- You have logged in to the Linux CVM of the partition and file system to be scaled out.

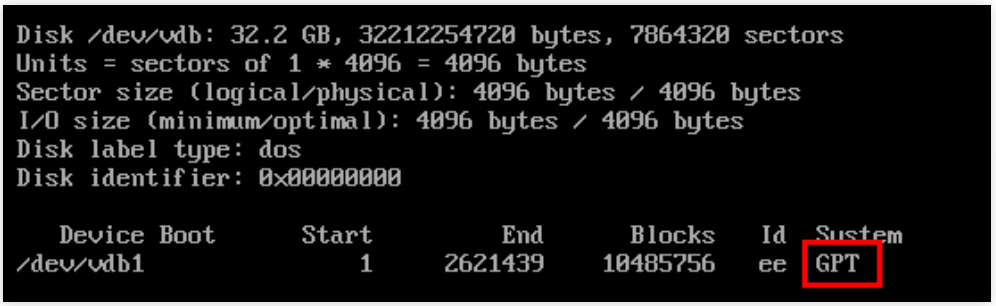
Operation Steps

Confirming Scale-Out Mode

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

   fdisk -l

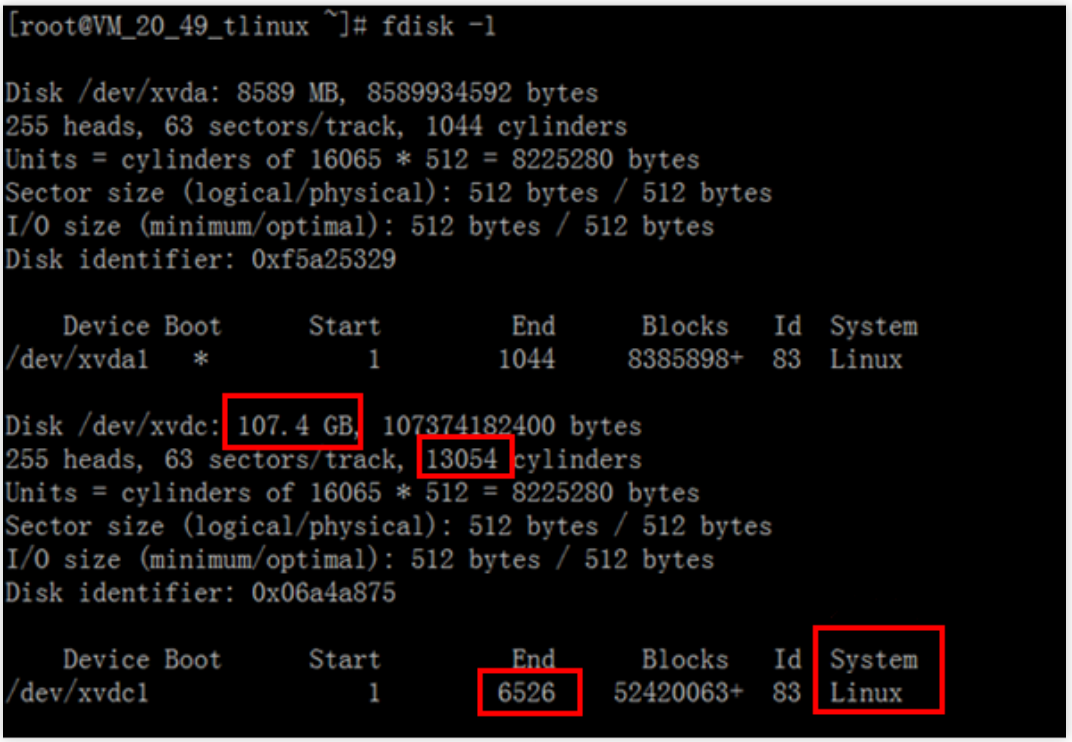
   - If the device displayed in the results has no partition (for example, only /dev/vdb is displayed), see Expanding File System.
   - If the results are shown in the following two figures (slightly different depending on the operating system), a GPT partition is used.





- If the results are shown in the following figures (slightly different depending on the operating system), an MBR partition is used.

> Note:
>
> The maximum disk capacity supported by the MBR partition is 2 TB. If your hard disk partition is in an MBR format and needs to be scaled out to exceed 2 TB, it is recommended that you create and mount a new data disk, use the GPT partition method, and copy the data to the new disk. For the Linux operating system, when GPT is selected as the disk partition form, the fdisk partition tool will not be available, and the parted tool needs to be used.



2. Based on the cloud disk partition format queried in step 1, select the corresponding operation guide.

Partition Form	Operation Guide	Description
-	Scaling out File System	It is applicable to the scene where no partition is created, and a file system is directly created on a raw device.
GPT	Allocating the Capacity of the Scaled-Out Part to the Original Partition (GPT)	It also applies to unpartitioned direct formatting scenes.
GPT	Formatting the Capacity of the Scaled-Out Part into a Separate New Partition (GPT)	The original partition can be kept unchanged.
MBR	Allocating the Capacity of the Scaled-Out Part to the Original Partition (MBR)	It also applies to unpartitioned direct formatting scenes.
MBR	Formatting the Capacity of the Scaled-Out Part into a Separate New Partition (MBR)	The original partition can be kept unchanged.

Scaling Out File System

1. Depending on the file system type, execute different commands to scale out.
   - For the EXT file system, execute the `resize2fs` command to scale-out.
   - For the XFS file system, execute the `xfs_growfs` command to scale-out.

   Take `/dev/vdb` as an example. The EXT file system executes the following command:

   ```
   resize2fs /dev/vdb
   ```

   Take `/dev/vdb` as an example. The XFS file system executes the following command:

   ```
   xfs_growfs /dev/vdb
   ```

2. Execute the following command to view the new partition.

   ```
   df -h
   ```

Allocating the Capacity of the Scaled-Out Part to the Original Partition (GPT)

1. Execute the following command as the root user to confirm the capacity change of the cloud disk.

   ```
   parted <disk_path> print
   ```

   This document takes the disk path of `/dev/vdb` as an example to execute the following command:

parted /dev/vdb print

If the following information is displayed in the process, input Fix .

```
Error: The backup GPT table is not at the end of the disk, as it should be.
This might mean that another operating system believes the disk is smaller.
Fix, by moving the backup to the end (and removing the old backup)?
Fix/Ignore/Cancel? Fix
Warning: Not all of the space available to /dev/vdb appears to be used, you can
fix the GPT to use all of the space (an extra 188743680 blocks) or continue with
the current setting?
Fix/Ignore? Fix
```

As shown in the figure below, the expanded cloud disk size is 107 GB, and the existing partition size is 10.7 GB.

```
[root@VM_164_30_centos ~]# parted /dev/vdb print
Model: Virtio Block Device (virtblk)
Disk /dev/vdb: 107GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt

Number  Start   End     Size    File system  Name     Flags
 1      1049kB  10.7GB  10.7GB  ext3         primary
```

2. Execute the command to confirm whether any partitions on this cloud disk are still attached.

mount | grep '<disk_path>'

This document takes the disk path of /dev/vdb as an example to execute the following command:

mount | grep '/dev/vdb'

As shown in the figure below, one partition (vdb1) on the cloud disk is attached to /data .

```
[root@VM_164_30_centos ~]# mount | grep '/dev/vdb'
/dev/vdb1 on /data type ext3 (rw)
```

3. Execute the following command to unmount the data disk:

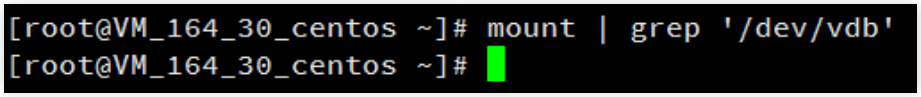
umount <mount_point>

This document takes the mount point of /data as an example to execute the following command:

umount /data

> Note:
>
> Detach the file systems of all partitions on the cloud disk, and then execute the operations in step 4. Repeat the following command to confirm that the file systems of all partitions on the hard disk have been detached.
>
> mount | grep '/dev/vdb'

File systems of all partitions on the cloud disk have been detached. As shown in the figure below:



4. Execute the following command to go to the parted partition tool:

   parted '<disk_path>'

   This document takes the disk path of /dev/vdb as an example to execute the following command:

   parted '/dev/vdb'

5. Execute the following command to change the display and manipulation units to sector (GB by default).

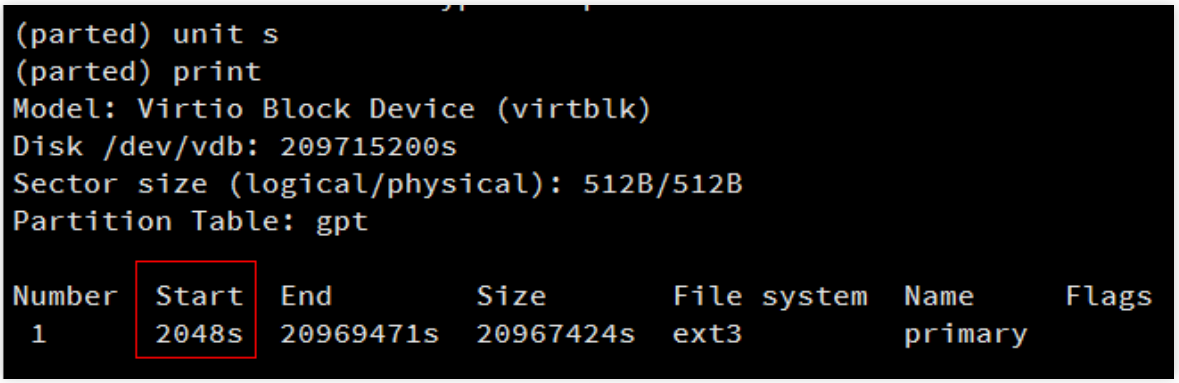
   unit s

6. Execute the following command to view partition information and record the Start value of existing partitions.

   > Note:
   >
   > After the partition is deleted and created, the Start value must remain unchanged; otherwise, data loss will occur.

   print

   The Start value in this document is 2,048s. As shown below:



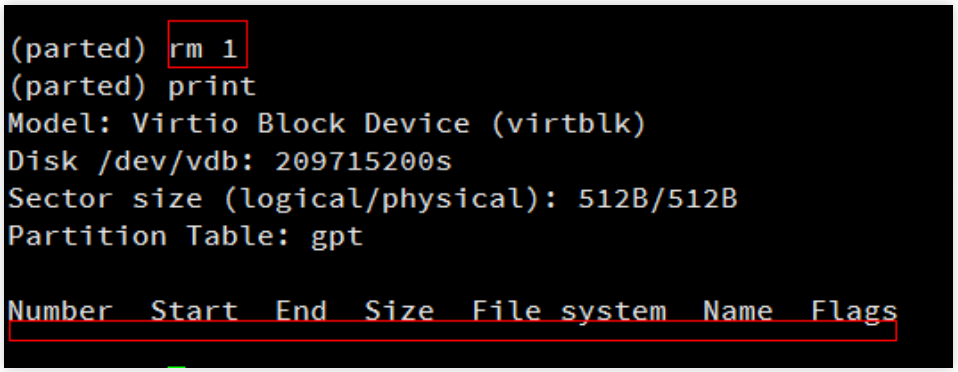
7. Execute the following command to delete the partition.

   rm <partition_number>

   For example, as shown in the figure above, there is one partition on the cloud disk with the Number as 1, so execute the command:

   rm 1

   The output is shown in the following figure.



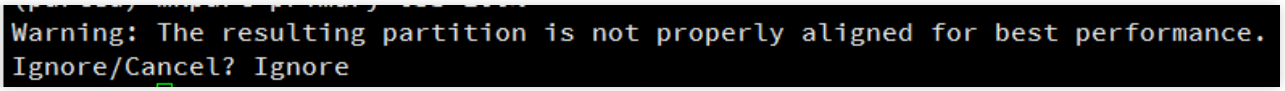
8. Execute the following command to create a new primary partition:

   mkpart primary <original_partition_originate_sector> 100%

   Among them, 100% means that this partition reaches the end of the disk.
   For example, if the primary partition starts from the 2,048th sector (which must be consistent with that before deletion, i.e., the Start value is 2,048s), execute the following command:

   mkpart primary 2048s 100%

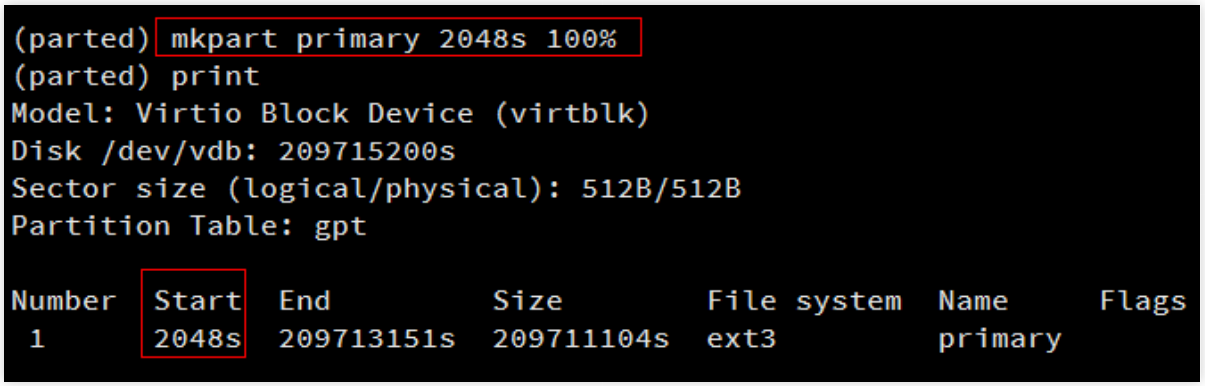
If the status shown in the following figure is displayed, input Ignore .



9. Execute the following command to check whether a new partition has been created.

   print

   The return result is shown in the following figure, indicating that the new partition has been created successfully.



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

    quit

11. Execute the following command to check the partitions after scale out:

    e2fsck -f <partition_path>

This document takes the new partition of 1 (i.e., the partition path is /dev/vdb1) as an example to execute the following command:

e2fsck -f /dev/vdb1

The results shown in the following figure are returned.

```
[root@VM_164_30_centos ~]# e2fsck -f /dev/vdb1
e2fsck 1.41.12 (17-May-2010)
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/vdb1: 11/655360 files (0.0% non-contiguous), 79696/2620928 blocks
```

12. Execute the following command to scale out the EXT file system on the new partition.

resize2fs <partition_path>

This document takes the partition path of /dev/vdb1 as an example to execute the following command:

resize2fs /dev/vdb1

The following figure shows the successful scale-out:

```
[root@VM_164_30_centos ~]# resize2fs /dev/vdb1
resize2fs 1.41.12 (17-May-2010)
Resizing the filesystem on /dev/vdb1 to 26213888 (4k) blocks.
The filesystem on /dev/vdb1 is now 26213888 blocks long.
```

13. Execute the following command to scale out the XFS file system on the new partition.

xfs_growfs <partition_path>

This document takes the partition path of /dev/vdb1 as an example to execute the following command:

xfs_growfs /dev/vdb1

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

mount <partition_path> <mount_point>

This document takes the partition path of /dev/vdb1 and the mount point of /data as an example to execute the following command:

mount /dev/vdb1 /data

15. Execute the following command to view the new partition.

df -h

If the following information is returned, the mounting is successful, and the data disk can be viewed.

```
[root@VM_164_30_centos ~]# mount /dev/vdb1 /data
[root@VM_164_30_centos ~]# df -hl
Filesystem            Size  Used Avail Use% Mounted on
/dev/vda1             7.9G  902M  6.6G  12% /
tmpfs                 1.9G     0  1.9G   0% /dev/shm
/dev/vdb1              99G  163M   94G   1% /data
```

Formatting the Capacity of the Scaled-Out Part into a Separate New Partition (GPT)

1. Execute the following command as the root user to confirm the capacity change of the cloud disk.

   parted <disk_path> print

   This document takes the disk path of /dev/vdb as an example to execute the following command:

   parted /dev/vdb print

   If the following information is prompted in the process, input Fix .

   ```
   Error: The backup GPT table is not at the end of the disk, as it should be.
   This might mean that another operating system believes the disk is smaller.
   Fix, by moving the backup to the end (and removing the old backup)?
   Fix/Ignore/Cancel? Fix
   Warning: Not all of the space available to /dev/vdb appears to be used, you can
   fix the GPT to use all of the space (an extra 188743680 blocks) or continue with
   the current setting?
   Fix/Ignore? Fix
   ```

   As shown in the figure below, the expanded cloud disk size is 107 GB, and the existing partition size is 10.7 GB.

   ```
   [root@VM_164_30_centos ~]# parted /dev/vdb print
   Model: Virtio Block Device (virtblk)
   Disk /dev/vdb: 107GB
   Sector size (logical/physical): 512B/512B
   Partition Table: gpt

   Number  Start   End     Size    File system  Name     Flags
    1      1049kB  10.7GB  10.7GB  ext3         primary
   ```

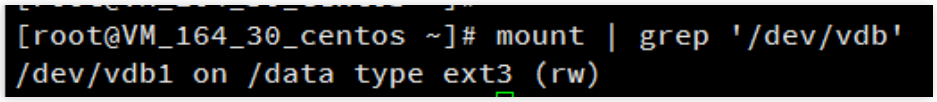
2. Execute the command to confirm whether any partitions on this cloud disk are still attached.

   mount | grep '<disk_path>'

This document takes the disk path of /dev/vdb as an example to execute the following command:

```
mount | grep '/dev/vdb'
```

As shown in the figure below, one partition (vdb1) on the cloud disk is attached to /data .



3. Execute the following command to unmount the data disk:

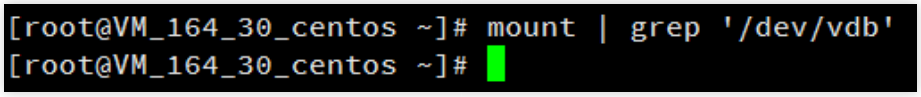
```
umount <mount_point>
```

This document takes the mount point of /data as an example to execute the following command:

```
umount /data
```

> Note:
>
> Detach the file systems of all partitions on the cloud disk, and then execute the operations in step 4. Repeat the following command to confirm that the file systems of all partitions on the hard disk have been detached.
>
> ```
> mount | grep '/dev/vdb'
> ```

File systems of all partitions on the cloud disk have been detached. As shown in the figure below:



4. Execute the following command to go to the parted partition tool:

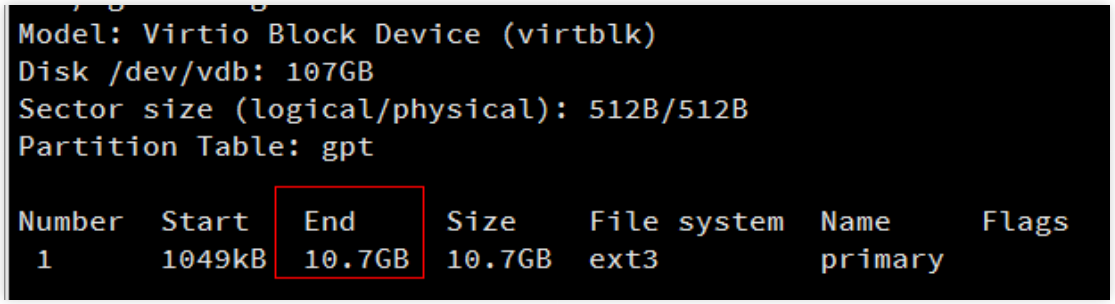
```
parted '<disk_path>'
```

This document takes the disk path of /dev/vdb as an example to execute the following command:

```
parted '/dev/vdb'
```

5. Execute the following command to view the partition information and record the End value of the existing partition, which is used as the start offset value of the next partition.

print



6. Execute the following command to create a new primary partition: This partition will start at the end of an existing partition and cover all new space on your hard disk.

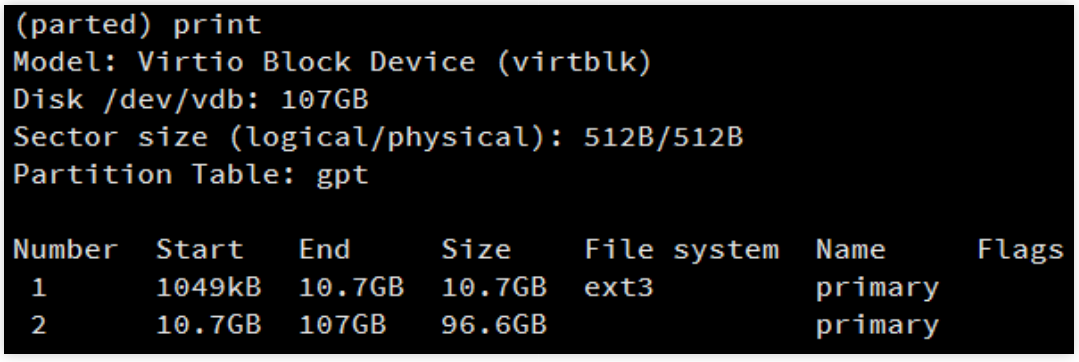
   mkpart primary start end

   This document takes the End value of 10.7 GB as an example to execute the following command:

   mkpart primary 10.7GB 100%

7. Execute the following command to check whether a new partition has been created.

   print



8. Execute the following command to exit the parted tool:

   quit

9. Execute the following command to format the new partition:

mkfs.<fstype> <partition_path>

You can choose the file system format yourself, such as EXT2 or EXT3.

This document takes the EXT3 file system as an example, so execute the command:

mkfs.ext3 /dev/vdb2

Allocating the Capacity of the Scaled-Out Part to the Original Partition (MBR)

The fdisk/e2fsck/resize2fs automatic expansion tool applies to Linux operating systems and is used to add the expanded cloud disk space to the existing file system. For the expansion to succeed, the following four conditions should be met:

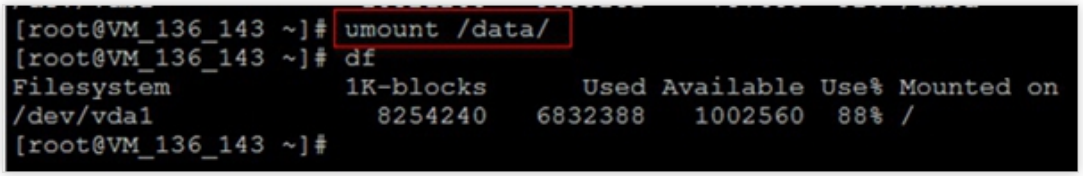
- The file system is EXT2/EXT3/EXT4/XFS.
- The current file system cannot contain errors.
- The disk size after scale-out must not exceed 2 TB.
- The current tool only supports the Python 2 version and does not support the Python 3 version.

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

   umount <mount_point>

   This document takes the mount point of /data as an example to execute the following command:

   umount /data



2. Open the browser, enter //imgcache.convergecloud.com/tcloud/cbs/devresize.py to download the tool script, and place it in the /tmp directory of the machine to be scaled out.

3. Execute the following command to scale out using the scale out tool:

   python /tmp/devresize.py <hard_disk_path>

   This document takes the hard disk path of /dev/vdb , and the file system is on vdb1 as an example to execute the following command:

python /tmp/devresize.py /dev/vdb



- If the output The filesystem on /dev/vdb1 is now XXXXX blocks long. is displayed, the scale-out succeeds. In this case, perform step 4.
- If the output is [ERROR]-e2fsck failed!!, perform the following steps:

  a. Execute the following command to restore the partition where the file system is located:

  fsck -a <partition_path>

  This document takes the hard disk path of /dev/vdb , and the file system is on vdb1 as an example to execute the following command:

  fsck -a /dev/vdb1

  b. After repairing successfully, execute the following command again to scale-out using the scale-out tool.

  python /tmp/devresize.py /dev/vdb

4. Execute the following command to manually mount the scaled-out partition:

mount <partition_path> <mount_point>

This document takes the mount point of /data as an example.

- If there are partitions before the scaling-out and the partition path of /dev/vdb1 is taken as an example, execute the following command:

mount /dev/vdb1 /data

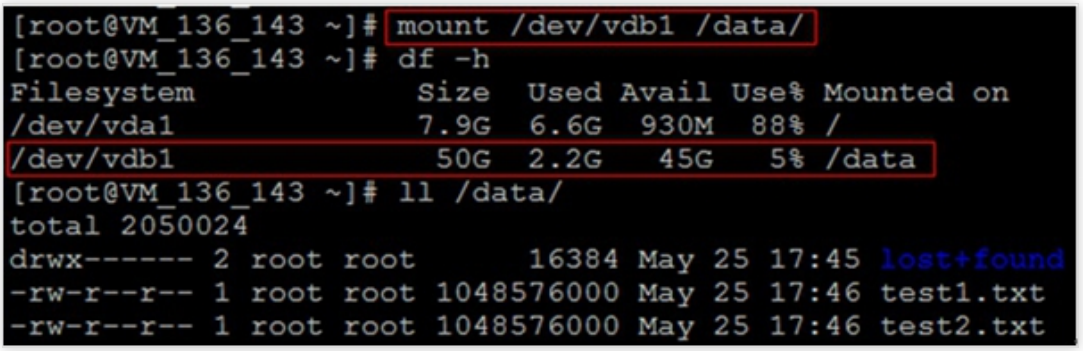
- If there is no partition before scaling out, execute the following command:

  mount /dev/vdb /data

5. Execute the following command to view the partition capacity after scaling out:

   df -h

   If information similar to that shown in the following figure is returned, the mounting is successful, and the data disk can be viewed:



6. Execute the following command to view the data information of the original partition after scaling out and check whether the newly added storage space is scaled out in the file system.

   ll /data

Formatting the Capacity of the Scaled-Out Part into a Separate New Partition (MBR)

1. Execute the following commands as user root to view the partition information of the mounted data disk.

   df -h

```
[root@VM_20_49_tlinux ~]# df -h
Filesystem            Size  Used Avail Use% Mounted on
/dev/xvda1            7.9G  2.1G  5.5G  28% /
tmpfs                 4.0G     0  4.0G   0% /dev/shm
/dev/xvdc1             50G  180M   47G   1% /data
[root@VM_20_49_tlinux ~]#
```

2. Execute the following command to view information about unpartitioned data disks after scaling out:

   fdisk -l

```
[root@VM_20_49_tlinux ~]# fdisk -l

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

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

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

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

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

   umount <mount_point>

   This document takes the mount point of /data as an example to execute the following command:

   umount /data

> Note:
>
> Detach all partitions on the cloud disk, and then execute step 4.

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

   fdisk <hard_disk_path>

   This document takes the disk path of /dev/xvdc as an example to execute the following command:

   fdisk /dev/xvdc

   According to the prompts on the interface, sequentially input p (to view existing partition information), n (to create partition), p (to create primary partition), and 2 (to create a second primary partition), press Enter twice (to use default configuration), and input w (to save partition table) to start partition. As shown in the figure below:

```
[root@VM_20_49_tlinux /]# fdisk /dev/xvdc

WARNING: DOS-compatible mode is deprecated. It's strongly recommended to
         switch off the mode (command 'c') and change display units to
         sectors (command 'u').

Command (m for help): p

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

Command (m for help): n
Command action
   e   extended
   p   primary partition (1-4)
p
Partition number (1-4): 2
First cylinder (6527-13054, default 6527):
Using default value 6527
Last cylinder, +cylinders or +size{K,M,G} (6527-13054, default 13054):
Using default value 13054

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

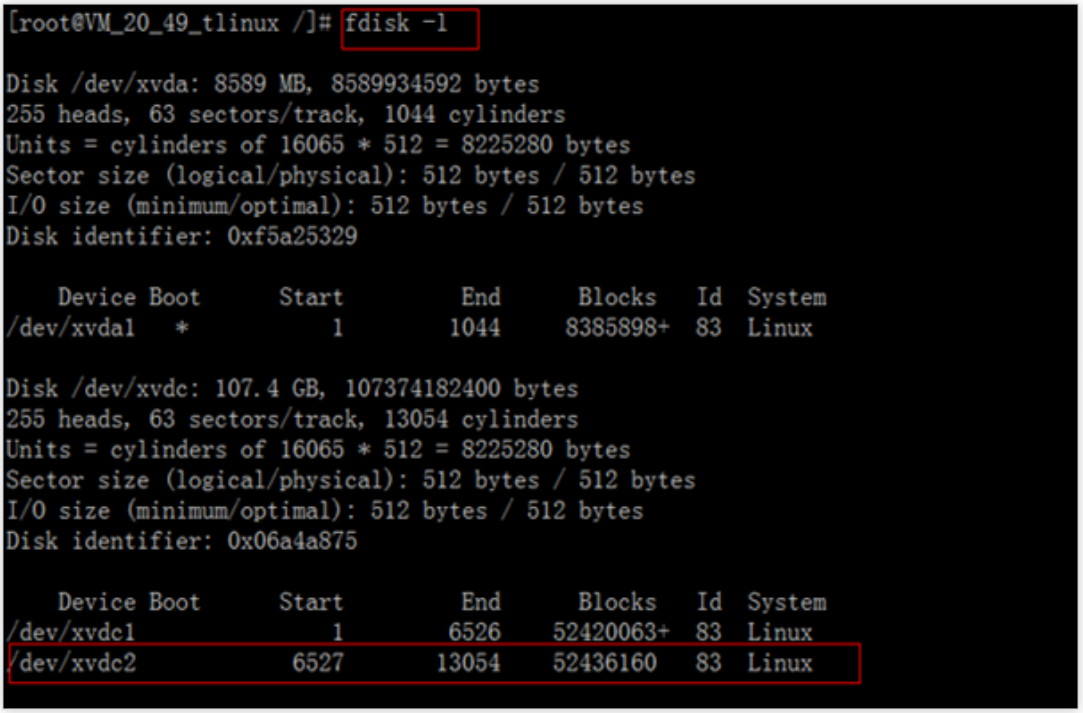
Calling ioctl() to re-read partition table.
Syncing disks.
[root@VM_20_49_tlinux /]#
```

> Note:
>
> This document takes the creation of one partition as an example. You can also create multiple partitions according to your actual needs.

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

fdisk -l

The following figure indicates that a new partition xvdc2 has been created.

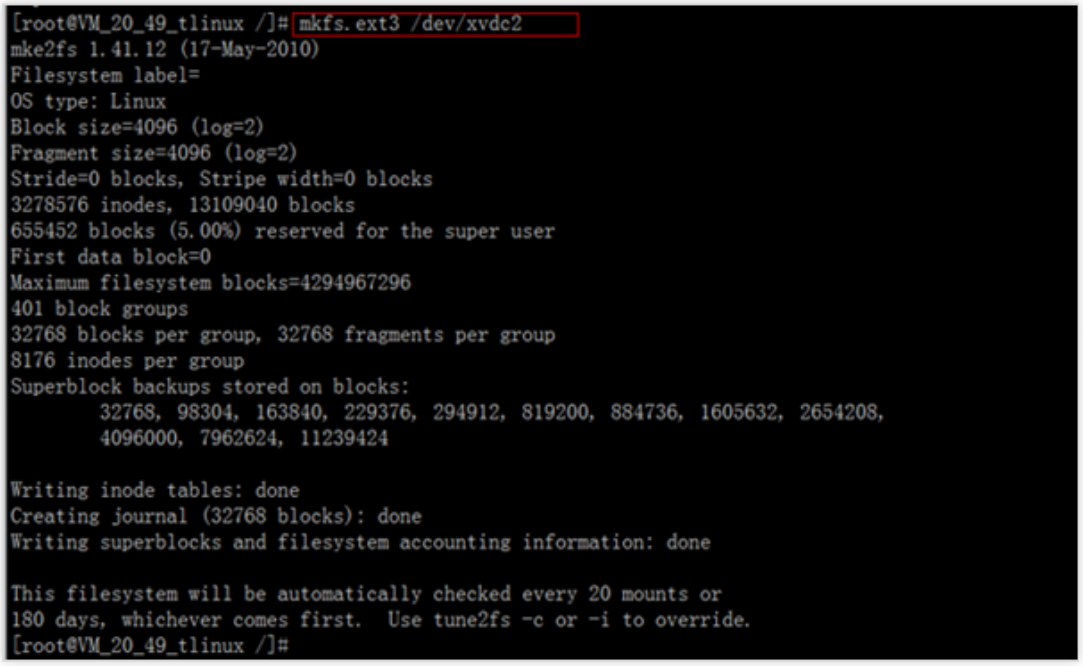


6. Execute the following commands to format the new partition and create a file system.

mkfs.<fstype> <partition_path>

You can choose the format of your file system, such as EXT2 and EXT3.
This document takes the file system of EXT3 as an example to execute the following command:

mkfs.ext3 /dev/xvdc2



7. Execute the following command to create a new mount point.

   mkdir <new_mount_point>

   This document takes the new mount point of /data1 as an example to execute the following command:

   mkdir /data1

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

   mount <new_partition_path> <new_mount_point>

   This document takes the new partition path of /dev/xvdc2 and the new mount point of /data1 as an example to execute the following command:

   mount /dev/xvdc2 /data1

9. Execute the following command to view information about the new partition:

   df -h

   If the following information is returned, the mounting is successful, and the data disk can be viewed.

```
[root@VM_20_49_tlinux /]# mkdir /data1
[root@VM_20_49_tlinux /]# mount /dev/xvdc2 /data1
[root@VM_20_49_tlinux /]# df -h
Filesystem            Size  Used Avail Use% Mounted on
/dev/xvda1            7.9G  2.1G  5.5G  28% /
tmpfs                 4.0G     0  4.0G   0% /dev/shm
/dev/xvdc1             50G  180M   47G   1% /data
/dev/xvdc2             50G  180M   47G   1% /data1
[root@VM_20_49_tlinux /]#
```

> Note:
>
> If you want the CVM to automatically mount on the data disk when it is restarted or started, you need to perform step 10 and step 11 to add new partition information to /etc/fstab
>
> .

10. Execute the following command to add a message.

```
echo '/dev/xvdc2 /data1 ext3 defaults 0 0' >> /etc/fstab
```

11. Execute the following command to view the message:

```
cat /etc/fstab
```

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

```
[root@VM_20_49_tlinux /]# echo '/dev/xvdc2  /data1   ext3   defaults  0 0' >> /etc/fstab
[root@VM_20_49_tlinux /]# cat /etc/fstab
/dev/xvda1              /                       ext3       noatime,acl,user_xattr 1 1
/dev/xvdb               swap                    swap       defaults 0 0
/dev/xvdc1              /data                   ext3       noatime,acl,user_xattr 1 2
tmpfs                   /dev/shm                tmpfs      defaults               0 0
proc                    /proc                   proc       defaults               0 0
sysfs                   /sys                    sysfs      noauto                 0 0
debugfs                 /sys/kernel/debug       debugfs    noauto                 0 0
devpts                  /dev/pts                devpts     mode=0620,gid=5        0 0
/dev/xvdc2              /data1                  ext3       defaults               0 0
[root@VM_20_49_tlinux /]#
```

Adjusting the Cloud Disk Type

Scenarios

CBS supports adjusting the storage hardware type while the business is running. You can upgrade the type while the service is online to meet the higher performance requirements for business development. Currently, the cloud disk type can only be upgraded and not downgraded. Details are as follows:

- A premium cloud disk can be upgraded to an SSD cloud disk.

Prerequisites

- Cloud Virtual Machine (CVM) Status: For CBS that has been mounted to a CVM, you can adjust the CBS type only when the CVM is in Running or Shutdown status.
- CBS Status: Adjustment is supported only when the CBS is In use or Not Mounted.
- Currently, the type of system disks and non-elastic data disks cannot be adjusted.
- The cloud disk type can only be adjusted when the current availability zone has available cloud disk types for upgrading and the current disk size is within the range supported by the target disk type.
- Adjusting the CBS type does not change the size of the disk. After adjustment, you can refer to Expanding Cloud Disk to adjust the disk size.
- Adjusting the CBS type does not change the cloud disk's lifecycle, disk ID, disk device name, and mount point.

Notes

- Adjusting the cloud disk type uses the data copying method to copy data from the source cloud disk to the target cloud disk. This operation may take a long time due to the size of the data and the data transfer speed. Please wait patiently. This process may slightly affect cloud disk performance, so it is recommended to perform the adjustment during a business downturn period if it involves critical services.
- Downgrading the cloud disk type is not supported.
- We recommend that after completing the adjustment operations, you start up and log in to the CVM to verify data integrity.

Directions

1. On the Cloud Disk page, select More > Adjust Cloud Disk Type on the right side of the target cloud disk.
2. In the pop-up dialog box, select the target disk type for adjustment, check to agree to the terms, and click OK.

Detaching Cloud Disks

Scenarios

When you need to attach an elastic cloud disk with a disk attribute of Data Disk to another CVM for use, you can actively detach the elastic cloud disk from the CVM and attach it to another CVM. Detaching the elastic cloud disk will not clear the data on the disk.

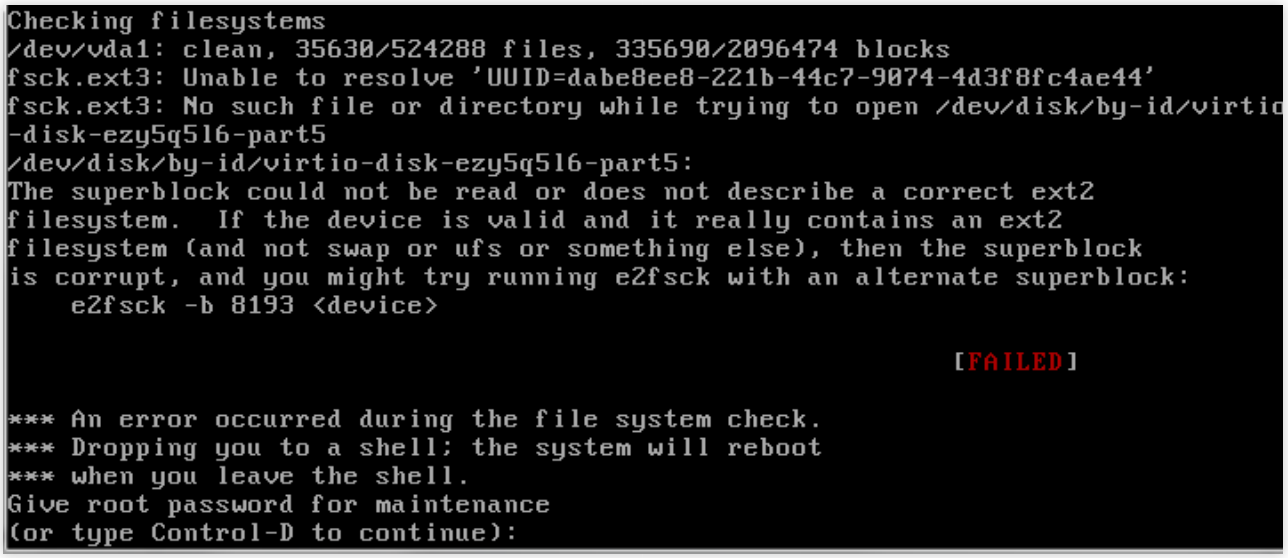
Currently, detaching an elastic cloud disk with a disk attribute of Data Disk is allowed. System disks and non-elastic cloud disks cannot be detached. Before detaching a cloud disk, you need to perform the umount (Linux) or offline (Windows) operation. Otherwise, the elastic cloud disk may not be identifiable when it is attached to the CVM again.

Prerequisites

Before detaching a data disk, make sure that you have understood the following precautions and prerequisites:

- In Windows operating system
  - To ensure data integrity, it is advised that reading and writing operations on all file systems in the disk be paused. Otherwise, the data not read or written will be lost.
  - When detaching an elastic cloud disk, you need to set the disk to offline first. Otherwise, without restarting the CVM, you may be unable to attach the elastic cloud disk again.

- In Linux operating system
  - You need to log in to the instance first and perform the unmount operation on the elastic cloud disk that needs to be detached. If the unmount operation is not performed and the disk is forcibly detached directly, the following issue may occur during shutdown and startup:



- If you have created a Logical Volume Manager (LVM) in CVM, detaching the cloud disk directly from the console will result in some device data remaining in instance memory. When an application in the instance tries to traverse or access the device, a system error will occur. Therefore, the following operations must be executed in advance (this example assumes that logical volume /dev/test/lv1 has been created based on/ dev/vdb1 and mounted under /data directory):
  a. Execute `umount /data` command at the corresponding umount point in the instance.
  b. Execute `lvremove /dev/test/lv1` command to remove Logical Volume (LV). If there is more than one LV, remove all LVs one by one.
  c. Execute `vgremove test` command to remove Volume Group.
  d. Execute `pvremove /dev/vdb1` command to remove physical volume.
  e. Change `/etc/fstab` file to avoid mounting the corresponding LV at the next startup.

Operation Steps

Detaching a cloud disk using the console

1. Log in to the CBS console.
2. You can detach a cloud disk using the following methods:
   - Single detachment: Click More > Unmount in the row of the target cloud disk in the state of In use.
   - Batch detachment: Select multiple target cloud disks in the state of In use and click Unmount above the list.
3. In the Unmount Cloud Disk pop-up prompt box, confirm warning items, and click OK to complete unmount.

Detaching a cloud disk using the API
You can use the `DetachDisks` API to detach a cloud disk.

Terminating Cloud Disks

Scenarios

When a cloud disk is no longer used and important data backup has been completed, you can release virtual resources by terminating the cloud disk. After the cloud disk is terminated, there will be no charge for it. 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 depends on the CVM and can only be terminated when the CVM is terminated. For specific operations, see Terminating Instances.
The lifecycle of an elastic cloud disk is independent of the CVM; therefore, it can be terminated independently of the CVM. This document mainly describes the operations for terminating elastic cloud disks.

Prerequisites

- The cloud disk is in the state of To be mounted. For cloud disks that are in use, you need to detach them first.
- Important data backup has been completed considering demands of business.

Operation Steps

1. Log in to the CBS console.

2. You can terminate a cloud disk through the following methods:
    - Individual termination: In the row where the target cloud disk at To be mounted status is located, select More > Terminate.
    - Batch termination: Select the target cloud disks in the To be mounted state, and click Terminate at the top of the list.

      > Note:
      >
      > When a cloud disk is terminated, all data on the cloud disk will be deleted and cannot be retrieved. The cloud disk that has been terminated is irrecoverable, so proceed with caution.

3. In the pop-up dialog box for terminating the cloud disk, click OK to complete the termination.

Billing for the target cloud disk will stop. The cloud disk will be permanently terminated and cannot be recovered.

Adjusting Cloud Disk Performance

Cloud disk performance is typically related to cloud disk capacity. You can achieve higher performance by adjusting the capacity when the cloud disk has not reached its maximum performance. Performance of an Enhanced SSD can exceed the baseline performance limits by configuring additional performance after the maximum baseline performance is reached. You can configure additional performance as needed and adjust it at any time when conditions are met.

> Note:
>
> - Currently, only Enhanced SSDs support independent performance adjustment.
> - Only when the benchmark performance has reached its maximum value can additional performance be adjusted independently.
> - During the process of adjusting cloud disk performance, business operations are not affected and can run normally.

Fee Description of Cloud Disk Performance Adjustment

Performance Upgrade

For pay-as-you-go cloud disks: The adjustment takes effect immediately with the billing price of the new configuration.

Performance Degrade

For pay-as-you-go cloud disks: The adjustment takes effect immediately with the billing price of the new configuration.

Performance Upgrade

Performance Upgrade Using the Console

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

1. Log in to the CBS console.
2. Select a region, and select the cloud disk for which you need to adjust performance.
3. Choose More > Adjust Performance for the target cloud disk.
4. In the Adjust Performance pop-up, select the target configuration that you need to adjust.
5. Check the description and start adjusting.

Performance Upgrade Using the API

You can use the ModifyDiskExtraPerformance API to upgrade the performance of the specified cloud disk.

Performance Degrade

Performance Degrade Using the Console

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

1. Log in to the CBS console.
2. Select a region, and select the cloud disk for which you need to adjust performance.
3. Choose More > Adjust Performance for the target cloud disk.
4. In the Adjust Performance pop-up, select the target configuration that you need to adjust.
5. Check the description and start adjusting.

Performance Degrade Using the API

You can use the ModifyDiskExtraPerformance API to degrade the performance of the specified cloud disk.

Managing Snapshots
Creating Snapshots

Scenarios

You can create a snapshot from a cloud disk to save the cloud disk data at a specified point in time. Snapshots are created incrementally, meaning only the data that has changed since the last snapshot is created. In this way, data changes are minimal, so that snapshots can be created quickly. Although snapshots are created incrementally, deleting a snapshot does not affect your use of any snapshot data. Any undeleted snapshot can restore the cloud disk to that snapshot's state.

You can create a snapshot from any state of a cloud disk, but a snapshot can only save data that has been written at the current time. If an application or process is writing data, such data may not be saved to the snapshot created at that time. Depending on your actual business situation, you can choose to temporarily stop all writes and create a snapshot promptly, or first detach the cloud disk from the CVM, create a snapshot, and then reattach it to obtain a complete data snapshot.

Prerequisites

- The cloud disk has been successfully created.

Notes

The snapshot only retains the data that has been written to the hard disk at that time and does not retain the data that has not been written to the hard disk in the memory at that time. It is strongly recommended that you shut down your computer before creating a snapshot or ensure that memory data has been written to the hard disk, pause reading and writing from the hard disk, and operate in the following two ways:

- Database Level
  For database business, it is recommended that all tables in the database be locked to read-only status first to prevent new data from being written into when creating a snapshot so that new data cannot be captured by the snapshot. This document takes mysql as an example to perform the following operations:
  Execute FLUSH TABLES WITH READ LOCK command to close all opened tables, and lock all tables in all databases using global read-lock. As shown below:

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

Create a snapshot of the cloud disk.
Execute UNLOCK TABLES command for unlocking. As shown below:

```
mysql> unlock tables;
Query OK, 0 rows affected (0.00 sec)
```

- System Level
  From the perspective of the system, data is usually stored in the memory buffer during processing and then written into the cloud disk at an appropriate time to improve the operation efficiency of the system. Therefore, when a snapshot is created, the data that has not been written into the cloud disk in the buffer cannot be written to and recovered from the snapshot, which affects data consistency. Execute the command sync to force the data in the file system memory buffer to be written to the cloud disk immediately, and avoid writing new data before creating the snapshot.
  After the command is executed, no error message will be returned; that is, the data in the cache has been written to the cloud disk. As shown below:

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

Operation Steps

Creating Snapshots Using the Console

1. Log in to the CBS console.
2. Click Create Snapshot in the target cloud disk row.
3. In the create snapshot dialog box, enter the snapshot name and click OK.

Creating Snapshots Using the API

You can create a snapshot using the CreateSnapshot interface.

Rolling Back Data from Snapshots

Scenarios

When certain changes cause data errors or data loss, you can roll back the snapshot to the cloud disk from which it was created, thereby restoring the data on the cloud disk to its state when the snapshot was created.

- The snapshot data can only be rolled back to the source cloud disk, not to other cloud disks.
- Scenes where rollback can be performed:
  - If the source cloud disk is in the state of To be mounted (that is, not attached to a CVM), you can directly perform the rollback operation.
  - If the source cloud disk is already attached to a CVM, the corresponding CVM needs to be shut down to perform the rolling back operation.

Operation Steps

Rolling Back Snapshots Using the Console

1. Log in to the CBS console.

2. In the left sidebar, choose Snapshot > Snapshot List.

3. In the Snapshot List page, click Roll Back in the target snapshot row.

   > Note:
   >
   > The data on the source cloud disk will be rolled back to the state when the snapshot was created. Data after this time will be cleared, so proceed with caution.

4. On the rollback data page, confirm rollback information and click OK to start rollback.

Rolling Back Snapshots Using the API

You can use ApplySnapshot API to create a snapshot.

Creating Cloud Disks from Snapshots

Scenarios

Snapshots are an important method for data sharing and migration. A cloud disk created from a snapshot contains all the data in the snapshot. You can use a snapshot to create a cloud disk with a capacity equal to or greater than that of the snapshot.

- When a snapshot is used to create a data disk of the same size, the new data disk can be read and written properly after direct mounting to the CVM, and no initialization is required.
- When you use a snapshot to create a data disk with a capacity larger than the snapshot, the system only completes the disk scaling out at the block device level, and does not implement the file system scaling out or automatic conversion of the partition format. When a new data disk is mounted, only the file system and data of the source snapshot can be used, and the new disk space cannot be directly used. You need to manually scale out the file system or even convert the partition format.

For example, use a data disk snapshot with an MBR partition format and a capacity of 1TB to create a new 3TB data disk. Since the maximum disk space supported by MBR is 2TB, you need to format the data disk and repartition it using GPT, which deletes the original data. Therefore, please operate with caution according to your actual needs.

This document describes how to create a cloud disk from a snapshot on the Snapshot List page. Additionally, when creating a cloud disk, you can configure the snapshot parameter to specify the corresponding snapshot for creating a cloud disk.

Operation Steps

Creating a Cloud Disk from a Snapshot Using the Console

1. Log in to console and enter the Snapshot List page.
2. In the row of the target snapshot, click More > Create Cloud Disk.
3. In the purchase data disk dialog box, set the following parameters:

Parameter	Parameter Description
AZ	This parameter is required. The availability zone where a cloud disk is located cannot be changed after a cloud disk is created.

Parameter	Parameter Description
Cloud Disk Type	This parameter is required. Cloud disk types include: - Premium Disk - Cloud SSD
Capacity	This parameter is required. When you create a cloud disk using a snapshot, the capacity cannot be smaller than the snapshot size. If you do not specify the capacity of a cloud disk, the capacity defaults to the same as the snapshot size.
Snapshot	This parameter is optional. When creating a cloud disk from a snapshot, check Create Cloud Disk with Snapshot and select the snapshot you want to use. - The capacity of a cloud disk created with a snapshot is equal to the snapshot size by default, and you can increase the capacity to a value larger than the default value. - When a cloud disk is created with a snapshot, the disk type is consistent with the type of the source cloud disk by default. You can adjust the cloud disk type.
Disk Name	This parameter is optional. The name of cloud disk supports up to 20 characters. After the cloud disk is created, it can be modified. - Create a single cloud disk: The disk name is the name of the cloud disk. - Create cloud disks in batches: When creating multiple cloud disks at once, the disk name will use the cloud disk name as a prefix. The final cloud disk name will be formed as Disk Name_ Number, ranging from "Disk Name_0" to "Disk Name_49".
Billing Mode	Pay-as-you-go by default.
Purchase Quantity	This parameter is optional. The default quantity is 1 , indicating that only one cloud disk is created. Currently, a maximum of 50 cloud disks can be created in batches.

4. Click OK.
5. You can check the created cloud disk on the Cloud Disk List page. The newly created elastic cloud disk is in the to be attached state. See Attaching Cloud Disks to attach the cloud disk to CVM in the same availability zone.

Creating a Cloud Disk from a Snapshot Using the API

You can use CreateDisks API to create a snapshot.

Scheduled Snapshot

Feature Overview

The scheduled snapshot feature is enabled for CBS snapshots so that developers can set backup task policies flexibly.

Using different scheduled snapshot policies for different businesses is advised. Recommended settings are given in the following table:

Business Scenario	Snapshot Frequency	Retention Type
Core Business	Use scheduled snapshots, with the policy set to once a day.	7-30 days
Non-core, Non-data Business	Use scheduled snapshots, with the policy set to once a week.	1w
Archive Business	Snapshots are manually created considering actual business requirements, and no fixed frequency is to be set.	One month to a few months
Test Business	Snapshots are manually created considering actual business requirements, and no fixed frequency is to be set.	Delete upon use

Policy Description

You can learn more about the contents and features of scheduled snapshot policies from the following table so that you can better use snapshots in your businesses.

Policy Item	Description
Object	All cloud disks, including system disks and data disks.
Execution Policy	Automatic snapshot time point can be every hour to every day, which is valid for a long time after setting the execution policy. When an existing execution policy is changed, the policy will become valid upon successful change.
Regular Termination (Important)	Scheduled snapshots allow regular termination. The lifecycle of a snapshot is preset (1-30 days). Upon expiration, automatically generated snapshots will be automatically deleted so that backup cost is effectively reduced. If no regular termination policy is set, automatic snapshots will be retained for a long time.

Policy Item	Description
Batch	Multiple cloud disks can be selected so that the same scheduled snapshot policy can be executed in batches for multiple cloud disks. No more than 80 cloud disks can be associated with the same snapshot policy in batches at a time.
Script pause	For a cloud disk that is out of use, no automatic snapshot policy will be executed. Out-of-use cloud disk means that the CVM associated with the cloud disk is shutdown or the data disk is not mounted.
Naming rules	Automatic snapshot name is snap_yyyyMMdd_HH. In which, yyyyMMdd is date and HH is hour. You can also manually change snapshot name. For example, snap_20180418_11 is an automatic snapshot created at 11:00 on April 18th, 2018.
Life Cycle (Important)	A snapshot has two types of lifecycle: - The lifecycle of a manually created snapshot is Long-term storage by default, and the snapshot can be preserved for a long time so long as the account balance is sufficient. - Scheduled snapshot: The time point of Regular Terminate can be set by following creation rules, and the snapshot can also be set as long-term retention.
Snapshot conflict	There is no conflict between automatic snapshots and custom snapshots in terms of use, but there may be time conflicts during creation. - When an automatic snapshot is executed on a disk, the user has to wait until the automatic snapshot is completed before a custom snapshot can be created (and vice versa). - If the disk has a large data size and the duration of a snapshot exceeds the interval between two automatic snapshot time points, no automatic snapshot will be created at the next time point. For example, if the user sets 9:00, 10:00, and 11:00 as automatic snapshot time points and the automatic snapshot use time at 9:00 is 70 minutes (it will be completed at 10:10), the automatic snapshot will not be executed at 10:00, and the next snapshot time point is 11:00.
Snapshot quota (available in later versions)	Every disk has a certain snapshot quota. If the number of snapshots in a disk is at the upper limit, the automatic snapshot task will be suspended and blocked. Snapshot quota is to avoid forgetting an automatic snapshot policy by the developer, otherwise, there will be an endless increase in storage cost.
ASP	Indicates the scheduled snapshot policy, i.e., Auto Snapshot Policy.
ASP quota (available in later versions)	No more than 30 ASP policies can be set for a region under a single account. A single ASP can be associated with up to 200 hard disks.

Policy Item	Description
Preservation cycle	- For automatic snapshots, the console will display a countdown to recovery. Changing the automatic snapshot retention period to permanent retention is allowed. - For a manually created snapshot, the display will be retained permanently.
ASP pause function	ASP Auto Snapshot policy provides Manual Trigger Pause function. After pausing, no new snapshot will be automatically created. While the lifecycle of automatic snapshots generated will not be affected by the pause function, and the snapshot will be regularly terminated or stored for a long time depending on the set rules.
Operation Logs	The process of creation of all automatic snapshots will be displayed, which is the same as manually added ones.

Operation Guide

Creating a Scheduled Snapshot Policy

> Note:
>
> No more than 30 scheduled snapshot policies can be created in the same region under a single account.

1. Log in to the console and go to the Scheduled Snapshot Policy page.
2. Select the Region.
3. Click Create.
4. On the Create Snapshot Policy page, set the following parameters and click OK.
   The parameters are described as follows:

Parameter	Parameter Description
Name	This parameter is required. Name of a scheduled snapshot policy, with no more than 60 characters.
Region	This parameter is required. This parameter cannot be changed on the current page.
Backup Cycle	This parameter is required. Set the cycle for executing regular snapshots. Backups can be scheduled by day, week, or month.
Backup Time	This parameter is required. The time point for executing a scheduled snapshot can be an hour in 00:00-23:00 (considering actual backend operation, the set time for timing snapshot may differ from

Parameter	Parameter Description
	the actual time when a snapshot is created through the console. For data in the snapshot, the time when the snapshot is created through the console will apply).
Retention Type	This parameter is required. - Supports retention by quantity. Automatically deletes when a fixed number is exceeded, up to 65,536 can be retained. - Supports retention by time. Automatically deletes after a fixed number of days, default retention is 1 month. - Permanently Save.

Associating Cloud Disks

> Note:
>
> The same scheduled snapshot policy can be associated with no more than 200 cloud disks.

1. Log in to the console and go to the Scheduled Snapshot Policy page.
2. Select the Region.
3. Click Associate Cloud Disk in the target policy row.
4. On the Associate Cloud Disks page, check the cloud disk to be associated.
5. Click OK.

Enabling/Disabling a Scheduled Snapshot Policy

1. Log in to the Scheduled Snapshot Policy page.
2. Select the Region.
3. Find the target policy row, and click the switch button under Scheduled Snapshots to enable or disable a scheduled snapshot policy.

Changing a Scheduled Snapshot Policy

1. Log in to the Scheduled Snapshot Policy page.
2. Select the Region.
3. Click Modify Policy in target policy row.
4. On the Change Snapshot Policy page, change the related parameters and click OK.

Deleting a Scheduled Snapshot Policy

1. Log in to the Scheduled Snapshot Policy page.
2. Select the Region.
3. You can delete a scheduled snapshot policy as stated below:
    - Individual deletion: Click Delete in the target policy row.

- Batch deletion: Select the scheduled snapshot policies to be deleted, and click Delete at the top of the list.

Converting Automatic Snapshots to Permanently Retained Snapshots

> Note:
>
> If Snapshot Retention Period in an automatic snapshot policy has been set to Permanently Retention, the following operations are not required for snapshots automatically generated under the policy.

1. Log in to the Snapshot List page.
2. Select the Region.
3. Click the ID of the target automatic snapshot.
4. On the details page, click Edit and set the retention type to Permanently Save.
5. Return to the snapshot list to see that the Snapshot Retention Period for that snapshot has been changed to Permanently Retention.

Cross-Region Snapshot Replication

At present, cross-region replication of cloud disk snapshot is supported, through which you can either migrate data and services to other regions or build your business cross-region disaster recovery system based on cross-region replication capability.

Operation Steps

1. Log in to console and enter Snapshot List page.

2. Click Cross-Region Replication in the row where the target snapshot is located.

3. Set the following parameters:

   - New Snapshot Name: Fill in the name of the new snapshot, which has no more than 60 characters and is an optional parameter.
     The default name will contain source snapshot ID and region information and will be in the format of Copied <Source snapshot ID> from <Source snapshot region> , e.g., Copied snap-oi5spwt2 from ap-city .
   - Region: Select the target region to which a snapshot is to be copied, which is a mandatory parameter.
     In selecting a region, pay attention to snapshot quota and regional restrictions.

4. Click OK to start replication. For a source snapshot, the status prompt will be added, and for the target region, a new snapshot will be added.

5. After replication, the new snapshot can be checked in the snapshot list of the target region.

   > Note:
   >
   > Deleting the source snapshot during cross-region replication is not allowed.

   During the process of cross-region replication:

   - Status of source snapshot: it can be checked in the status bar of the source snapshot row in the source region snapshot list.
   - Status of target snapshot: it can be checked in the snapshot list page of the target region.

Deleting Snapshots

Scenarios

When a snapshot is no longer needed, you can delete it to release virtual resources.

Notes

- When a snapshot is deleted, only the data specific to that snapshot is deleted, and the cloud disk from which the snapshot was created will not be affected.
- Each snapshot of data provided can independently restore the cloud disk to its data state at that moment. Deleting a snapshot created earlier for a cloud disk does not affect your continued use of a snapshot created later.
- When a snapshot is deleted, all data in the snapshot will be deleted, and deleted data is irrecoverable. As the deleted snapshot is irrecoverable, act with caution.

Operation Steps

Deleting Snapshots Using the Console

1. Log in to console and enter the Snapshot List page.
2. You may delete snapshots as stated below:
    - Single deletion: Click Delete in the row of the snapshot to be deleted.
    - Batch deletion: Select all snapshots to be deleted (make sure they are not in task status) and click Delete above the list.
3. In the Delete Snapshot dialog box, after confirming the information as prompted, click OK.

Deleting Snapshots Using the API

You can use DeleteSnapshots API to delete snapshots.

Sharing Snapshots

Scenarios

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

Notes

- Each snapshot can be shared with up to 50 users.
- The system disk snapshots and encrypted snapshots cannot be shared.
- 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 unshare it.
- Snapshots can be shared with other accounts in the same region. If you need to share to different regions, you need to copy the snapshots to the corresponding regions before sharing.
- It is not possible to share the obtained shared snapshot with other users again.

Operation Steps

Getting the Unique ID of the Peer Account

The shared snapshot is identified by the unique ID of the other party's root account, which means that you first need to get the unique ID of the root account of the other party's account. The ways to obtain it are as follows:

1. Log in to the CVM console.
2. Click the account center in the upper right corner and select Account Info.
3. On the Account Info page, check and record the Account ID of the root account.
4. Inform the other party to send the obtained account ID to you.

Sharing Snapshots Using the Console

1. Log in to console and enter the Snapshot List page.
2. Choose More > Shared Snapshots on the right of the row of the snapshot to be shared.
3. In the pop-up Share Snapshot window, enter the unique ID of the peer account.
4. Click Share to share the snapshot.
5. Inform the other party to go to the snapshot list page and select the Shared Snapshots tab to view

the shared snapshots.

To share with multiple users, repeat the above steps.

Sharing Snapshots Using the API

You can use the ModifySnapshotsSharePermission API to share snapshots.

Managing Snapshot Groups
Creating Snapshot Groups

Scenarios

You can use Snapshot Group to create snapshots for multiple cloud disks in an instance at the same time.

Prerequisites

- You have successfully created the CVM.
- The quantities and total capacity of your snapshots in the current region have not reached the maximum value. For details, see Snapshot Usage Limitations.

Notes

The snapshot only retains the data that has been written to the hard disk at that time and does not retain the data that has not been written to the hard disk in the memory at that time (for example, files under the Linux system /run directory). It is strongly recommended that you shut down your computer before creating a snapshot or ensure that memory data has been written to the hard disk, pause reading and writing from the hard disk, and operate in the following two ways:

- Database Level
  For database business, it is recommended that all tables in the database be locked to read-only status first to prevent new data from being written into when creating a snapshot so that new data cannot be captured by the snapshot. This document takes mysql as an example to perform the following operations:
  Execute the FLUSH TABLES WITH READ LOCK command to close all opened tables and use the global read lock to lock all tables in all databases. As shown below:

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

  Create a snapshot of the cloud disk.
  Execute UNLOCK TABLES to unlock. As shown below:

```
mysql> unlock tables;
Query OK, 0 rows affected (0.00 sec)
```

- System Level
  From the perspective of the system, data is usually stored in the memory buffer during processing and then written into the cloud disk at an appropriate time to improve the operation efficiency of the system. Therefore, when a snapshot is created, the data that has not been written into the cloud disk in the buffer cannot be written to and recovered from the snapshot, which affects data consistency. Execute the `sync` command to force the data in the memory buffer of the file system to be written into the cloud disk immediately and avoid writing new data before creating a snapshot.
  After the command is executed, no error message will be returned; that is, the data in the cache has been written to the cloud disk. As shown below:

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

Operation Steps

Creating Snapshot Groups Using the Console

1. Log in to the CVM console.
2. On the Instances page, click More > Create Snapshot Group on the right side of the row containing the target instance.
3. In the pop-up Create snapshot group dialog box, input the snapshot name and click Confirm.

Creating Snapshot Groups Using the API

You can use the `CreateSnapshotGroup` API to create snapshot groups.

Rolling Back Data from Snapshot Groups

Scenarios

When certain changes cause data errors or data loss, you can roll back the snapshot to the cloud disk that the snapshot group was created from, thereby restoring the data on the cloud disk to the state when the snapshot was created.

The snapshot data can only be rolled back to the source cloud disk, not to other cloud disks.
Scenes where rollback can be performed:

- If the source cloud disk is in the state of To be mounted (that is, not attached to a CVM), you can directly perform the rollback operation.
- If the source cloud disk is already attached to a CVM, the corresponding CVM needs to be shut down to perform the rolling back operation.

Operation Steps

Rolling Back Snapshots Using the Console

1. Log in to the CBS console.
2. In the left sidebar, choose Snapshot > Snapshot List.
3. On the Snapshot Group tab, click Roll Back in the row of the target snapshot.

   > Note:
   >
   > The data on the source cloud disk will be rolled back to the state when the snapshot was created. Data after this time will be cleared, so proceed with caution.

4. On the rollback data page, confirm rollback information and click OK to start rollback.

Rolling Back Snapshots Using the API

You can use the ApplySnapshotGroup API to roll back snapshot groups.

Creating Cloud Disks from Snapshot Groups

Scenarios

Snapshots are an important method for data sharing and migration. A cloud disk created from a snapshot contains all the data in the snapshot. You can use a snapshot to create a cloud disk with a capacity equal to or greater than that of the snapshot.

- When you use a snapshot to create a data disk of the same size, the new data disk does not need to be initialized. You can mount it directly and right-click > Disk Management. After it is connected to the CVM in Disk Management, it can be read and written normally.
- When you use a snapshot to create a data disk with a capacity larger than the snapshot, the system only completes the disk expansion at the block device level and does not extend the file system or convert the partition format automatically. After a new data disk is attached, only the file system and data from the source snapshot can be used, and the new disk space cannot be directly used. You need to manually extend the file system or even convert the partition format.
  For example, use a data disk snapshot with an MBR partition format and a capacity of 1TB to create a new 3TB data disk. Since the maximum disk space supported by MBR is 2TB, you need to format the data disk and repartition it using GPT, which deletes the original data. Therefore, please operate with caution according to your actual needs.
  This document describes how to create a cloud disk from a snapshot group on the snapshot group list page. Additionally, when creating a cloud disk, you can configure the snapshot parameter to specify the corresponding snapshot for creating a cloud disk.

Operation Steps

Creating Cloud Disks from Snapshot Groups Using the Console

1. Log in to the CBS console.
2. In the left sidebar, select Snapshot > Snapshot List.
3. On the Snapshot Group tab, select More > Create Cloud Disk in the row containing the target snapshot group.
4. In the Purchase Data Disk dialog box, set the following parameters:

Parameter	Parameter Description
Disk	This parameter is optional.

Parameter	Parameter Description
name	The name of cloud disk supports up to 20 characters. After the cloud disk is created, it can be modified. - Create a single cloud disk: The disk name is the name of the cloud disk. - Create cloud disks in batches: When multiple cloud disks are created at a time, the disk name is the prefix of the cloud disk name. The final cloud disk name is in the format of disk name_number, from hard disk name_0 to hard disk name_49.
Project	This parameter is required. When creating a cloud disk, you can set the project it belongs to. The default project is Default Project.
Tag	This parameter is optional. You can bind tags to cloud disks when creating them. Tags are used to identify cloud resources and enable classification and search of cloud resources.

5. The default billing mode is Pay-as-you-go. You can view the newly created fee and click OK.
6. You can check the created cloud disk on the Cloud Block Storage page. The new elastic cloud disk is in the To be mounted state. To attach it to a CVM instance in the same AZ, see Attaching Cloud Disks .

Creating Custom Images from Snapshot Groups

Scenarios

You can create a new custom image using a snapshot group that contains a system disk.

Operation Steps

Creating Custom Images from Snapshot Groups Using the Console

1. Log in to the CBS console.
2. In the left sidebar, choose Snapshot > Snapshot List.
3. On the Snapshot Group tab, select More > Create Custom Image in the row containing the target snapshot group.
4. In the "Create Custom Image" dialog box, set the image name, and click OK.
5. View the created custom images on the Images > Custom Image page.
   You can create new instances using custom images.

Deleting Snapshot Groups

Scenarios

When a snapshot is no longer needed, you can delete it to release virtual resources.

Notes

1. When a snapshot group is deleted, only the data unique to the snapshots it contains is deleted, and the cloud disks from which the snapshots were created will not be affected.
2. Deleting a snapshot created earlier for a cloud disk does not affect your continued use of a snapshot created later.
3. When a snapshot group is deleted, all snapshot data in the snapshot group are deleted, and deleted data is irrecoverable. As the deleted snapshot is irrecoverable, please operate with caution.

Operation Steps

Deleting Snapshots Group Using the Console

1. Log in to the CBS console.
2. In the left sidebar, choose Snapshot > Snapshot List.
3. On the Snapshot Group tab, click Delete in the row of the snapshot group to be deleted. If your snapshot has been associated with a related image, you can choose to delete the custom image together or keep it.

Deleting Snapshots Using the API

You can use the DeleteSnapshotGroup API to delete snapshot groups.

Monitoring and Alarming

Providing a good monitoring environment for cloud disks is an important part of maintaining high data reliability. You can use Cloud Monitor service to monitor the cloud disk mounted on CVM instance to view the metric data of cloud disk and analyze and set alarms related to cloud disk. Cloud monitor collects the raw data of hard disk from the running CVM instance and displays the data in a readable chart form. The statistical data is saved for one month by default. You can observe the situation of disks in different time periods within a month to better understand information such as usage, reading, and writing.

You can obtain data through the cloud monitor console or cloud monitor API. For more details, see Getting the Monitoring Data of Specified Metrics and Get monitoring views and reports.

Currently, Cloud Monitor provides the following monitoring metrics for CBS:

Metric Name	Calculation Method	Meanings of Metrics	Unit	Statistical Granularity (Period)
Disk Read IOPS	The average value of hard disk reading IOPS in the statistical period	Frequency of IO reading from hard disk to memory per second	Frequency	10s, 60s, and 300s
Disk Read Traffic	The average value of hard disk reading throughput in the statistical period	Rate of data reading from hard disk to memory	KB/s	10s, 60s, and 300s
Disk Write IOPS	The average value of hard disk writing IOPS in the statistical period	Frequency of IO writing from memory to hard disk per second	Frequency	10s, 60s, and 300s
Disk Write Traffic	The average value of hard disk writing throughput in the statistical period	Rate of data writing from memory to hard disk	KB/s	10s, 60s, and 300s
IO Await	The average value of ioawait of hard disks in the statistical period	Percentage of time when the CPU is idle and there are still I/O requests that are not completed during the sampling period	ms	10s, 60s, and 300s
IO Svctm	The average value of svctm of hard disks in the statistical period	IO service time	ms	10s, 60s, and 300s

Metric Name	Calculation Method	Meanings of Metrics	Unit	Statistical Granularity (Period)
IO %util	The average value of io_util of hard disks in the statistical period	Ratio of time when the disk has IO operations (i.e., non-idle time)	%	10s, 60s, and 300s

Tag
Using Tags to Manage Cloud Disks

Scenarios

Tag is a key-value pair to identify cloud resources. Tags can help you conveniently classify and manage Cloud Block Storage (CBS) resources from various dimensionalities (such as business, purpose, owner.). It should be noted that Converge Cloud does not use the tags you set. The tags are only used for your management of server resources.

Use Limits

You need to pay attention to the following use limits of tags:

- Quantity Limit: Each cloud resource can have up to 50 tags.
- Tag Key Limit: It can only be digits , letters , or +=.@- , and the maximum length of the tag key is 255 characters.
- Tag Value Limit: It can only be null character strings or digits , letters , and +=.@- , and the maximum length of the tag value is 127 characters.

Directions and Cases

Case Description

Case: A company purchased 6 cloud disks. The information regarding the usage departments, business scopes, and responsible persons for these 6 cloud disks is as follows.

Cloud Disk instance-id	Department	Business Scope	Owner
ins-abcdef1	E-commerce	Marketing Campaigns	Tom
ins-abcdef2	E-commerce	Marketing Campaigns	Harry
ins-abcdef3	Game	Game A	Jane
ins-abcdef4	Game	Game B	Harry
ins-abcdef5	Entertainment	Post-production	Harry
ins-abcdef6	Entertainment	Post-production	Tom

Taking ins-abcdef1 as an example, we can add the following three sets of tags to this cloud disk:

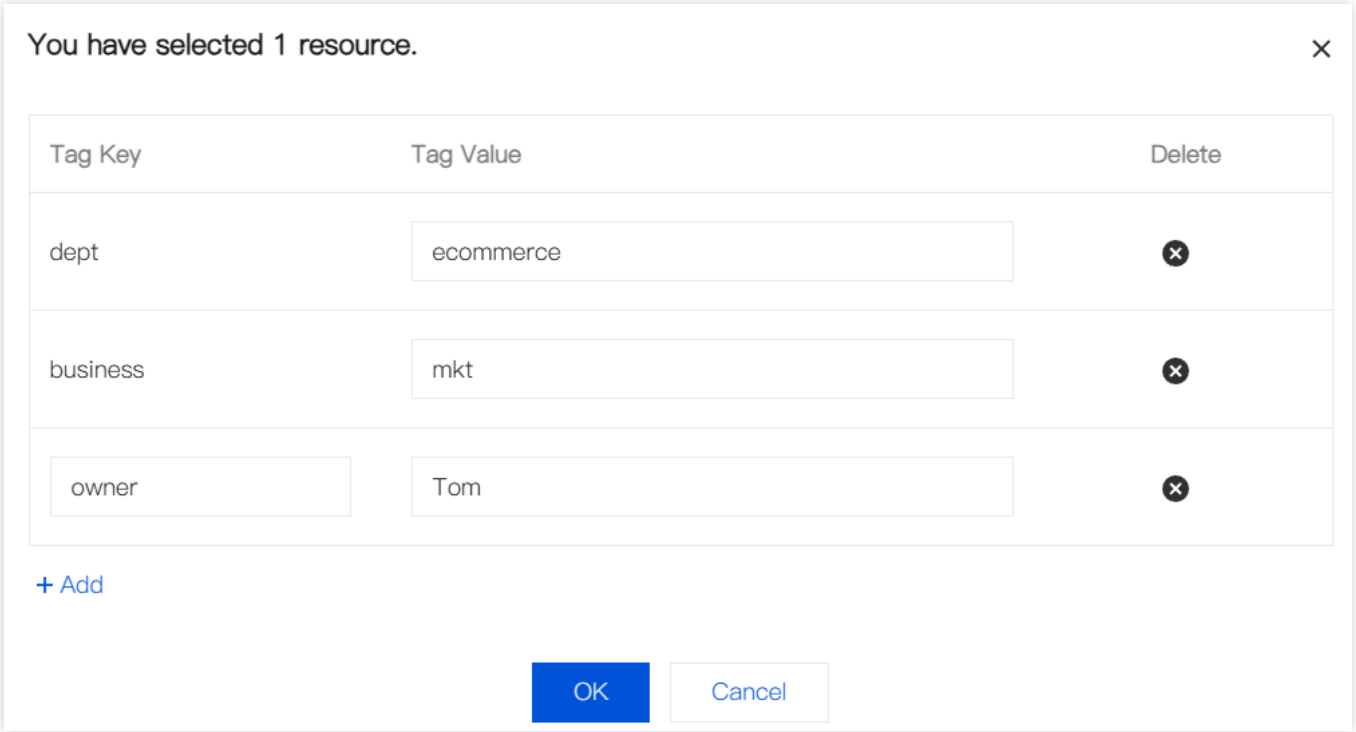
Tag Key	Tag Value
dept	ecommerce
business	mkt
owner	zhangsan

Similarly, we can set tags according to usage department, business scope, and the responsible person for other cloud disks.

Setting Tags in the CBS Console

After the tag keys and values are designed as detailed above, you can log in to the CBS console to set tags.

1. Log in to the CBS console.
2. On the cloud disk management page, select the cloud disk whose tag you want to edit and click More > Modify Tag.
3. Set the tag in the pop-up You have selected 1 resource window.
   For example, add three sets of tags to the cloud disk ins-abcdef1.



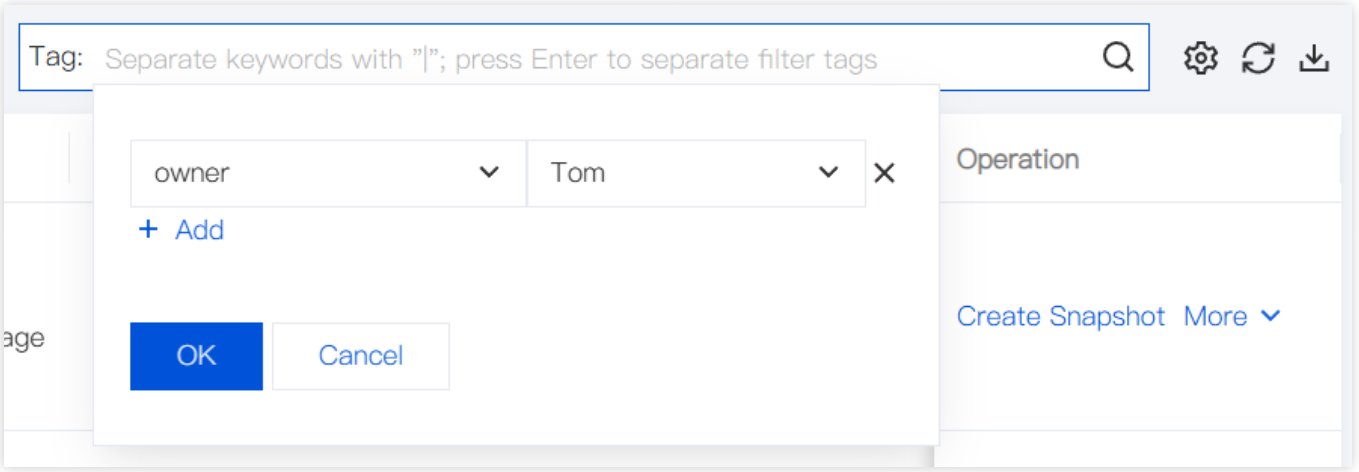
4. Click OK. Then, a prompt pops up indicating that the modification is successful.

Filtering Cloud Disks by Tag

You can filter cloud disks with a certain type of tag through the following operations.

1. Select Tag in the search box.

2. Enter the tag key and tag value after Tag: and click to search.
   For example, if you want to filter out cloud disk resources with the responsible person as Tom, you can input Tom to filter.



Editing Tags

Scenarios

This document guides you in editing tags for resources.

Use Limits

When you edit tags, you need to be aware of the following limits:

- Quantity Limit: Each cloud resource can have up to 50 tags.
- Tag Key Limit: It can only be digits , letters , or +=.@- , and the maximum length of the tag key is 255 characters.
- Tag Value Limit: It can only be null character strings or digits , letters , and +=.@- , and the maximum length of the tag value is 127 characters.

Prerequisites

Logged in to the Cloud Block Storage (CBS) console.

Operation Steps

Modifying Tags for a Single Cloud Disk

1. On the cloud disk management page, select the cloud disk whose tag you want to edit and click More > Modify Tag.
2. In the pop-up window, you can add, modify, or delete tags based on actual needs.

Modifying Tags for Multiple Cloud Disks

> Note:
>
> Batch editing of tags is supported for up to 20 resources.

1. On the cloud disk management page, select the cloud disk whose tag you want to edit and click More > Modify Tag at the top.
2. In the pop-up window, you can add, modify, or delete tags based on actual needs.

Best Practices
How to Measure the Cloud Disk Performance

Measuring Metrics

Since different applications have different workloads, if the I/O requests provided are not sufficient to fully utilize the cloud disk, the maximum performance of the cloud disk may not be achieved.

Generally, the following metrics are used to measure cloud disk performance:

- IOPS: Number of reads/writes per second with units in counts. The underlying driver type of the storage device determines different IOPS.
- Throughput: the amount of read and write data per second in MB/s.
- Delay: the time from sending I/O operation to receiving and confirmation in seconds.

Testing Tool

FIO is a tool for testing disk performance, which is used for stress tests on hardware and verification. This document takes FIO as an example.

When using FIO, it is recommended to use libaio's I/O engine for testing. Install FIO and libaio yourself.

> Warning:
>
> - Do not perform an FIO test on the system disk to avoid damaging important files of the system.
> - To avoid data corruption caused by metadata corruption of the underlying file system, do not test on the business data disk.
> - Ensure that there is no mounting configuration for the hard disk under test in the `/etc/fstab` file configuration item. Otherwise, the CVM will fail to start.

Suggestions for Test Object

- It is recommended to perform the FIO test on a free hard disk that has not saved important data and rebuild the file system of the tested hard disk after the test.
- When testing the performance of hard disks, it is recommended to directly test raw data disks (such as /dev/vdb).
- When testing file system performance, it is recommended to specify a specific file test (such as /data/file).

Test Examples

The test formulas for different scenes are basically the same, except for rw, iodepth, and bs (block size). For example, the optimal iodepth for each workload varies depending on how sensitive your particular application is to IOPS and delay.

Parameter description:

Parameter Name	Description	Example of Value
bs	Block size per request. Values include 4k, 8k, and 16k.	4k
ioengine	I/O engine. Linux's asynchronous I/O engine is recommended.	libaio
iodepth	Requested I/O queue depth.	1
direct	Specify the direct mode. - True (1) indicates that the O_DIRECT marker B transparency is specified, the I/O cache is ignored, and data is written directly. - False (0) indicates that the O_DIRECT marker B transparency is not specified. The default value is True (1).	1
rw	Read and write mode. The values include read, write, randread, randwrite, randrw, and rw, read/write.	read
time_based	Specifies that the time mode is used. There is no need to set this parameter value as long as the FIO runs based on time.	N/A
runtime	Specifies the test duration, i.e., FIO run time.	600
refill_buffers	The FIO refills the I/O buffer with each submission. The default setting is to fill and reuse the data only initially.	N/A
norandommap	When performing random I/O, FIO will overwrite each block of the file. If this parameter is given, a new offset will be selected without viewing the I/O history.	N/A
randrepeat	Whether the random sequence is repeatable. - True (1) indicates that the random sequence is repeatable. - False (0) indicates that the random sequence is not repeatable. The default value is True (1).	0
group_reporting	When multiple jobs are running concurrently, the statistics of the whole group will be printed.	N/A
name	Name of job.	fio-read
size	Addressing space for I/O testing.	100GB
filename	The test object, i.e., the name of the disk device to be tested.	/dev/sdb

Common use cases are as follows:

- bs = 4k iodepth = 1: random read/write test, which can reflect the delay performance of hard disk.
  Execute the following command to test the random read delay of hard disks:

  ```
  fio -bs=4k -ioengine=libaio -iodepth=1 -direct=1 -rw=randread -time_based -runtime=600  -refill_buf
  fers -norandommap -randrepeat=0 -group_reporting -name=fio-randread-lat --size=10G -filename=/
  dev/vdb
  ```

  Execute the following command to test the random write delay of hard disks:

  ```
  fio -bs=4k -ioengine=libaio -iodepth=1 -direct=1 -rw=randwrite -time_based -runtime=600  -refill_bu
  ffers -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 the Cloud SSD.

  ```
  fio --bs=4k --ioengine=libaio --iodepth=1 --direct=1 --rw=randrw --time_based --runtime=100 --refil
  l_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 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
  ```

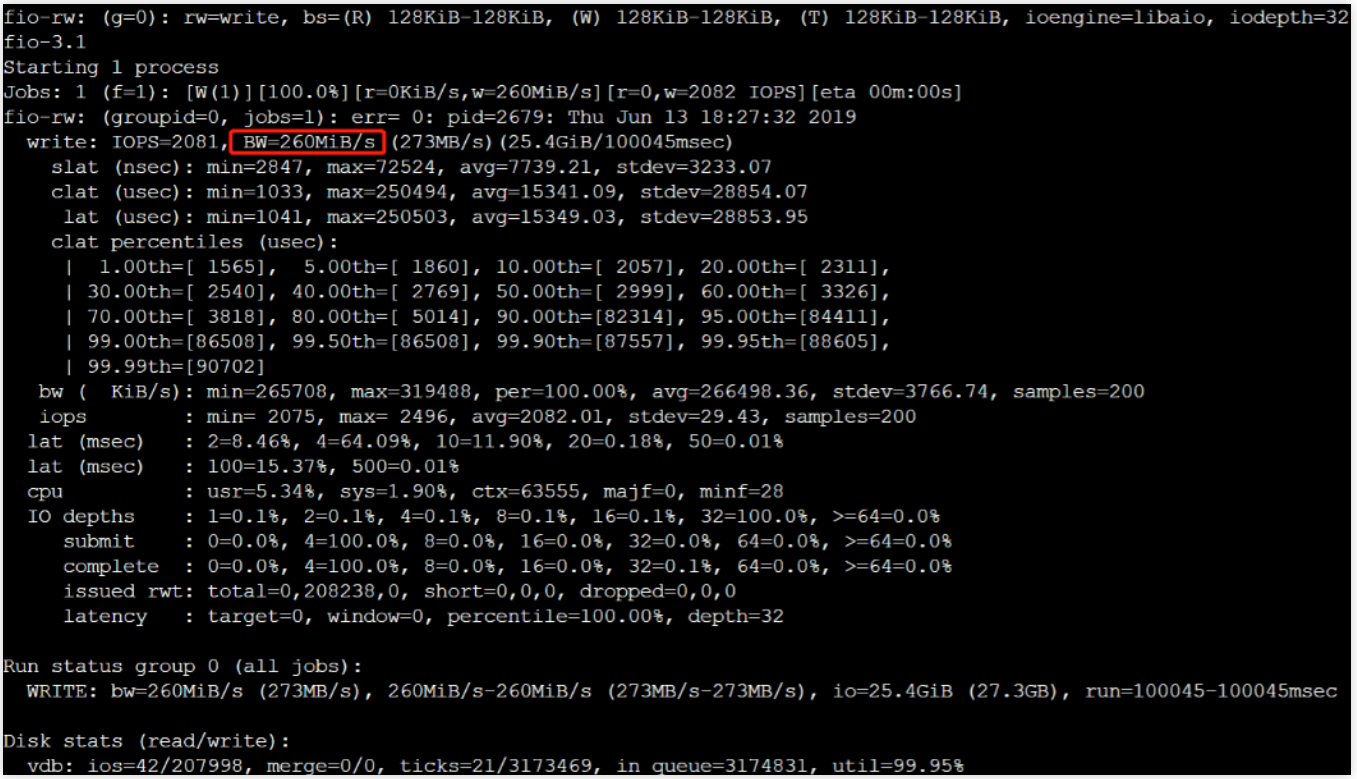
  Execute the following command to test the sequential write throughput bandwidth of the hard disk:

  ```
  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 the Cloud SSD.

  ```
  fio --bs=128k --ioengine=libaio --iodpth=32 --direct=1 --rw=write --time_based --runtime=100 --refill_buffers --norandommap --randrepeat=0 --group_reporting -rw --size=1G --filname=/dev/vdb
  ```

The test results are shown in the following figure:



- bs = 4k iodepth = 32: random read/write test, which can reflect the IOPS performance of 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_b
  uffers -norandommap -randrepeat=0 -group_reporting -name=fio-randread-iops --size=10G -filenam
  e=/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_b
  uffers -norandommap -randrepeat=0 -group_reporting -name=fio-randwrite-iops --size=10G -filenam
  e=/dev/vdb
  ```

  Test the random read IOPS performance of a Cloud SSD. As shown in the figure below:

```
[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 rwt: 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 LVM Logic Volumes with Multiple Elastic Cloud Disks

Introduction to LVM

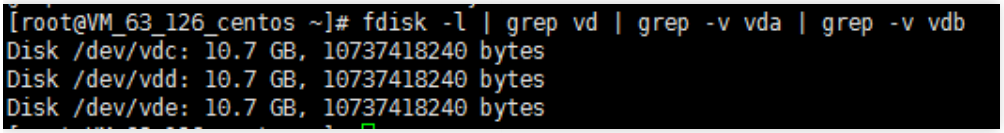
Logical Volume Manager (LVM) divides disks or partitions into Physical Extents (PE) units of the same size by establishing a logical layer above hard disks and partitions. Different disks or partitions can be assigned to the same volume group (VG), logical volume (LV) can be created on VG, and file systems can be created on LV.

Compared with directly using disk partitions, LVM has the strength of flexibly adjusting the capacity of 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 purchase three 4-TB elastic cloud disks and use LVM to create a nearly 12-TB large file system.
- You can dynamically resize your LV without having to repartition the disk.
  When the space of the LVM volume group cannot meet your requirements, you can purchase an elastic cloud disk separately, attach it to the corresponding CVM, and then add it to the LVM volume group for expansion.

Building LVM

> Note:
>
> This document takes how to create a dynamically resizable file system using 3 elastic cloud disks via LVM as an example. As shown in the figure below:



Step 1: Creating a Physical Volume (PV)

1. Log in to the Linux CVM as root user.

2. Execute the following command to create a PV.

pvcreate <disk path 1> ... <Disk path N>

This document takes /dev/vdc , /dev/vdd , and /dev/vde as examples to execute the following command:

pvcreate /dev/vdc /dev/vdd /dev/vde

The successful creation is shown in the following figure:

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

3. Execute the following command to view the PVs in the system.

lvmdiskscan | grep LVM

```
[root@VM_63_126_centos ~]# lvmdiskscan  | 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
```

Step 2 Creating a VG

1. Execute the following command to create a VG.

vgcreate [-s <specified PE size>] <volume group name> <physical volume path>

This document takes creating a volume group named lvm_demo0 as an example to execute the following command:

vgcreate lvm_demo0 /dev/vdc /dev/vdd

The successful creation is shown in the following figure:

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

When the message Volume Group "" Successfully Created is prompted, the VG has been created successfully.

- After the VG is created, execute the following command to add a new PV to the VG.

  vgextend Volume group name New physical volume path

  The following figure shows the successful addition:

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

- After the VG is created, you can execute vgs and vgdisplay commands to view the VG information in the system. As shown below:

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

Step 3 Creating an LV

Execute the following command to create an LV.

lvcreate [-L <logical volume size>][-n <logical volume name>] <VG name>

This document takes the creation of an 8 GB LV named lv_0 as an example to execute the following command:

lvcreate -L 8G -n lv_0 lvm_demo0

The successful creation is shown in the following figure:

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

> Note:
>
> Execute the pvs command to view that only 8 GB is used by /dev/vdc at this time. As shown below:

```
[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
```

Step 4 Creating and Mounting the File System

1. Execute the following command to create a file system on the created LV.

   mkfs.ext3 /dev/lvm_demo0/lv_0

2. Execute the following command to mount the file system:

   mount /dev/lvm_demo0/lv_0 vg0/

   The following figure shows the successful mounting:

```
[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)
```

Step 5 Dynamically Scale out the Size of LV and File System

> Note:
>
> LV capacity can be dynamically scaled out only when there is surplus VG capacity. After scaling out the LV capacity, it is necessary to scale out the size of the file system created on the LV.

1. Execute the following command to scale out the LV size:

   lvextend [-L +/- <increase or decrease capacity>]<logical volume path>

   This document takes scaling out the capacity of 4 GB to LV lv_0 as an example to execute the following command:

   lvextend -L +4G /dev/lvm_demo0/lv_0

   The following figure shows the successful scale-out:

```
[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
```

> Note:
>
> Execute the pvs command to view that /dev/vdc is fully used, and /dev/vdd is used 2GB. As shown below:

```
[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
```

2. Execute the following command to scale out the file system:

   resize2fs /dev/lvm_demo0/lv_0

   The following figure shows the successful scale-out:

```
[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
```

   After the scale-out is successful, execute the following command to check whether the capacity of the LV has changed to 12 GB.

   df -h

Troubleshooting
Disk Not Mounting upon Linux CVM Restart

Phenomenon

The file system has been created and configured for automatic attachment to the cloud disk on the Linux CVM, but after the CVM restarted, the cloud disk is not automatically attached.

Possible Causes

- Cause 1: The automatic attachment of the cloud disk is not configured in the system's fstab configuration file.
- Cause 2: The fstab configuration file is not properly configured.
  For example, if it is configured to use the device name for automatic mounting, but the device name changes when the CVM restarts, resulting in that the system cannot be started normally.

Solutions

- Solution 1 (corresponding to cause 1):
  By reconfiguring the /etc/fstab file using the following configuration methods, you can enable automatic attachment of the corresponding cloud disk after the CVM restarts.
  - Use the soft links from disk (recommended)
  - Use universally unique identifier (UUID) of the file system
  - Use the device name (not recommended)
    For details, see Configuring the /etc/fstab File.
- Solution 2 (corresponding to cause 2):
  Log in to the Linux CVM through VNC and enter the single-user mode. Restore the /etc/fstab configuration file in single-user mode, and reconfigure it. For details, see Repairing the /etc/fstab File.

Processing Procedures

Configuring the /etc/fstab File

1. Log in to the Linux CVM as instructed in Using the Standard Login Method to Log In to Linux Instances (Recommended).
2. Select the configuration method and obtain the corresponding information.

Using the Soft Link of the Elastic Cloud Disk (Recommended)

Configuration Mode Analysis

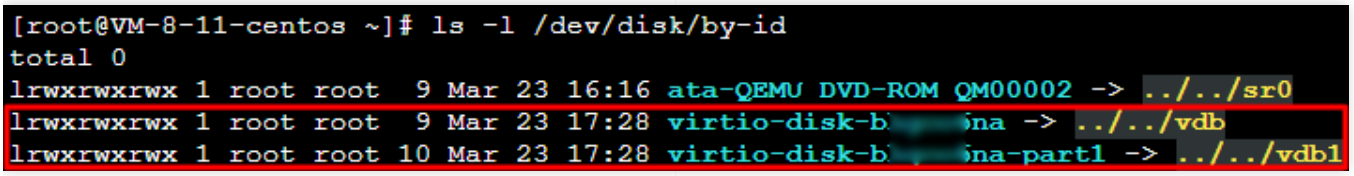
- Advantages: The soft link of each elastic cloud disk is fixed and unique, remaining unchanged even during operations such as unmounting and remounting, or formatting partitions.
- Disadvantages: Only elastic cloud disks have soft links, and the system cannot detect formatting operations on partitions.

Ways to Obtain Information

Execute the following command to check the soft link of elastic cloud disks.

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

The returned result is as shown in the following figure:



Using the UUID of the File System

Configuration Mode Analysis

Automatic mount settings may be invalidated due to file system UUID changes.
For example, the UUID of a file system will change after it is reformatted.

Ways to Obtain Information

Execute the following command to check the UUID of the file system.

```
blkid /dev/vdb1
```

The returned result is as shown in the following figure:



Using the Device Name (Not Recommended)

Configuration Mode Analysis

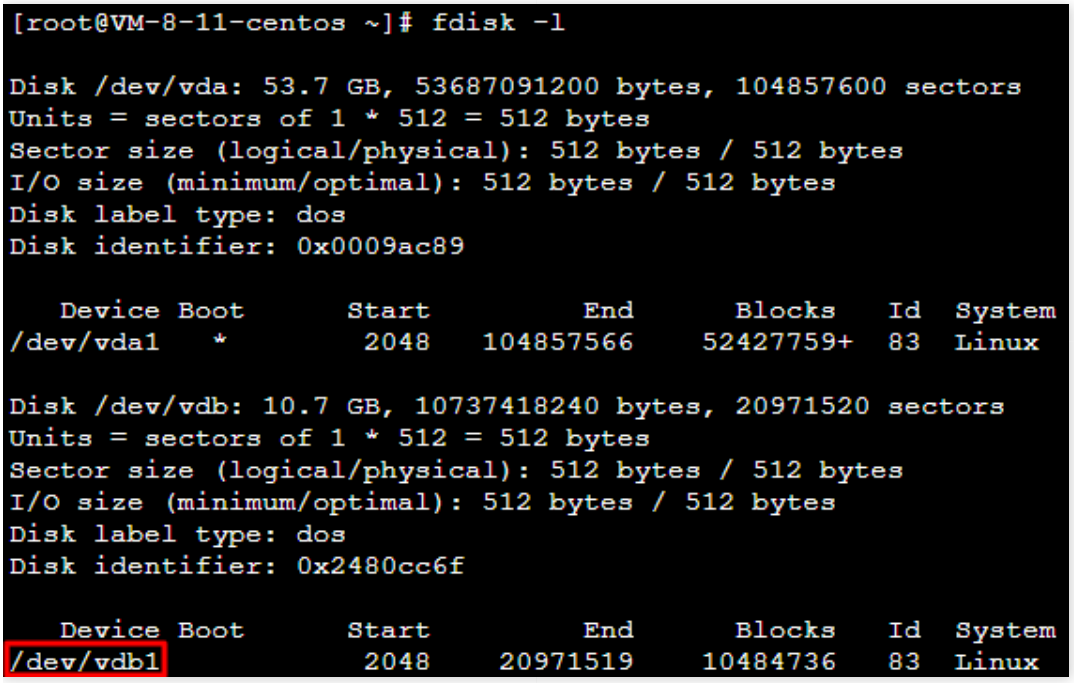
The automatic mount setting may be invalidated due to the change of device name.
For example, when the elastic cloud disks on the CVM are unmounted and remounted during data migration, the name may change when the operating system recognizes the file system again.

Ways to Obtain Information

Execute the following command to check the device name.

```
fdisk -l
```

The returned result is as shown in the following figure:



1. Execute the following command to back up the /etc/fstab file. This document takes backup to the /home directory as an example:

   ```
   cp /etc/fstab /home
   ```

2. Execute the following command to open the /etc/fstab file with the VI editor.

   ```
   vi /etc/fstab
   ```

3. Press i to enter edit mode, and enter the following configuration information on a new line at the end of the file.

   <Device information> <Mount point> <File system format> <Installation options of file system> <File system dump frequency> <File system check sequence at startup>

   You can configure by referring to the following example according to the configuration method selected in Step 2.

- Configuration example of using the soft link from elastic CBS (recommended):

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

- Configuration example of using UUID of file system:

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

- Configuration example of using device name (not recommended):

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

4. Press Esc to enter :wq, and then press Enter to save the settings and exit the editor.
5. Execute the following command to check whether the /etc/fstab file is written successfully.

   ```
   mount -a
   ```

   If the returned result is as shown in the following figure, the file has been written successfully. The file system is automatically mounted when the CVM starts, and you can restart the CVM for verification.

   ```
   [root@VM-8-11-centos ~]# mount -a
   [root@VM-8-11-centos ~]#
   ```

Repairing the /etc/fstab File

1. Log in to the CVM via VNC as instructed in Using VNC to Log in to Linux Instances.
2. Enter single-user mode. For details, see Setting Linux CVM to Enter Single-User Mode.
3. Execute the following command to back up the /etc/fstab file. This document takes backup to the /home directory as an example:

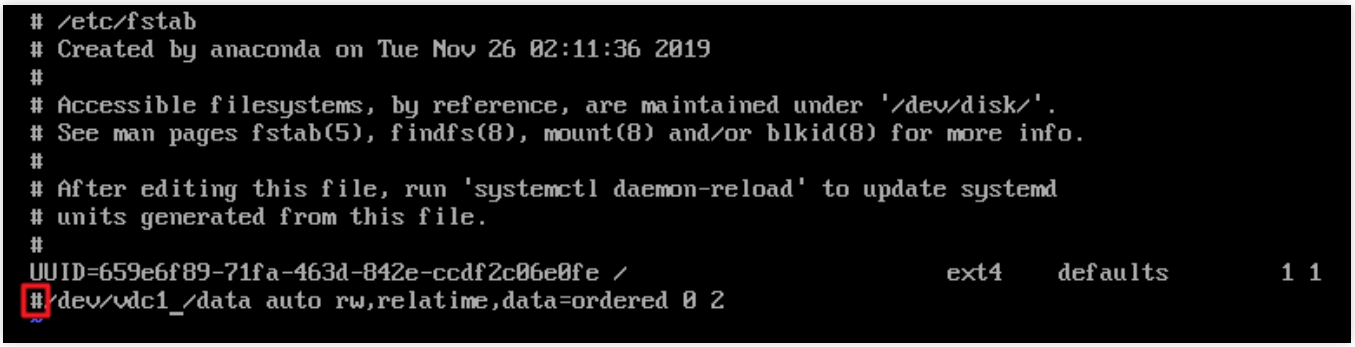
   ```
   cp /etc/fstab /home
   ```

4. Execute the following command to open the /etc/fstab file with the VI editor.

   ```
   vi /etc/fstab
   ```

5. Press i to enter the edit mode, move the cursor to the beginning of the line with configuration error, and enter # to configure a note for this line. As shown below:

> Note:
>
> This row is configured with automatic data disk mounting, but due to incorrect configuration, the actual CVM fails to mount when it restarts.



6. Press Esc to enter :wq, and then press Enter to save the settings and exit the editor.

7. Enter exit in the command line of single-user mode to exit the single-user mode.

8. Wait for the CVM to restart, and then log in to the CVM.

9. Refer to Configure /etc/fstab File and re-configure.

FAQs
Performance-related Issues

How to Measure the Performance of a Cloud Disk?

The following metrics are generally used to describe the performance of storage devices:

- IOPS: Number of reads/writes per second with units in counts. The underlying driver type of the storage device determines different IOPS.
- Throughput: the amount of read and write data per second in MB/s.
- Delay: the time from sending I/O operation to receiving and confirmation in seconds.

How to Test Disk Performance?

It is recommended to use FIO to perform stress testing and verification on the cloud disk. For specific operations, see How to Measure the Performance of a Cloud Disk.

Will the I/O Size Read and Written by the Application Affect My IOPS Performance?

Yes. For a given resource, the IOPS you get depends on the I/O size of your application's read-and-write operations. Usually, the IOPS performance of a hard disk can be fully applied when reading and writing in small blocks (for example, I/O size is 256 KB).

Will the I/O Size Read and Written by the Application Affect My Throughput Performance?

Yes. For a given resource, the throughput you get depends on the I/O size of your application's read-and-write operations. Usually, the throughput performance of a hard disk can be fully applied when reading and writing in large blocks (for example, I/O size is 1 MB).

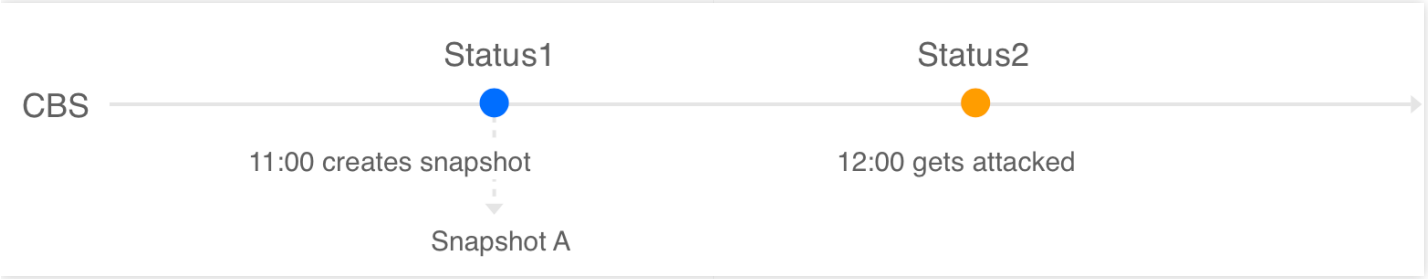
Can Multiple Hard Disks Be Logically Combined into One for Better Performance?

Yes, you can stripe multiple cloud disks attached to a CVM and balance the I/O load across multiple disks to enhance I/O parallelism and achieve higher performance beyond that of a single disk. For details, see Building a Logical Volume in LVM with Multiple Elastic Cloud Disks.

Snapshot-related Issues

Why Is Snapshot Protection Still Needed When CBS Already Uses a Triple-Replica Mechanism to Protect Data Security?

When data exceptions occur at the logical layer, such as accidental data deletion by users, data corruption caused by system virus infection, or file system exceptions, the actual stored three data replicas will be modified simultaneously, making it impossible to recover normal historical data. If snapshots have been created at historical time points, you can use these snapshots to restore the data to its normal state at that historical time point.



For example, the administrator created a snapshot A for a cloud disk at 11:00, and the cloud disk was infected by a virus at 12:00, making the data unusable. At this point, the data on the three replicas had also been updated to state 2, making data recovery impossible. Only by using the snapshot A created at 11:00 can the data be restored to the uninfected state 1.

What is the difference between the snapshot and image? What is the connection between the two?

Difference between the snapshot and image

For example, if no data disk is mounted to an instance and all data is written to the system disk, the system disk cannot be protected by only creating an image. Because the image cannot be created regularly, if the system disk data is damaged, it can only be traced back to the initial data when the image was created, which cannot protect the data. The detailed differences are shown in the following table:

Name	Snapshot	Mirror
Nature	A backup of cloud disk data at a certain point	Template of Cloud Virtual Machine (CVM) software configuration (operating system, pre-installed program, etc.)
Applicable Scenario	- Regularly back up the important business data - Back up data before major operations	- Back up system that does not change in the short term. - Deploy applications in batch - Migrate System

Name	Snapshot	Mirror
	- Produce multi-replica application of data	

Connection Between the Snapshot and Image

- The underlying data storage of the image uses the CBS snapshot service, and the associated snapshot generated by the custom image is charged according to the actual capacity of the snapshot. You can view the total capacity of the snapshot in Snapshot Overview.
- An image is bound to its associated snapshot. To delete the associated snapshot, click Delete Image first. You can view the snapshot information associated with a custom image in Image Details.
- Snapshots can be created in both system disk and data disk. However, you can only use the system disk snapshot to create a custom image, and then use the created image to create one or more instances with the same environment.
- When a shared image is used, only the creator account of the image needs to pay for the associated snapshot, and the sharee does not need to pay.

What is the Regular Snapshot?

A scheduled snapshot is a snapshot automatically created by the system for a cloud disk associated with the scheduled snapshot policy you created. To use this feature, you need to first create a scheduled snapshot policy and associate it with a cloud disk. For specific operations, see Scheduled Snapshots.

Are There Any Regional Limits on the Snapshot?

All availability zones are supported by the current snapshot feature.

Do I Need To Uninstall Hard Disks or Interrupt All Read/Write Operations To Create Snapshots?

No. You can create snapshots in real time while the cloud disk is connected and in use, without affecting your normal business. However, snapshots can only capture data that has been written to the cloud disk and do not include data cached in memory by applications or the operating system. To ensure that all application data is captured in the snapshot, it is recommended that you completely suspend I/O operations on the hard disk before creating the snapshot. For cloud disks used as system disks, it is recommended that you shut down the CVM before creating a snapshot to create a more complete snapshot.

Does Creating Snapshots Affect the Performance of the Hard Disk?

Creating a snapshot consumes a small amount of I/O on the cloud disk. It is recommended to perform snapshot operations during relatively idle business periods.

How Long Does It Take To Create a Snapshot Until the Snapshot Become Available?

The time taken to create a snapshot is affected by various factors such as the amount of data written to the cloud disk and the underlying read-write conditions, making it difficult to predict. However, creating a snapshot does not affect the normal operation of the hard disk.

How to Create a Cloud Disk from Snapshots?

For specific operations, see Creating a Cloud Disk from a Snapshot.

How To Roll Back a Snapshot?

For specific operations, see Rolling Back Data from a Snapshot.

Do I Need To Shut Down for Snapshot Rollback?

- For a cloud disk that is already attached to a CVM, the CVM needs to be shut down before performing the rollback.
- For a cloud disk that is not attached to a CVM, you can directly perform the rollback operation.

Can a Cloud Disk Be Restored by Reading a Previous Snapshot?

Yes. You can use an existing snapshot at any point in time to restore data, regardless of the snapshot time.

Can the Source Snapshot Be Deleted When It Is Copied?

Not supported. You can delete it only after the copy is complete.

Does the Previous Disk Association Still Exist After the Copy Is Completed?

After the cross-regional copy, there will be no association between the new snapshot and the source disk of the source snapshot, and the rollback feature of the new snapshot is unavailable after copying.

How Do I Delete a Snapshot?

- Snapshots on the CBS can be deleted directly in the console or through API. For specific operations, see Delete Snapshot.
- For the associated snapshot of a custom image, you need to delete the custom image before deleting the snapshot.

Usage-related Issues

What Are the Applicable Scenarios for Different Cloud Disks?

- After purchasing a CVM, if you find that the disk space is insufficient, you can purchase and attach an elastic cloud disk as a data disk to meet storage needs.
- When purchasing a CVM, if you do not want to buy additional data disks, you can purchase an elastic cloud disk when you have storage needs and attach it as a data disk.
- Server A has 10 GB of important data stored in an elastic cloud disk. When the data needs to be shared with server B, the disk can be directly detached from server A and then attached to server B.
- When a single cloud disk with the maximum specification cannot meet storage needs, you can purchase multiple cloud disks of the same capacity and configure an LVM logical volume to provide larger disk space.
- When the I/O performance of a single disk cannot meet business requirements, you can purchase multiple cloud disks and configure RAID 0, RAID 10, or other levels to enhance the I/O performance.

How to Choose a Cloud Disk Type?

Before selecting the type of hard disk, confirm what type of work data you need to store:

- For scenarios with relatively infrequent requests, such as system logs, enterprise work files, data repositories, small blogs, and BBS, it is recommended that you use Premium Disk to reduce costs.
- For common scenarios such as small and medium-sized databases, Web/App applications, and transactional workloads, we recommend you to use Premium Disks for better cost-effectiveness.
- For scenarios with high workload and high performance requirements, such as large core databases, OLTP business, and NoSQL databases, we recommend you to use a Cloud SSD for higher performance.

What Are the Must-Knows for Using Cloud Disks?

- For an independently purchased cloud disk, when configuring static file system information using fstab, you should use the file system's UUID or label as the file system identifier to prevent the kernel name of the cloud disk in the CVM from changing due to repeated attachments and detachments of multiple cloud disks on the same CVM.
- If a cloud disk expires before the CVM does, then that cloud disk will be throttled, detached, or even recycled within a certain period after expiration. Renew the cloud disk in time to prevent impact on your business.
- If detaching a cloud disk from the CVM does not seriously affect your critical business, you can consider using the nofail option when configuring fstab to prevent errors during system reboot caused by cloud disk detachment from the CVM.
- We recommend that you execute the command san policy=OnlineAll in diskpart before using a cloud

disk in the Windows operating system.
- When detaching a cloud disk from the Windows operating system, we recommend that you interrupt all read-write operations to the disk first and perform an offline operation.

How to Automatically Mount the Data Disk When Starting a New Instance Using Custom Image and Data Disk Snapshot?

For specific operations, see the Auto Attachment section in Attaching Cloud Disks.

How to Purchase a Cloud Disk?

You can create a cloud disk in the console or via API. For specific operations, see Creating a Cloud Disk.

Why Do I Fail to Find the CVM I Want to Attach to When Attaching a Cloud Disk?

Cloud disks cannot be attached across availability zones. Ensure that your CVM instance and cloud disk are in the same availability zone of the same region, and that your CVM has not been released.

Why Do I Fail to See the New Disk Space in the CVM's Operating System After Attaching a Cloud Disk?

Some Linux CVMs may not be able to recognize the elastic cloud disk. First, you need to enable the disk hot plug feature in the CVM. For details, see Enabling Hot-Swappable Disk Feature.

After manually attaching a cloud disk, you need to select and perform the corresponding subsequent operations to make it available.

Creation Model	Subsequent Operations
Create	- Cloud disk with a capacity of less than 2 TB: Initializing cloud disks (< 2 TB) - Cloud disks with a capacity of greater than or equal to 2 TB: Initializing cloud disks (≥ 2 TB)
Creating from Snapshot	- Cloud disks with a capacity that equals the snapshot capacity: No further operations are required. The cloud disk can be used directly after being attached. - Cloud disks with a capacity that is greater than the snapshot capacity and less than or equal to 2 TB, or cloud disks with a capacity that is greater than the snapshot capacity, with the snapshot capacity greater than 2 TB: - Mount to Windows CVM: scale out the partition and file system (Windows) - Mount to Linux CVM: scale out the partition and file system (Linux) - Cloud disks with a capacity of greater than 2 TB, with the snapshot capacity less than or equal to 2 TB: - If MBR partition is used in the snapshot: See Initializing Cloud Disks (≥ 2 TB) to repartition using GPT. This operation will delete existing data. - If the GPT partition is used in the snapshot: - Mount to Windows CVM: scale out the partition and file system (Windows)

Creation Model	Subsequent Operations
	- Mount to Linux CVM: scale out the partition and file system (Linux)

After Attaching a Cloud Disk, How to Partition and Format It?

For specific operations, see Initializing Cloud Disks (< 2 TB) or Initializing Cloud Disks (≥ 2 TB).

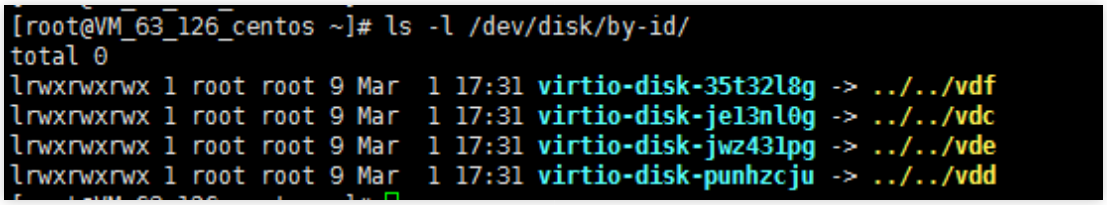
Is Multiple CVMs Accessing the Same Cloud Disk Supported?

Not supported currently. You can attach up to 32 cloud disks to the same CVM, but currently, multiple CVMs cannot share the same cloud disk simultaneously. Data sharing can only be achieved by detaching the cloud disk from CVM A and then attaching it to CVM B.

When Multiple Cloud Disks of the Same Size and Type Are Attached to the Same CVM, How Can I Distinguish Them in the Operating System?

- For a Linux operating system, you can execute the following command to view the correspondence between elastic cloud disks and device names:

  ls -l /dev/disk/by-id



- For the Windows operating system, you can view it by executing the following command:

  wmic diskdrive get caption,deviceid,serialnumber

  or

  wmic path win32_physicalmedia get SerialNumber,Tag



When I Detach a Cloud Disk, Will Data Be Lost?

Data on a cloud disk does not change due to attachment or detachment. To maintain data consistency, we strongly recommend:

- Under the Linux operating system, log in to the CVM instance and umount the disk. After the command is successfully executed, go to the console to unmount the disk.
- Under the Windows operating system, pause the read and write operations of all file systems on this disk before unmounting; otherwise, the data that has not been read or written will be lost.

How to Detach an Elastic Cloud Disk?

For specific operations, see Detaching a Cloud Disk.

After a Cloud Disk Is Successfully Purchased, Can Its Capacity Be Adjusted?

Yes. Cloud disk capacity can be adjusted, and it can be expanded, but not reduced.

What Are the Conditions for Scaling-out the File System?

Only cloud disks can be expanded. Local disks cannot be expanded. For specific operation instructions, see Introduction to Scaling Scenarios.

> Note:
>
> - It is strongly recommended that you take a snapshot before the scale-out to ensure data security.
> - If the maximum capacity of a cloud disk still cannot meet your business needs, you can use multiple elastic cloud disks to build a RAID group or an LVM logical volume.
> - 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.

How to Expand a Cloud Disk?

For a specific operation guide, see Introduction to Scaling Scenarios.

Does the CVM Support CPU/Memory Scale-out?

When the system disk is a cloud disk, the CVM supports CPU and memory adjustments.

What Should I Do If the Cloud Disk Partition Is in MBR Format and Cannot Be Expanded

Further?

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.

What Should I Do If the Cloud Disk Has Been Expanded to Its Maximum Capacity but Still Cannot Meet the Business Needs?

We recommend that you use multiple elastic cloud disks to build an LVM logical volume.

How to Build LV in LVM with Multiple Elastic Cloud Disks?

For specific operations, see Building a Logical Volume in LVM with Multiple Elastic Cloud Disks.

What Happens to the Data When a CVM Is Terminated?

- The lifecycle of the system disk completely follows that of the CVM. When the CVM is destroyed, the data stored in the system disk will also be terminated.
- The lifecycle of a data disk (that is, an elastic cloud disk) is independent of the CVM. You can freely choose whether to retain the elastic cloud disk and its data beyond the lifecycle of the CVM. Therefore, we recommend that you use an elastic cloud disk to store data that needs to be retained for a long time.

How to Recover a Cloud Disk after It Is Formatted?

The data on a cloud disk cannot be recovered after it is formatted. We recommend that you first create a snapshot before formatting it.

How to Delete a Cloud Disk?

- The lifecycle of the system disk follows that of the CVM and can only be deleted when Terminating Instances.
- The lifecycle of a data disk (that is, an elastic cloud disk) is independent of the CVM and can be deleted independently. For specific operations, see Terminating a Cloud Disk.

Can the System Disk Be Partitioned?

The system disk does not support partitioning.

Glossary

Elastic Cloud Disk

An elastic cloud disk is a cloud disk with an independent lifecycle (billing cycle) that can be freely attached to and detached from different CVM instances in the same region and availability zone. Currently, an elastic cloud disk cannot be attached to multiple CVM instances.

Non-Elastic Cloud Disk

A non-elastic cloud disk is a cloud disk that is created together with the CVM and whose lifecycle depends on the CVM. A non-elastic cloud disk does not support elastic attachment.

GPT

See Global Partition Table.

Rollback

It refers to the operation used to restore a program or data to its last correct status when program or data processing error occurs. Rollback includes program rollback and data rollback.

IOPS

Input/Output Per Second (IOPS) is the number of input/output operations (or read-write operations) per second, one of the key metrics for measuring hard disk performance. IOPS refers to the number of I/O requests that a system can process per unit of time, usually measured in the number of I/O requests processed per second. I/O requests are usually read or write data operation requests.
Traditional disks are essentially mechanical devices, such as FC, SAS, and SATA disks, with rotational speeds of 5,400/7,200/10,000/15,000 rpm. The key factor affecting disk performance is disk service time, the time it takes for a disk to complete an I/O request. Disk service time consists of three parts: seek time, rotational latency, and data transfer time.
Typically, a single mechanical hard disk with a rotational speed of 7200 rpm provides 75 to 150 IOPS, while a single mechanical hard disk with a rotational speed of 15000 rpm provides 175 to 210 IOPS. The specific values depend on factors such as access patterns (for example, sequential or random) and I/O size.

Snapshot Chain

Snapshot chain refers to the relationship chain composed of all snapshots on the same disk, and each node represents a snapshot of the disk.

MBR

See Master Boot Record Partitioning.

Guid Partition Table

The Guid Partition Table (GPT) is a partition structure of the physical hard disk, which is part of the Extensible Firmware Interface (EFI) standard and used to replace the MBR partition table of most existing disks.

Full Snapshot

It refers to the snapshot for disk data created for the first time, and it saves the full amount of disk data.

Sequential I/O

Sequential I/O refers to read/write operations that successively access data from adjacent addresses and is performed one by one on a logical block basis. With sequential I/O access, the seek time of hard disk is greatly reduced because the read-write heads rarely need to move to access the next block. Most services, such as data backup and log flow, generate sequential I/O.

Random I/O

Random I/O refers to that access address is not continuous, but randomly distributed in the addressing space of disk LUNs. The services that mainly generate random I/O include: OLTP service, SQL and instant messaging service.

Striping

Striping is a technique that automatically balances the load of I/O across multiple physical disks.

Throughput

It refers to the quantity of data successfully transmitted per unit time for the network, device, port, virtual circuit or other facilities.

Cloud Disk Snapshot

It saves a replica of the cloud disk at a certain point, and can be used to restore the cloud disk to the state when the snapshot was created.

Incremental Snapshot

Incremental snapshot refers to the snapshot where only the files added or modified compared with the previous one need to be backed up for each subsequent snapshot after the full snapshot or the previous incremental snapshot.

Main Boot Record

The Main Boot Record (MBR), also known as the main boot sector, is the first sector that must be read when accessing the hard disk after the computer is started. The main boot sector records the relevant information of the hard disk itself and the size and location information of each partition of the hard disk, and is also an important entrance for data information.

API Documentation
CBS(cbs)
version(2017-03-12)
API overview

API version

V3

Cloud Disk APIs

Interface name	Interface function
AttachDisks	AttachDisks
CreateDisks	CreateDisks
DescribeBlockStorages	DescribeBlockStorages
DescribeDiskAssociatedSnapshots	DescribeDiskAssociatedSnapshots
DescribeDiskConfigQuota	DescribeDiskConfigQuota
DescribeDiskSupportFeatures	DescribeDiskSupportFeatures
DescribeDisks	DescribeDisks
DescribeDisksDeniedActions	DescribeDisksDeniedActions
DescribeInstancesDiskNum	DescribeInstancesDiskNum
DescribeUserDiskResources	DescribeUserDiskResources
DescribeZoneDiskConfigInfos	DescribeZoneDiskConfigInfos
DetachDisks	DetachDisks
InitializeDisks	InitializeDisks
InquirePriceModifyDiskExtraPerformance	InquirePriceModifyDiskExtraPerformance

Interface name	Interface function
InquiryPriceCreateDisks	InquiryPriceCreateDisks
InquiryPriceModifyDiskAttributes	InquiryPriceModifyDiskAttributes
InquiryPriceResizeDisk	InquiryPriceResizeDisk
ModifyDiskAttributes	ModifyDiskAttributes
ModifyDiskExtraPerformance	ModifyDiskExtraPerformance
ResizeDisk	ResizeDisk
SwitchParameterCreateDisks	SwitchParameterCreateDisks
SwitchParameterResizeDisk	SwitchParameterResizeDisk
TerminateDisks	TerminateDisks

Other APIs

Interface name	Interface function
DescribeDedicatedClusterDiskStatistics	DescribeDedicatedClusterDiskStatistics
ModifyDisksSynchronizeFlag	ModifyDisksSynchronizeFlag

Snapshot APIs

Interface name	Interface function
ApplySnapshot	ApplySnapshot
ApplySnapshotGroup	ApplySnapshotGroup
BindAutoSnapshotPolicy	BindAutoSnapshotPolicy
CopyAutoSnapshotPolicyCrossAccount	CopyAutoSnapshotPolicyCrossAccount
CopySnapshotCrossAccount	CopySnapshotCrossAccount
CopySnapshotCrossRegions	CopySnapshotCrossRegions
CreateAutoSnapshotPolicy	CreateAutoSnapshotPolicy

Interface name	Interface function
CreateSnapshot	CreateSnapshot
CreateSnapshotGroup	CreateSnapshotGroup
DeleteAutoSnapshotPolicies	DeleteAutoSnapshotPolicies
DeleteSnapshotGroup	DeleteSnapshotGroup
DeleteSnapshots	DeleteSnapshots
DescribeAutoSnapshotPolicies	DescribeAutoSnapshotPolicies
DescribeCopyAutoSnapshotPolicy	DescribeCopyAutoSnapshotPolicy
DescribeDiskAssociatedAutoSnapshotPolicy	DescribeDiskAssociatedAutoSnapshotPolicy
DescribeSnapshotGroups	DescribeSnapshotGroups
DescribeSnapshotGroupsDeniedActions	DescribeSnapshotGroupsDeniedActions
DescribeSnapshotSharePermission	DescribeSnapshotSharePermission
DescribeSnapshots	DescribeSnapshots
DescribeSnapshotsDeniedActions	DescribeSnapshotsDeniedActions
GetSnapOverview	GetSnapOverview
ModifyAutoSnapshotPolicyAttribute	ModifyAutoSnapshotPolicyAttribute
ModifySnapshotAttribute	ModifySnapshotAttribute
ModifySnapshotsSharePermission	ModifySnapshotsSharePermission
UnbindAutoSnapshotPolicy	UnbindAutoSnapshotPolicy

CallingMethod signaturev1

Converge Cloud authenticates each access request, i.e. each request needs to include authentication information (Signature) in the common parameters to verify the identity of the requester.
The Signature is generated by the security credentials which include SecretId and SecretKey. If you don't have the security credentials yet, go to the API Key page to apply for them; otherwise, you cannot invoke the API.

1. Applying for Security Credentials

Before using the API for the first time, go to the API Key page to apply for security credentials.
Security credentials consist of SecretId and SecretKey:

- SecretId is used to identify the API requester.
- SecretKey is used to encrypt the signature string and verify it on the server.
- You must keep your security credentials private and avoid disclosure.

You can apply for the security credentials through the following steps:

1. Log in to the Converge Cloud Console.
2. Go to the Cloud API Key page.
3. On the API Key Management page, click Create Key to create a SecretId/SecretKey pair.

Note: Each account can have up to two pairs of SecretId/SecretKey.

2. Generating a Signature

With the SecretId and SecretKey, a signature can be generated. The following describes how to generate a signature:

Assume that the SecretId and SecretKey are:

- SecretId: AKIDz8krbsJ5yKBZQpn74WFkmLPx3EXAMPLE
- SecretKey: Gu5t9xGARNpq86cd98joQYCN3EXAMPLE

Note: This is just an example. For actual operations, please use your own SecretId and SecretKey.

Take the Cloud Virtual Machine's request to view the instance list (DescribeInstances) as an example. When

you invoke this API, the request parameters may be as follows:

Parameter name	Description	Parameter value
Action	Method name	DescribeInstances
SecretId	Key ID	AKIDz8krbsJ5yKBZQpn74WFkmLPx3EXAMPLE
Timestamp	Current timestamp	1465185768
Nonce	Random positive integer	11886
Region	Region where the instance is located	luzon
InstanceIds.0	ID of the instance to query	ins-09dx96dg
Offset	Offset	0
Limit	Allowed maximum output	20
Version	API version number	2017-03-12

2.1. Sorting Parameters

First, sort all the request parameters in an ascending lexicographical order (ASCII code) by their names. Notes: (1) Parameters are sorted by their names instead of their values; (2) The parameters are sorted based on ASCII code, not in an alphabetical order or by values. For example, InstanceIds.2 should be arranged after InstanceIds.12. You can complete the sorting process using a sorting function in a programming language, such as the ksort function in PHP. The parameters in the example are sorted as follows:

```
{
    'Action' : 'DescribeInstances',
    'InstanceIds.0' : 'ins-09dx96dg',
    'Limit' : 20,
    'Nonce' : 11886,
    'Offset' : 0,
    'Region' : 'luzon',
    'SecretId' : 'AKIDz8krbsJ5yKBZQpn74WFkmLPx3EXAMPLE',
    'Timestamp' : 1465185768,
    'Version': '2017-03-12',
}
```

When developing in another programming language, you can sort these sample parameters and it will work as long as you obtain the same results.

2.2. Concatenating a Request String

This step generates a request string.
Format the request parameters sorted in the previous step into the form of "parameter name"="parameter value". For example, for the Action parameter, its parameter name is "Action" and its parameter value is "DescribeInstances", so it will become Action=DescribeInstances after formatted.
Note: The "parameter value" is the original value but not the value after URL encoding.

Then, concatenate the formatted parameters with "&". The resulting request string is as follows:

```
Action=DescribeInstances&InstanceIds.0=ins-09dx96dg&Limit=20&Nonce=11886&Offset=0&Region=luzon&SecretId=AKIDz8krbsJ5yKBZQpn74WFkmLPx3EXAMPLE&Timestamp=1465185768&Version=2017-03-12
```

2.3. Concatenating the Signature Original String

This step generates a signature original string.
The signature original string consists of the following parameters:

1. HTTP method: POST and GET modes are supported, and GET is used here for the request. Please note that the method name should be in all capital letters.
2. Request host: the domain name of the request to view the list of instances (DescribeInstances) is cvm.convergecloud.com. The actual request domain name varies by the module to which the API belongs. For more information, see the instructions of the specific API.
3. Request path: The request path in the current version of Cloud API is fixed to / .
4. Request string: the request string generated in the previous step.

The concatenation rule of the signature original string is: Request method + request host + request path + ? + request string

The concatenation result of the example is:

```
GETcvm.convergecloud.com/?Action=DescribeInstances&InstanceIds.0=ins-09dx96dg&Limit=20&Nonce=11886&Offset=0&Region=luzon&SecretId=AKIDz8krbsJ5yKBZQpn74WFkmLPx3EXAMPLE&Timestamp=1465185768&Version=2017-03-12
```

2.4. Generating a Signature String

This step generates a signature string.
First, use the HMAC-SHA1 algorithm to sign thesignature original stringobtained in the previous step, and then encode the generated signature using Base64 to obtain the final signature.

The specific code is as follows with the PHP language being used as an example:

```
$secretKey = 'Gu5t9xGARNpq86cd98joQYCN3EXAMPLE';
$srcStr = 'GETcvm.convergecloud.com/?Action=DescribeInstances&InstanceIds.0=ins-09dx96dg&Limit
=20&Nonce=11886&Offset=0&Region=luzon&SecretId=AKIDz8krbsJ5yKBZQpn74WFkmLPx3EXAMPL
E&Timestamp=1465185768&Version=2017-03-12';
$signStr = base64_encode(hash_hmac('sha1', $srcStr, $secretKey, true));
echo $signStr;
```

The final signature is:

```
EliP9YW3pW28FpsEdkXt/+WcGeI=
```

When developing in another programming language, you can sign and verify the original in the example above and it works as long as you get the same results.

3. Encoding a Signature String

The generated signature string cannot be directly used as a request parameter and must be URL encoded.

For example, if the signature string generated in the previous step is EliP9YW3pW28FpsEdkXt/+WcGeI= , the final signature string request parameter (Signature) is EliP9YW3pW28FpsEdkXt%2f%2bWcGeI%3d, which will be used to generate the final request URL.

**Note: If your request method is GET, or the request method is POST and the Content-Type is application/x-www-form-urlencoded, then all the request parameter values need to be URL encoded (except the parameter key and the symbol of =) when sending the request. Non-ASCII characters need to be encoded with UTF-8 before URL encoding. **

Note: The network libraries of some programming languages automatically URL encode all parameters, in which case there is no need to URL encode the signature string; otherwise, two rounds of URL encoding will cause the signature to fail.

Note: Other parameter values also need to be encoded using RFC 3986. Use %XY in percent-encoding for special characters such as Chinese characters, where "X" and "Y" are hexadecimal characters (0-9 and uppercase A-F), and using lowercase will cause an error.

4. Signature Failure

The following situational error codes for signature failure may occur. Please resolve the errors accordingly.

Error code	Error description
AuthFailure.SignatureExpire	The signature is expired

Error code	Error description
AuthFailure.SecretIdNotFound	The key does not exist
AuthFailure.SignatureFailure	Signature error
AuthFailure.TokenFailure	Token error
AuthFailure.InvalidSecretId	Invalid key (not a Cloud API key type)

5. Signature Demo

When calling API 3.0, you are recommended to use the corresponding Converge Cloud SDK 3.0, Swhich encapsulates the signature process, enabling you to focus on only the specific APIs provided by the product when developing. See SDK Center for more information. Currently, the following programming languages are supported:

- Python
- Java
- PHP
- Go
- Node

To further explain the signing process, we will use a programming language to implement the process described above. The request domain name, API and parameter values in the sample are used here. This goal of this example is only to provide additional clarification for the signature process, please see the SDK for actual usage.

The final output URL might be: https://cvm.convergecloud.com/?Action=DescribeInstances&InstanceIds.0=ins-09dx96dg&Limit=20&Nonce=11886&Offset=0&Region=luzon&SecretId=AKIDz8krbsJ5yKBZQpn74WFkmLPx3EXAMPLE&Signature=EliP9YW3pW28FpsEdkXt%2F%2BWcGeI%3D&Timestamp=1465185768&Version=2017-03-12

Note: The key in the example is fictitious, and the timestamp is not the current time of the system, so if this URL is opened in the browser or called using commands such as curl, an authentication error will be returned: Signature expired. In order to get a URL that can work properly, you need to replace the SecretId and SecretKey in the example with your real credentials and use the current time of the system as the Timestamp.

Note: In the example below, even if you use the same programming language, the order of the parameters in the URL may be different for each execution. However, the order does not matter, as long as all the parameters are included in the URL and the signature is calculated correctly.

Note: The following code is only applicable to API 3.0. It cannot be directly used in other signature processes. Even with an older API, signature calculation errors may occur due to the differences in details. Please refer to the corresponding documentation.

Java

```
import java.io.UnsupportedEncodingException;
import java.net.URLEncoder;
import java.util.Random;
import java.util.TreeMap;
import javax.crypto.Mac;
import javax.crypto.spec.SecretKeySpec;
import javax.xml.bind.DatatypeConverter;

public class CloudAPIDemo {
    private final static String CHARSET = "UTF-8";

    public static String sign(String s, String key, String method) throws Exception {
        Mac mac = Mac.getInstance(method);
        SecretKeySpec secretKeySpec = new SecretKeySpec(key.getBytes(CHARSET), mac.getAlgorithm());
        mac.init(secretKeySpec);
        byte[] hash = mac.doFinal(s.getBytes(CHARSET));
        return DatatypeConverter.printBase64Binary(hash);
    }

    public static String getStringToSign(TreeMap<String, Object> params) {
        StringBuilder s2s = new StringBuilder("GETcvm.convergecloud.com/?");
        // When signing, the parameters need to be sorted in lexicographical order. TreeMap is used here
to guarantee the correct order.
        for (String k : params.keySet()) {
            s2s.append(k).append("=").append(params.get(k).toString()).append("&");
        }
        return s2s.toString().substring(0, s2s.length() - 1);
    }

    public static String getUrl(TreeMap<String, Object> params) throws UnsupportedEncodingException
{
        StringBuilder url = new StringBuilder("https://cvm.convergecloud.com/?");
        // There is no requirement for the order of the parameters in the actual request URL.
        for (String k : params.keySet()) {
            // The request string needs to be URL encoded. As the Key is all in English letters, only the valu
e is URL encoded here.
            url.append(k).append("=").append(URLEncoder.encode(params.get(k).toString(), CHARSET)).app
end("&");
        }
        return url.toString().substring(0, url.length() - 1);
```

```
    }

    public static void main(String[] args) throws Exception {
        TreeMap<String, Object> params = new TreeMap<String, Object>(); // TreeMap enables automati
c sorting
        // A random number should be used when actually calling, for example: params.put("Nonce", new
Random().nextInt(java.lang.Integer.MAX_VALUE));
        params.put("Nonce", 11886); // Common parameter
        // The current time of the system should be used when actually calling, for example:   params.put
("Timestamp", System.currentTimeMillis() / 1000);
        params.put("Timestamp", 1465185768); // Common parameter
        params.put("SecretId", "AKIDz8krbsJ5yKBZQpn74WFkmLPx3EXAMPLE"); // Common parameter
        params.put("Action", "DescribeInstances"); // Common parameter
        params.put("Version", "2017-03-12"); // Common parameter
        params.put("Region", "luzon"); // Common parameter
        params.put("Limit", 20); // Business parameter
        params.put("Offset", 0); // Business parameter
        params.put("InstanceIds.0", "ins-09dx96dg"); // Business parameter
        params.put("Signature", sign(getStringToSign(params), "Gu5t9xGARNpq86cd98joQYCN3EXAMPLE
", "HmacSHA1")); // Common parameter
        System.out.println(getUrl(params));
    }
}
```

Python

Note: If running in a Python 2 environment, the following requests dependency package must be installed first: pip install requests .

```
# -*- coding: utf8 -*-
import base64
import hashlib
import hmac
import time

import requests

secret_id = "AKIDz8krbsJ5yKBZQpn74WFkmLPx3EXAMPLE"
secret_key = "Gu5t9xGARNpq86cd98joQYCN3EXAMPLE"

def get_string_to_sign(method, endpoint, params):
    s = method + endpoint + "/?"
    query_str = "&".join("%s=%s" % (k, params[k]) for k in sorted(params))
    return s + query_str

def sign_str(key, s, method):
```

```
    hmac_str = hmac.new(key.encode("utf8"), s.encode("utf8"), method).digest()
    return base64.b64encode(hmac_str)

if __name__ == '__main__':
    endpoint = "cvm.convergecloud.com"
    data = {
        'Action' : 'DescribeInstances',
        'InstanceIds.0' : 'ins-09dx96dg',
        'Limit' : 20,
        'Nonce' : 11886,
        'Offset' : 0,
        'Region' : 'luzon',
        'SecretId' : secret_id,
        'Timestamp' : 1465185768, # int(time.time())
        'Version': '2017-03-12'
    }
    s = get_string_to_sign("GET", endpoint, data)
    data["Signature"] = sign_str(secret_key, s, hashlib.sha1)
    print(data["Signature"])
    # n actual invocation would occur here, which may incur fees after success
    # resp = requests.get("https://" + endpoint, params=data)
    # print(resp.url)
```

signaturev3

Converge Cloud The API will authenticate each access request, meaning that each request must include signature information (Signature) in the public request parameters to verify the requester's identity.
The signature information is generated by the security credentials, which include SecretId and SecretKey . If the user does not have security credentials, please apply for them on the Cloud API Key page, otherwise they cannot call the Cloud API interfaces.

1. Apply for security credentials.

Before using the Cloud API for the first time, please go to the Cloud API Key page to apply for a security credential.
Security credentials include SecretId and SecretKey:

- SecretId The security credentials are used to identify the identity of the API caller.
- SecretKey The secret key is used to encrypt and sign the string for server-side verification of the signature.
- Users must strictly protect their security credentials and avoid disclosure them.

The specific steps to apply for security credentials are as follows:

1. Log in to Converge Cloud Administration Center Console
2. Go to the console page for cloud API keys
3. On the Cloud API Key page, click New to create a pair of SecretId/SecretKey

Note: Developer accounts can have up to two pairs SecretId / SecretKey.

2. TC3-HMAC-SHA256 signature method

Note: For GET method, only Content-Type: application/x-www-form-urlencoded protocol format is supported. For POST method, Content-Type: application/json and Content-Type: multipart/form-data are currently supported. By default, all business interfaces are supported in json format. Only specific business interfaces are supported in multipart format. In this case, the interface cannot be called in json format. Refer to the specific business interface document description.

The following takes Cloud Virtual Machine querying Region-1 instance list as an example to introduce the signature calculation process step by step. We only use two parameters for the query instance list: Limit and Offset, called using the GET method.

Assume that the SecretId and SecretKey of the user are AKIDz 8krbsJ 5yKBZQpn 74WFkmLPx 3EXAMPLE and Gu 5t9xGARNpq 86cd 98joQYCN3EXAMPLE, respectively.

2.1. splice specification request string

Concatenate CanonicalRequest strings in the following format:

```
CanonicalRequest =
  HTTPRequestMethod + '\n' +
  CanonicalURI + '\n' +
  CanonicalQueryString + '\n' +
  CanonicalHeaders + '\n' +
  SignedHeaders + '\n' +
  HashedRequestPayload
```

- HTTPRequestMethod: HTTP request methods (GET, POST), GET in this example;
- CanonicalURI: URI parameter, API 3.0 fixed to forward slash (/);
- CanonicalQueryString: Query string in the URL that initiates the HTTP request. For POST requests, it is fixed as an empty string. For GET requests, it is the question mark (?) in the URL. The following string content, this example value is: Limit=10&Offset=0. Note: CanonicalQueryString needs to be URL encoded.
- CanonicalHeaders: The header information of participating signature, at least including host and content-type, can also be added with customized header to improve the uniqueness and security of its own request. Splicing rules: 1) head key and value are uniformly converted to lowercase, and the first and last spaces are removed, and spliced according to the format of key:value\n;2) multiple heads are spliced according to the dictionary sorting of head key (lowercase). In this example: content-type: application/x-www-form-urlencoded\nhost: cvm.convergecloud.com\n
- SignedHeaders: The header information involved in the signature indicates which headers participate in the signature of this request, and the header content contained in CanonicalHeaders is one-to-one corresponding. content-type and host are required headers. constructed as: 1) The head key is uniformly converted to lowercase; 2) Multiple header keys (lowercase) are concatenated according to dictionary sorting and separated by semicolons (;). In this case: content-type;host
- HashedRequestPayload: The hash value of the request body, calculated as Lowercase(HexEncode(Hash.SHA256(RequestPayload))), SHA256 hash the entire payload of the HTTP request body, then hexadecimal encoding, and finally converting the encoding string to lowercase letters. Note: For GET requests, RequestPayload is fixed as an empty string. For POST requests, RequestPayload is the payload of the HTTP request body.

According to the above rules, the resulting canonical request string in the example is as follows (for clarity, \nnewlines are replaced by a new line printed separately):

```
GET
```

```
/
Limit=10&Offset=0
content-type:application/x-www-form-urlencoded
host:cvm.convergecloud.com

content-type;host
e3b0c44298fc1c149afbf4c8996fb92427ae41e4649b934ca495991b7852b855
```

2.2. Concatenate the string to be signed

Concatenate the string to be signed in the following format:

```
StringToSign =
  Algorithm + \n +
  RequestTimestamp + \n +
  CredentialScope + \n +
  HashedCanonicalRequest
```

- Algorithm: Signature algorithm, currently fixed as TC3-HMAC-SHA256;
- RequestTimestamp: Request timestamp, i.e. the X-TC-Timestamp value of the request header. The request in the example above is 1539084154;
- CredentialScope: Credential range in the format Date/service/tc3_request, containing date, requested service and termination string (tc3_request). **Date is the date of UTC standard time, the value needs to be consistent with UTC standard time converted by common parameter X-TC-Timestamp **;service is the product name, which must be consistent with the product domain name called, such as cvm. The above example request, the value is 2018-10-09/cvm/tc3_request;
- HashedCanonicalRequest: The hash value of the canonical request string obtained by concatenating the above steps is calculated by Lowercase(HexEncode(Hash.SHA256(CanonicalRequest))).

note:

> 1. Date must be calculated from timestamp X-TC-Timestamp and time zone UTC+0. Adding local time zone information to the system, such as East Zone 8, will cause calls to succeed during the day and night, but calls to fail in the early hours of the morning. Assuming the timestamp is 1551113065, the time in the East Eight District is 2019-02-26 00:44:25, but the calculated Date with UTC+0 should be 2019-02-25 instead of 2019-02-26.
> 2. Timestamp must be the current system time and ensure that the system time and standard time are synchronized. If the difference exceeds five minutes, it will fail. If it is not synchronized with the standard time for a long time, it may cause the request to fail after running for a period of time (return signature expiration error).

According to the above rules, the string to be signed in the example is as follows (for clarity,\nnewlines are

replaced by new lines printed):

```
TC3-HMAC-SHA256
1539084154
2018-10-09/cvm/tc3_request
91c9c192c14460df6c1ffc69e34e6c5e90708de2a6d282cccf957dbf1aa7f3a7
```

2.3. calculate the signature

1. Compute the derived signature key, pseudocode as follows

```
SecretKey = "Gu5t9xGARNpq86cd98joQYCN3EXAMPLE"
SecretDate = HMAC_SHA256("TC3" + SecretKey, Date)
SecretService = HMAC_SHA256(SecretDate, Service)
SecretSigning = HMAC_SHA256(SecretService, "tc3_request")
```

- SecretKey: Original SecretKey;
- Date: That is, the Date field information in Credential, as shown in the example above, is 2018-10-09.;
- Service: That is, the Service field information in Credential, as shown in the example above, is cvm;

2. Calculate the signature, pseudocode as follows

```
Signature = HexEncode(HMAC_SHA256(SecretSigning, StringToSign))
```

- SecretSigning: That is, the derivative signature key calculated above;
- StringToSign: That is, the character string to be signed calculated in step 2;

2.4. Splicing Authorization

Splice as follows Authorization:

```
Authorization =
    Algorithm + ' ' +
    'Credential=' + SecretId + '/' + CredentialScope + ', ' +
    'SignedHeaders=' + SignedHeaders + ', '
    'Signature=' + Signature
```

- Algorithm: Signature method, fixed as TC3-HMAC-SHA256;
- SecretId: SecretId in key pair;
- CredentialScope: See above, Scope of Voucher;
- SignedHeaders: See above, header information for participating signatures;
- Signature: signature value

According to the above rules, the values obtained in the example are:

```
TC3-HMAC-SHA256 Credential=AKIDEXAMPLE/Date/service/tc3_request, SignedHeaders=content-type;host, Signature=5da7a33f6993f0614b047e5df4582db9e9bf4672ba50567dba16c6ccf174c474
```

The final complete invocation information is as follows:

```
https://cvm.convergecloud.com/?Limit=10&Offset=0

Authorization: TC3-HMAC-SHA256 Credential=AKIDz8krbsJ5yKBZQpn74WFkmLPx3EXAMPLE/2018-10-09/cvm/tc3_request, SignedHeaders=content-type;host, Signature=5da7a33f6993f0614b047e5df4582db9e9bf4672ba50567dba16c6ccf174c474
Content-Type: application/x-www-form-urlencoded
Host: cvm.convergecloud.com
X-TC-Action: DescribeInstances
X-TC-Version: 2017-03-12
X-TC-Timestamp: 1539084154
X-TC-Region: luzon
```

3. signing failed

According to the actual situation, there are the following signature failure error codes, please handle according to the actual situation

error code	error description
AuthFailure.SignatureExpire	signature expiration
AuthFailure.SecretIdNotFound	key is not present
AuthFailure.SignatureFailure	signature error
AuthFailure.TokenFailure	token error
AuthFailure.InvalidSecretId	Illegal key (not cloud API key type)

4. Signature Demo

Java

```
import java.io.BufferedReader;
import java.io.InputStream;
```

```
import java.io.InputStreamReader;
import java.net.URL;
import java.text.SimpleDateFormat;
import java.util.Date;
import java.util.Map;
import java.util.TimeZone;
import java.util.TreeMap;
import javax.crypto.Mac;
import javax.crypto.spec.SecretKeySpec;
import javax.net.ssl.HttpsURLConnection;
import javax.xml.bind.DatatypeConverter;

import org.apache.commons.codec.digest.DigestUtils;

public class CloudAPITC3Demo {
    private final static String CHARSET = "UTF-8";
    private final static String ENDPOINT = "cvm.convergecloud.com";
    private final static String PATH = "/";
    private final static String SECRET_ID = "AKIDz8krbsJ5yKBZQpn74WFkmLPx3EXAMPLE";
    private final static String SECRET_KEY = "Gu5t9xGARNpq86cd98joQYCN3EXAMPLE";
    private final static String CT_X_WWW_FORM_URLENCODED = "application/x-www-form-urlencode
d";
    private final static String CT_JSON = "application/json";
    private final static String CT_FORM_DATA = "multipart/form-data";

    public static byte[] sign256(byte[] key, String msg) throws Exception {
        Mac mac = Mac.getInstance("HmacSHA256");
        SecretKeySpec secretKeySpec = new SecretKeySpec(key, mac.getAlgorithm());
        mac.init(secretKeySpec);
        return mac.doFinal(msg.getBytes(CHARSET));
    }

    public static void main(String[] args) throws Exception {
        String service = "cvm";
        String host = "cvm.convergecloud.com";
        String region = "luzon";
        String action = "DescribeInstances";
        String version = "2017-03-12";
        String algorithm = "TC3-HMAC-SHA256";
        String timestamp = "1539084154";
        //String timestamp = String.valueOf(System.currentTimeMillis() / 1000);
        SimpleDateFormat sdf = new SimpleDateFormat("yyyy-MM-dd");
        // Pay attention to the time zone, otherwise it is easy to make mistakes.
        sdf.setTimeZone(TimeZone.getTimeZone("UTC"));
        String date = sdf.format(new Date(Long.valueOf(timestamp + "000")));

        // ************* Step 1: Assemble the specification request string *************
```

```
        String httpRequestMethod = "GET";
        String canonicalUri = "/";
        String canonicalQueryString = "Limit=10&Offset=0";
        String canonicalHeaders = "content-type:application/x-www-form-urlencoded\n" + "host:" + host
+ "\n";
        String signedHeaders = "content-type;host";
        String hashedRequestPayload = DigestUtils.sha256Hex("");
        String canonicalRequest = httpRequestMethod + "\n" + canonicalUri + "\n" + canonicalQueryStri
ng + "\n"
                + canonicalHeaders + "\n" + signedHeaders + "\n" + hashedRequestPayload;
        System.out.println(canonicalRequest);

        // ************* Step 2: Stitching the string to be signed *************
        String credentialScope = date + "/" + service + "/" + "tc3_request";
        String hashedCanonicalRequest = DigestUtils.sha256Hex(canonicalRequest.getBytes(CHARSET));
        String stringToSign = algorithm + "\n" + timestamp + "\n" + credentialScope + "\n" + hashedCan
onicalRequest;
        System.out.println(stringToSign);

        // ************* Step 3: Calculate the signature *************
        byte[] secretDate = sign256(("TC3" + SECRET_KEY).getBytes(CHARSET), date);
        byte[] secretService = sign256(secretDate, service);
        byte[] secretSigning = sign256(secretService, "tc3_request");
        String signature = DatatypeConverter.printHexBinary(sign256(secretSigning, stringToSign)).toLow
erCase();
        System.out.println(signature);

        // ************* Step 4: Splicing Authorization *************
        String authorization = algorithm + " " + "Credential=" + SECRET_ID + "/" + credentialScope + ", "
                + "SignedHeaders=" + signedHeaders + ", " + "Signature=" + signature;
        System.out.println(authorization);

        TreeMap<String, String> headers = new TreeMap<String, String>();
        headers.put("Authorization", authorization);
        headers.put("Host", host);
        headers.put("Content-Type", CT_X_WWW_FORM_URLENCODED);
        headers.put("X-TC-Action", action);
        headers.put("X-TC-Timestamp", timestamp);
        headers.put("X-TC-Version", version);
        headers.put("X-TC-Region", region);
    }
}
```

Python

```
# -*- coding: utf-8 -*-
```

```
import hashlib, hmac, json, os, sys, time
from datetime import datetime

# key parameters
secret_id = "AKIDz8krbsJ5yKBZQpn74WFkmLPx3EXAMPLE"
secret_key = "Gu5t9xGARNpq86cd98joQYCN3EXAMPLE"

service = "cvm"
host = "cvm.convergecloud.com"
endpoint = "https://" + host
region = "luzon"
action = "DescribeInstances"
version = "2017-03-12"
algorithm = "TC3-HMAC-SHA256"
timestamp = 1539084154
date = datetime.utcfromtimestamp(timestamp).strftime("%Y-%m-%d")
params = {"Limit": 10, "Offset": 0}

# ************* Step 1: Assemble the specification request string *************
http_request_method = "GET"
canonical_uri = "/"
canonical_querystring = "Limit=10&Offset=0"
ct = "x-www-form-urlencoded"
payload = ""
if http_request_method == "POST":
    canonical_querystring = ""
    ct = "json"
    payload = json.dumps(params)
canonical_headers = "content-type:application/%s\nhost:%s\n" % (ct, host)
signed_headers = "content-type;host"
hashed_request_payload = hashlib.sha256(payload.encode("utf-8")).hexdigest()
canonical_request = (http_request_method + "\n" +
                     canonical_uri + "\n" +
                     canonical_querystring + "\n" +
                     canonical_headers + "\n" +
                     signed_headers + "\n" +
                     hashed_request_payload)
print(canonical_request)

# ************* Step 2: Stitching the string to be signed *************
credential_scope = date + "/" + service + "/" + "tc3_request"
hashed_canonical_request = hashlib.sha256(canonical_request.encode("utf-8")).hexdigest()
string_to_sign = (algorithm + "\n" +
                  str(timestamp) + "\n" +
                  credential_scope + "\n" +
                  hashed_canonical_request)
print(string_to_sign)
```

```
# ************* Step 3: Calculate the signature *************
# Compute signature digest function
def sign(key, msg):
    return hmac.new(key, msg.encode("utf-8"), hashlib.sha256).digest()
secret_date = sign(("TC3" + secret_key).encode("utf-8"), date)
secret_service = sign(secret_date, service)
secret_signing = sign(secret_service, "tc3_request")
signature = hmac.new(secret_signing, string_to_sign.encode("utf-8"), hashlib.sha256).hexdigest()
print(signature)

# ************* Step 4: Splicing Authorization *************
authorization = (algorithm + " " +
                 "Credential=" + secret_id + "/" + credential_scope + ", " +
                 "SignedHeaders=" + signed_headers + ", " +
                 "Signature=" + signature)
print(authorization)

# Common parameters added to request header
headers = {
    "Authorization": authorization,
    "Host": host,
    "Content-Type": "application/%s" % ct,
    "X-TC-Action": action,
    "X-TC-Timestamp": str(timestamp),
    "X-TC-Version": version,
    "X-TC-Region": region,
}
```

request

1. Service Address

Region refers to the geographical area of a physical data center. Converge Cloud ensures complete isolation between different regions, providing maximum stability and fault tolerance. To reduce access latency and improve download speed, it is recommended to choose the region closest to your customers.

You can view the complete region list by using the API interface "DescribeRegions".

2. Communications Protocol

All the Converge Cloud APIs communicate via HTTPS, providing highly secure communication tunnels.

3. Request Methods

Supported HTTP request methods:

- POST (recommended)
- GET

PThe Content-Type types supported by POST requests:

- application/json (recommended). The TC3-HMAC-SHA256 signature algorithm must be used.
- application/x-www-form-urlencoded. The HmacSHA1 or HmacSHA256 signature algorithm must be used.
- multipart/form-data (only supported by certain APIs). You must use TC3-HMAC-SHA256 to calculate the signature.

The size of a GET request packet is up to 32 KB. The size of a POST request is up to 1 MB when the HmacSHA1 or HmacSHA256 signature algorithm is used, and up to 10 MB when TC3-HMAC-SHA256 is used.

4. Character Encoding

Only UTF-8 encoding is used.

result

Response for Successful Requests

For example, when calling CAM API (version: 2017-03-12) to view the status of instances (DescribeInstancesStatus), if the request has succeeded, you may see the response as shown below:


```
{
    "Response": {
        "TotalCount": 0,
        "InstanceStatusSet": [],
        "RequestId": "b5b41468-520d-4192-b42f-595cc34b6c1c"
    }
}
```


- The API will return Response, which contains RequestId, as long as it processes the request. It does not matter if the request is successful or not.
- RequestId is the unique ID of an API request. Contact us with this ID when an exception occurs.
- Except for the fixed fields, all fields are action-specified. For the definitions of action-specified fields, see the corresponding API documentation. In this example, TotalCount and InstanceStatusSet are the fields specified by the API DescribeInstancesStatus. 0 TotalCount means that the requester owns 0 CVM instance so the InstanceStatusSet is empty.

Response for Failed Requests

If the request has failed, you may see the response as shown below:


```
{
    "Response": {
        "Error": {
            "Code": "AuthFailure.SignatureFailure",
            "Message": "The provided credentials could not be validated. Please check your signature is cor
rect."
        },
        "RequestId": "ed93f3cb-f35e-473f-b9f3-0d451b8b79c6"
    }
}
```


- The presence of the Error field indicates that the request has failed. A response for a failed request will include Error, Code and Message fields.

- Code is the code of the error that helps you identify the cause and solution. There are two types of error codes so you may find the code in either common error codes or API-specified error codes.
- Message explains the cause of the error. Note that the returned messages are subject to service updates. The information the messages provide may not be up-to-date and should not be the only source of reference.
- RequestId is the unique ID of an API request. Contact us with this ID when an exception occurs.

Common Error Codes

If there is an Error field in the response, it means that the API call failed. The Code field in Error indicates the error code. The following table lists the common error codes that all actions can return.

Error Code	Description
AuthFailure.InvalidSecretId	Invalid key (not a API key type).
AuthFailure.MFAFailure	MFA failed.
AuthFailure.SecretIdNotFound	The key does not exist.
AuthFailure.SignatureExpire	Signature expired.
AuthFailure.SignatureFailure	Signature error.
AuthFailure.TokenFailure	Token error.
AuthFailure.UnauthorizedOperation	The request does not have CAM authorization.
DryRunOperation	DryRun Operation. It means that the request would have succeeded, but the DryRun parameter was used.
FailedOperation	Operation failed.
InternalError	Internal error.
InvalidAction	The API does not exist.
InvalidParameter	Incorrect parameter.
InvalidParameterValue	Invalid parameter value.
LimitExceeded	Quota limit exceeded.
MissingParameter	A parameter is missing.
NoSuchVersion	The API version does not exist.

Error Code	Description
RequestLimitExceeded	The number of requests exceeds the frequency limit.
ResourceInUse	Resource is in use.
ResourceInsufficient	Insufficient resource.
ResourceNotFound	The resource does not exist.
ResourceUnavailable	Resource is unavailable.
UnauthorizedOperation	Unauthorized operation.
UnknownParameter	Unknown parameter.
UnsupportedOperation	Unsupported operation.
UnsupportedProtocol	hHTTPS request method error. Only GET and POST requests are supported.
UnsupportedRegion	API does not support the requested region.

PublicParameter

Common parameters are parameters used to identify users and interface authentication purposes. If not necessary, these parameters are not described in the separate interface document of each interface, but these parameters need to be carried in each request before the request can be normally initiated.

signature method v3

When using TC3-HMAC-SHA256 signature method, the common parameters need to be uniformly placed in the HTTP Header request header, as follows:

parameter name	type	required	described
X-TC-Action	String	Yes	The name of the operation interface. Refer to the description of the public parameter Action in the input parameters of the interface documentation for the value. For example, the value for the query instance list interface of the cloud server is DescribeInstances.
X-TC-Region	String	Yes	region parameters, Data used to identify which region you wish to operate on. The region value accepted by the interface refers to the description of the input parameter common parameter Region in the interface document. Note: Some interfaces do not need to pass this parameter, which is specifically stated in the interface documentation. In this case, even passing this parameter will not take effect.
X-TC-Timestamp	Integer	Yes	The current UNIX timestamp, which records the time when the API request was initiated. for example 1529223702. Note: If the time difference from the server exceeds 5 minutes, it will cause a signature expiration error.
X-TC-Version	String	Yes	Version of API for operation. Refer to the description of the public parameter Version in the interface document for the value. For example, the version of the cloud server. 2017-03-12.
Authorization	String	Yes	HTTP The standard authentication header field is "Authorization". For example: TC3-HMAC-SHA256 Credential=AKIDEXAMPLE/Date/service/tc3_request, SignedHeaders=content-type;host, Signature=fe5f80f77d5fa3beca038a248ff027d0445342fe2855ddc963176630326f1024 Here, - TC3-HMAC-SHA256: Signature method, currently fixed as this value, - Credential: Credential: Signature credential; AKIDEXAMPLE is the SecretId; Date is a date in UTC time, and this value must match the value of X-TC-Timestamp (a common parameter) in UTC time format; service is the name of the product/service, and is generally a domain name prefix. For example,cvm; - SignedHeaders: The headers that contains the authentication information; content-type and host are the required headers; - Signature: Signature digest.
X-TC-Token	String	No	The token used for temporary certificates needs to be used in conjunction with a temporary key. Temporary keys and tokens need to be obtained by calling the access management service interface. Long-term keys do not require tokens.

signature method v1

When using HmacSHA1 and HmacSHA256 signature methods, common parameters need to be placed uniformly in the request string, as follows

parameter name	type	required	described
Action	String	Yes	The name of the interface to be operated. The value refers to the explanation of the public parameter Action in the interface documentation. For example, for the cloud server's query instance list interface, the value is DescribeInstances .
Region	String	Yes	The region parameter is used to identify which region's data you want to operate on. The region value accepted by the interface is specified in the input parameter of the public parameter Region in the interface documentation. Note: Some interfaces do not require this parameter to be passed, which will be specifically mentioned in the interface documentation. Even if this parameter is passed, it will not take effect in this case.
Timestamp	Integer	Yes	The current UNIX timestamp, which can record the time when the API request is initiated. For example, 1529223702. If the difference between the current time is too large, it will cause a signature expiration error.
Nonce	Integer	Yes	A random positive integer combined with Timestamp is used to prevent replay attacks.
SecretId	String	Yes	The SecretId is an identifier applied for the cloud API key to identify the user's identity. A SecretId corresponds to a unique SecretKey , and the SecretKey is used to generate the request signature Signature .
Signature	String	Yes	Request signature is used to verify the legitimacy of this request. Users need to calculate it based on the actual input parameters. The specific calculation method can be found in the interface authentication documentation.
Version	String	Yes	The version of the API being operated on. The value refers to the description of the public parameter Version in the interface documentation. For example, the version of the cloud server. 2017-03-12.
SignatureMethod	String	No	Signature method, currently supports HmacSHA256 and HmacSHA1. HmacSHA256 algorithm is used to verify signatures only if this parameter is specified as HmacSHA 256, otherwise HmacSHA1 is used to verify signatures.
Token	String	No	Token used for temporary certificate needs to be used together with temporary key. The temporary key and Token need to be obtained from the Access Management service invocation interface. Long-term keys do not require tokens.

zone list

A region (Region) refers to the geographical area of a physical data center. Converge Cloud Translation: To deliver verification that completely isolates different regions, ensuring maximum stability and fault tolerance between regions. In order to reduce access latency and improve download speed, it is recommended that you choose the region closest to your customers.

You can view the complete region list by using the API interface Query Region List.

Cloud Disk APIs
AttachDisks

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API is used to mount a cloud disk.

- Batch operations are supported for mounting multiple cloud disks to the same CVM. If any one in multiple cloud disks cannot be mounted, the operation will not be executed and a specific error code will be returned.
- This API is asynchronous. When a successful response is returned for a request to mount a cloud disk, it indicates that the backend has initiated an operation to mount the cloud disk. The status of the corresponding cloud disk can be queried through the DescribeDisks API. If the status of the cloud disk changes from ATTACHING to ATTACHED, it indicates that the cloud disk is mounted successfully.

Default interface request rate limit: 20 times per second.

Interface update time: 2025-08-28 01:23:24.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: AttachDisks
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
DiskIds	Yes	No	Array of String	ID of the elastic cloud disk to be mounted. It can be queried through the DescribeDisks API. Up to 10 elastic cloud disks can be mounted in a single request. Example: ["disk-7q2iy7gp"]
InstanceId	Yes	No	String	ID of the CVM instance to which the cloud disk will be mounted. It can be queried through the DescribeInstances API. Example: ins-9oxguxa
AlignType	No	No	String	Optional. If this parameter is not input, only the mount operation will be performed. When this parameter is input, the lifetime of the cloud disk will be aligned with that of the host to be mounted when mounting. Value range: • AUTO_RENEW: When auto-renew is not enabled for the cloud disk, this value can be input to enable auto-renew. • DEADLINE_ALIGN: When the cloud disk's expiry date is earlier than that of the host to be mounted, this value can be input to align the expiry date of the cloud disk with that of the host. Example: AUTO_RENEW
DeleteWithInstance	No	No	Bool	Optional parameter. If this parameter is not input, only the mounting operation is executed. If True is input, the cloud disk will be destroyed with the cloud host upon successful mounting. It is only valid for pay-as-you-go cloud disks. Example: TRUE
InstanceIds	No	No	Array of String	Optional. This parameter is used for the console to batch mount shared disks to multiple CVMs by specifying the instance IDs. If multiple InstanceIds are input, all cloud disks specified in DisksId must be shared cloud disks; otherwise an error is returned. Example: ["ins-locaa0iu"]

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
ZoneNotMatch	The cloud disk and the instance are not in the same availability zone.
InvalidDisk.NotPortable	Non-elastic cloud disks are not supported.
InvalidDisk.TypeError	Incorrect cloud disk type.
InvalidInstance.NotSupported	The Cloud Virtual Machine does not support mounting a cloud disk.
InvalidInstanceId.NotFound	The input InstanceId does not exist.
InvalidParameterValue.LimitExceeded	The parameter value is out of range
InvalidDisk.Attached	The cloud disk has been mounted.
InternalError.ResourceOpFailed	The operation on the resource failed. For specific error information, view the Message field in the error description. Try again later or contact customer service for help.
InvalidDiskId.NotFound	The input DiskId does not exist.
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
InvalidDisk.NotSupported	The cloud disk does not support this operation.
InternalError.ComponentError	Third-party Component Error

CreateDisks

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API is used to create a cloud disk.

- The purchase of a prepaid cloud disk will deduct the amount required for the current purchase in advance. Ensure sufficient account balance before this API is called.
- This API allows you to create a cloud disk by inputting a data disk snapshot and then copy snapshot data to the purchased cloud disk.
- This API is asynchronous. When a creation request is successfully issued, a new cloud disk ID list is returned, but creation of the cloud disk is not immediately completed. The status of the corresponding cloud disk can be queried by DiskId through the DescribeDisks API. If the cloud disk is found in the UNATTACHED or ATTACHED state, it indicates that the cloud disk is created successfully.

Default interface request rate limit: 20 times per second.

Interface update time: 2025-08-28 17:42:50.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: CreateDisks
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
DiskType	Yes	No	String	Hard disk media type. Value range: • CLOUD_BASIC: common cloud disk • CLOUD_PREMIUM: high-performance cloud disk • CLOUD_SSD: SSD cloud disk. Example: CLOUD_SSD
DiskName	No	No	String	Display name of the cloud disk. If not specified, it defaults to "unnamed". The maximum length cannot exceed 60 bytes. Example: HCC_System disk
DiskCount	No	No	Uint64	Number of cloud disks created. If this parameter is not input, the default value is 1. There is a limit on the maximum number of cloud disks that can be created in a single request. For details, see the cloud disk use limits. Example: 1
DiskChargeType	Yes	No	String	Payment mode. Currently, the only option is PREPAID. Example: PREPAID
DiskChargePrepaid	No	No	DiskChargePrepaid	Prepaid mode, that is, the settings for the monthly subscription-related parameters. Through this parameter, you can specify the purchase duration of the monthly subscribed cloud disk, whether to enable auto-renewal, and other attributes. It is a required parameter for creating a prepaid cloud disk. Example:View
Placement	Yes	No	Placement	Location of the instance. By specifying this parameter, you can determine the instance availability zone and project. If the project is not specified, the instance will be created under the default project. Example:View
DiskSize	No	No	Uint64	Cloud disk size, in GB. • If SnapshotId is input, DiskSize is optional. In this case, the size of the new cloud disk will be the size of the snapshot. • If SnapshotId and DiskSize are both input, then the cloud disk size must be greater than or equal to the snapshot size. • The range of cloud disk sizes: Common cloud disk: 10 GB to 4000 GB; high performance cloud disk: 50 GB to 4000 GB; SSD cloud disk: 100 GB to 4000 GB. The step is 10 GB. Example: 500
SnapshotId	No	No	String	Snapshot ID. If this parameter is input, a cloud disk is created based on this snapshot. The snapshot type must be a data disk snapshot. Snapshots can be queried through the DescribeSnapshots API. For details, see the explanation of the output parameter DiskUsage. Example: snap-4zgc8cnk
ClientToken	No	No	String	A string used to ensure request idempotency. This string is generated by the client and must be unique among requests, with a maximum of 64 ASCII characters. If this parameter is not specified, the request idempotency cannot be guaranteed. Example: 5034dbd5ec5446161df12497be78934d1e89100910001
Encrypt	No	No	String	Input this parameter to create an encrypted cloud disk. The value is fixed to ENCRYPT. Example: ENCRYPT
Tags	No	No	Array of Tag	Tag to which cloud disks are bound. Example:View
Shareable	No	No	Bool	Optional. The default value is False. When True is input, the cloud disk will be created as a shared cloud disk. Example: false
AutoSnapshotPolicyId	No	No	String	Regular snapshot policy ID. When this parameter is input, the cloud disk will be automatically bound to the regular snapshot policy upon successful creation. Example: asp-mrsm243
DiskStoragePoolGroup	No	No	String	Storage Resource Pool Group Example: pool-ec
PlatformProjectId	No	No	String	Project ID Example: pr-42749a13
ThroughputPerformance	No	No	Uint64	Optional. You can use this parameter to purchase additional performance for cloud disk. Currently, only ultra-fast cloud disk (CLOUD_TSSD) and enhanced SSD cloud disk (CLOUD_HSSD) are supported. Example: 100
DiskBackupQuota	No	No	Uint64	Number of cloud disk backup points. Example: 1
BackupDisk	No	No	Bool	Whether to back up the data by using the snapshot during cloud disk termination. Example: false

3. Output Parameters

Parameter name	Type	Description
DiskIdSet	Array of String	ID list of created cloud disks. Example: ["disk-lzrg2pwi"]
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameter.DiskConfigNotSupported	The current region does not support the cloud disk of the current configuration.
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.
TradeDealConflict	Order conflict.
InternalError.ComponentError	Third-party Component Error
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
ResourceInsufficient.OverQuota	Insufficient quota

DescribeBlockStorages

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API (DescribeBlockStorages) is used to query the list of block storage, including both local disks and cloud disks.

Default interface request rate limit: 20 times per second.

Interface update time: 2020-06-03 09:07:28.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DescribeBlockStorages
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
Limit	No	No	Uint64	Number of returned values. The default value is 100, and the maximum value is 1000. For more information about Limit , refer to the relevant sections in API overview. Example: 20
Offset	No	No	Uint64	Offset. The default value is 0. For more information about offset, refer to the relevant sections in API overview. Example: 0
DiskIds	No	No	Array of String	Query by one or more cloud disk IDs. Example: ["disk-6rag8jef"]
InnerSearch	No	No	String	Used to support search box search, internal parameters. Example: ins-9oxguxa

3. Output Parameters

Parameter name	Type	Description
TotalCount	Uint64	Number of block storages that meet the criteria. Example: 0
DiskSet	Array of BlockStorage	Detailed information list of block storage. Example:View
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.

DescribeDiskAssociatedSnapshots

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API (DescribeDiskAssociatedSnapshots) is used to query snapshots associated with a cloud disk.

Default interface request rate limit: 20 times per second.

Interface update time: 2024-07-16 06:42:29.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DescribeDiskAssociatedSnapshots
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
DiskId	Yes	No	String	ID of the cloud disk to be queried. Example: "disk-7q2iy7gp"
Offset	No	No	Uint64	Offset. The default value is 0. Example: 0
Limit	No	No	Uint64	Number of returned values. The default value is 20, and the maximum value is 100. Example: 20

3. Output Parameters

Parameter name	Type	Description
SnapshotSet	Array of Snapshot	List of Snapshots Associated with Cloud Disk Example:View
TotalSnapSize	Uint64	Total Size of Snapshots Associated with Cloud Disk Example: 100
TotalCount	Uint64	Total Number of Snapshots Associated with Cloud Disk Example: 1
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.

DescribeDiskConfigQuota

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API (DescribeDiskConfigQuota) is used to query the cloud disk quota.

Default interface request rate limit: 20 times per second.

Interface update time: 2024-08-09 13:00:09.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DescribeDiskConfigQuota
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
InquiryType	Yes	No	String	Query category. Value range: • INQUIRY_cloud disk_CONFIG: Query cloud disk configuration list • INQUIRY_CVM_CONFIG: Query configuration list of cloud disk with instances. Example: INQUIRY_CVM_CONFIG
Zones	No	No	Array of String	Queried configurations under one or more availability zones. Example: ["ap-region1-7"]
DiskChargeType	No	No	String	Payment mode. Value range: • PREPAID: prepaid • POSTPAID_BY_HOUR: pay-as-you-go. Example: POSTPAID_BY_HOUR
DiskTypes	No	No	Array of String	Hard disk media type. Value range: • CLOUD_BASIC: common cloud disk • CLOUD_PREMIUM: high-performance cloud disk • CLOUD_SSD: SSD cloud disk. Example: ["CLOUD_PREMIUM"]
DiskUsage	No	No	String	System disk or data disk. Value range: • SYSTEM_DISK: system disk • DATA_DISK: data disk. Example: DATA_DISK
DeviceClasses	No	No	Array of String	Instance model. Example: ["VSELF_5"]
InstanceFamilies	No	No	Array of String	Filter by instance model series like S1, I1, and M1. For details, see the instance types. Example: ["S3"]

Parameter name	Required	Allow NULL	Type	Description
CPU	No	No	Uint64	Number of instance CPU cores. Example: 2
Memory	No	No	Uint64	Instance memory size. Example: 4
InnerInquiryType	No	No	String	INQUIRY_RESIZE and INQUIRY_CREATE Example: INQUIRY_CBS_CONFIG
DedicatedClusterId	No	No	String	ID of a local dedicated cluster. Example: cluster-ay3cdplf

3. Output Parameters

Parameter name	Type	Description
DiskConfigSet	Array of DiskConfig	Cloud disk configuration list. Example:View
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
InternalError.ComponentError	Third-party Component Error

DescribeDiskSupportFeatures

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API is used to query the features supported by the cloud disk.

Default interface request rate limit: 20 times per second.

Interface update time: 2020-01-11 03:04:13.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DescribeDiskSupportFeatures
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface

3. Output Parameters

Parameter name	Type	Description
DiskFeature	DiskFeature	Detailed cloud disk features. Example:View
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

There is no error code related to business logic for this interface, For other error codes, see public ErrorCode.

DescribeDisks

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API (DescribeDisks) is used to query the cloud disk list.

- You can query details of the cloud disk by its ID, type, or status. These conditions are in an AND relationship. For detailed filter information, refer to the Filter .
- If parameters are empty, a certain number (specified by Limit , default is 20) of the cloud disk list for the current user is returned.

Default interface request rate limit: 20 times per second.

Interface update time: 2024-08-14 08:50:33.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DescribeDisks
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
DiskIds	No	No	Array of String	Query by one or more cloud disk IDs in the format like disk-11112222 . The parameters DiskIds and Filters cannot be specified at the same time. Example: ["disk-6rag8jef"]
Filters	No	No	Array of Filter	Filter criteria. The parameters DiskIds and Filters cannot be specified at the same time. • disk-usage - Array of String - Required: No - (Filter criteria) Filter by cloud disk type. (SYSTEM_DISK: system disk
Offset	No	No	Uint64	Offset. The default value is 0. Example: 0
Limit	No	No	Uint64	Number of returned values. The default value is 20, and the maximum value is 100. Example: 20
Order	No	No	String	Output cloud disk list sorting order. Value range: • ASC: ascending • DESC: descending. Example: ASC
OrderField	No	No	String	Field based on which cloud disk list is sorted. Value range: • CREATE_TIME: Sort by the cloud disk creation time • DEADLINE: Sort by the cloud disk expiration time Cloud disks are sorted by the creation time by default. Example: CREATE_TIME
ReturnBindAutoSnapshotPolicy	No	No	Bool	It indicates whether the cloud disk details need to return the ID of the regular snapshot policy to which the cloud disk is bound. TRUE indicates to return, and FALSE indicates not to return. Example: FALSE
InnerSearch	No	No	String	Used to support search box search, internal parameters. Example: disk-6rag8jef

3. Output Parameters

Parameter name	Type	Description
TotalCount	Uint64	Number of qualified cloud disks. Example: 1
DiskSet	Array of Disk	Detailed information list of cloud disks. Example:View
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
InvalidFilter	The specified filter is not supported.

DescribeDisksDeniedActions

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API (DescribeDisksDeniedActions) is used to query the operation mask of cloud disks.

- Based on the cloud disk IDs provided by the input parameter DiskIds, the operation mask of the corresponding cloud disk is returned.

Default interface request rate limit: 20 times per second.

Interface update time: 2024-07-16 07:40:36.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DescribeDisksDeniedActions
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
DiskIds	Yes	No	Array of String	Query by one or more cloud disk IDs in the format like disk-11112222 . Example: ["disk-6rag8jef"]

3. Output Parameters

Parameter name	Type	Description
DiskDeniedActionSet	Array of DiskDeniedAction	Set of Operation Masks for the Querying Cloud Disk Example:View
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

There is no error code related to business logic for this interface, For other error codes, see public ErrorCode.

DescribeInstancesDiskNum

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API (DescribeInstancesDiskNum) is used to query the number of cloud disks mounted on an instance.

- It supports batch operation. When multiple Cloud Virtual Machine instance IDs are input, the results will list the number of cloud disks mounted to each Cloud Virtual Machine.

Default interface request rate limit: 20 times per second.

Interface update time: 2025-08-28 17:51:25.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DescribeInstancesDiskNum
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
InstanceIds	Yes	No	Array of String	CVM instance ID. It can be queried through the DescribeInstances API. Example: ["ins-9w5d2buw","ins-jw0vit58"]
AttachMode	No	Yes	String	Mounting mode for I/O virtualization. Example: “VF”

3. Output Parameters

Parameter name	Type	Description
AttachDetail	Array of AttachDetail	Number of elastic cloud disks that are already mounted and can be mounted on each Cloud Virtual Machine. Example:View
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.

Error Code	Description
MissingParameter	Missing parameter. The request does not include the mandatory parameter.

DescribeUserDiskResources

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API (DescribeUserTotalResources) is used to query the user's number of cloud disks, number of snapshots, and quotas across all regions.

Default interface request rate limit: 20 times per second.

Interface update time: 2025-08-28 18:25:05.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DescribeUserDiskResources
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
Regions	No	No	Array of String	It specifies the region to query. Multiple regions can be specified. If this parameter is not input, all regions will be queried by default. Example: region1

3. Output Parameters

Parameter name	Type	Description
ResourcesDetailSet	Array of ResourcesDetail	Description of resource details in the Nth region. Example:View
DiskOverview	DiskOverview	It describes the overall usage of the cloud disk in the specified region. Example:View
SnapOverview	SnapOverview	It describes overall usage of snapshots in the specified region. Example:View
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.
MissingParameter	Missing parameter. The request does not include the mandatory parameter.

DescribeZoneDiskConfigInfos

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API is used to query whether the corresponding cloud disk resources can be created.

Default interface request rate limit: 20 times per second.

Interface update time: 2025-08-28 18:29:07.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DescribeZoneDiskConfigInfos
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
DiskChargeType	Yes	No	String	Payment mode. Example: POSTPAID_BY_HOUR
DiskConfigs	Yes	No	Array of DiskConfigSet	Cloud disk configuration list. Example:View
Zone	Yes	No	String	Availability zone. Example: az1

3. Output Parameters

Parameter name	Type	Description
DiskConfigSetQuota	Int64	Number of Cloud Disks That Can Be Created Example: 20
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.
MissingParameter	Missing parameter. The request does not include the mandatory parameter.

DetachDisks

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API is used to unmount a cloud disk.

- Batch operations are supported for unmounting multiple cloud disks from the same CVM. If any one in multiple cloud disks cannot be unmounted, the operation will not be executed and a specific error code will be returned.
- This API is asynchronous. When a successful response is returned for a request, the cloud disk is not immediately unmounted from the CVM. The status of the corresponding cloud disk can be queried through the DescribeDisks API. If the status of the cloud disk changes from ATTACHED to UNATTACHED, it indicates that the cloud disk is unmounted successfully.

Default interface request rate limit: 20 times per second.

Interface update time: 2025-04-03 18:30:31.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DetachDisks
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
DiskIds	Yes	No	Array of String	ID of the cloud disk to be unmounted. It can be queried through the DescribeDisks API. Up to 10 elastic cloud disks can be unmounted in a single request. Example: ["disk-l7mxecdc"]
InstanceId	No	No	String	For non- shared cloud disks, this parameter will be ignored. For shared cloud disks, this parameter indicates from which CVM instance the cloud disk will be unmounted. Example: ins-arxix3xm

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidDiskId.NotFound	The input DiskId does not exist.
InvalidDisk.TypeError	Incorrect cloud disk type.
InvalidDisk.NotPortable	Non-elastic cloud disks are not supported.
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
InternalError.ResourceOpFailed	The operation on the resource failed. For specific error information, view the Message field in the error description. Try again later or contact customer service for help.
InvalidDisk.NotSupported	The cloud disk does not support this operation.
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.
InvalidParameterValue.LimitExceeded	The parameter value is out of range
ResourceUnavailable.TypeError	Disk type error.

InitializeDisks

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API is used to reinitialize cloud disks.

Default interface request rate limit: 20 times per second.

Interface update time: 2025-08-28 18:38:39.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: InitializeDisks
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
DiskIds	Yes	No	Array of String	List of cloud disk IDs to be reinitialized. Up to 20 disks can be initialized for the first time. Example: ["disk-12sdf234"]

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

There is no error code related to business logic for this interface, For other error codes, see public ErrorCode.

InquirePriceModifyDiskExtraPerformance

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API (InquirePriceModifyDiskExtraPerformance) is used to query the price for adjusting cloud disk extra performance.

Default interface request rate limit: 20 times per second.

Interface update time: 2024-07-16 07:04:50.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: InquirePriceModifyDiskExtraPerformance
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
ThroughputPerformance	Yes	No	Uint64	Additional cloud disk performance value to be purchased (in MB/s). Example: 200
DiskId	Yes	No	String	Cloud disk ID. It can be queried through the DescribeDisks API. Example: disk-dritwhhm

3. Output Parameters

Parameter name	Type	Description
DiskPrice	Price	It describes the price when adjusting extra performance for cloud disks. Example:View
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
InvalidDisk.Expire	The cloud disk has expired.
InvalidParameter	Parameter error

InquiryPriceCreateDisks

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API (InquiryPriceCreateDisks) is used to query the price for creating a cloud disk.

- It supports querying the price for creating multiple cloud disks, in which case the returned result is the total price.

Default interface request rate limit: 20 times per second.

Interface update time: 2025-08-28 18:42:50.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: InquiryPriceCreateDisks
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
DiskType	Yes	No	String	Cloud disk type. Value range: • Common cloud disk: CLOUD_BASIC • High-performance cloud disk: CLOUD_PREMIUM • SSD Cloud Disk: CLOUD_SSD. Example: CLOUD_TSSD
DiskSize	Yes	No	Uint64	Cloud disk size. Value range: Common cloud disk: 10 GB to 4000 GB; High-performance cloud disk: 50 GB to 4000 GB; SSD cloud disk: 100 GB to 4000 GB. The step is 10 GB. Example: 500
DiskChargeType	Yes	No	String	Payment mode. Currently, the only option is PREPAID. Example: PREPAID
DiskChargePrepaid	No	No	DiskChargePrepaid	Settings related to the prepaid parameter. Through this parameter, you can specify the purchase duration of the monthly subscribed cloud disk. It is a required parameter for creating a prepaid cloud disk. Example:View
DiskCount	No	No	Uint64	Number of cloud disks to be purchased. If this parameter is not specified, the default value is 1. Example: 1
ProjectId	No	No	Uint64	Project ID of the cloud disk. Example: 0
ThroughputPerformance	No	No	Uint64	Performance size additionally purchased. Example: 100
Zone	No	No	String	Availability Zone Example: az1
DiskBackupQuota	No	No	Uint64	Quota of cloud disk backup points. Example: 10

3. Output Parameters

Parameter name	Type	Description
DiskPrice	Price	It describes the price of purchasing new cloud disks. Example:View
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.
MissingParameter	Missing parameter. The request does not include the mandatory parameter.

InquiryPriceModifyDiskAttributes

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API is used to query the price for modifying the cloud disk type.

- Currently, type modification is supported only for elastic cloud disks. If the Portable field returned by calling the DescribeDisks API is true, it indicates an elastic cloud disk.
- Currently, cloud disk types can only be upgraded but cannot be downgraded, specifically as follows:
  - CLOUD_BASIC changes to CLOUD_PREMIUM;
  - CLOUD_BASIC changes to CLOUD_SSD;
  - CLOUD_PREMIUM changes to CLOUD_SSD.

Default interface request rate limit: 20 times per second.

Interface update time: 2024-07-16 07:01:21.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: InquiryPriceModifyDiskAttributes
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
DiskIds	Yes	No	Array of String	List of cloud disk instance IDs to be migrated. Currently, only one cloud disk can be input at a time. Example: ["disk-6rag8jef"]
DiskType	Yes	No	String	Target type of cloud disk change. Value range: • CLOUD_PREMIUM: high-performance cloud disk • CLOUD_SSD: SSD cloud disk. Type downgrading is not currently supported. Example: CLOUD_SSD
ProjectId	No	No	Uint64	Cloud disk's project ID. If this parameter is input, it is only used for authentication. Example: 1002189

3. Output Parameters

Parameter name	Type	Description
DiskPrice	Price	It describes the price of changing cloud disk types. Example:View
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidDisk.NotPortable	Non-elastic cloud disks are not supported.
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.
InvalidDisk.Expire	The cloud disk has expired.
ResourceInUse.DiskMigrating	The cloud disk is migrating.

InquiryPriceResizeDisk

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API (InquiryPriceResizeDisk) is used to query the price for cloud disk scaling.

- Only scaling price query of prepaid cloud disks is supported.

Default interface request rate limit: 20 times per second.

Interface update time: 2024-07-16 07:43:52.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: InquiryPriceResizeDisk
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
DiskId	Yes	No	String	Cloud disk ID. It can be queried through the DescribeDisks API. Example: disk-dw0bbzws
DiskSize	Yes	No	Uint64	Size after cloud disk scaling, in GB. It cannot be smaller than the current cloud disk size. Value range: Common cloud disk: 10 GB to 4000 GB; high performance cloud disk: 50 GB to 4000 GB; SSD cloud disk: 100 GB to 4000 GB. The step is 10 GB. Example: 200
ProjectId	No	No	Uint64	Cloud disk's project ID. If this parameter is input, it is only used for authentication. Example: 1002189

3. Output Parameters

Parameter name	Type	Description
DiskPrice	Price	It describes the price for scaling cloud disks. Example:View
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
InvalidDisk.NotPortable	Non-elastic cloud disks are not supported.
InvalidInstanceId.NotFound	The input InstanceId does not exist.
InvalidDisk.Expire	The cloud disk has expired.

ModifyDiskAttributes

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API is used to modify the attributes of a cloud disk.

*Only the project ID of cloud disks can be modified. The project ID of cloud disks created together with CVM instances is the same as that of the CVM instances. It can be queried by calling the DescribeDisks API. For details, see the description of the output parameter Portable.
*The cloud disk name is used only for snapshot management by users. The cloud platform does not use this name for submitting tickets or performing snapshot management operations.
*Batch operations are supported. If multiple cloud disk IDs are passed in, the attribute after modification applies to all corresponding cloud disks. If any cloud disk does not support the operation, the operation will not be performed, and a specific error code will be returned.

Default interface request rate limit: 20 times per second.

Interface update time: 2025-03-29 00:42:53.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: ModifyDiskAttributes
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
DiskIds	Yes	No	Array of String	ID of one or more cloud disks to be operated on. If multiple cloud disk IDs are input, only modification to the same attribute for all cloud disks is supported. Example: ["disk-7q2iy7gp"]
ProjectId	No	No	Uint64	New project ID of the cloud disk. Modification of the project ID is supported only for elastic cloud disks. The available projects and their IDs can be queried through the DescribeProjects API. Example: 0
DiskName	No	No	String	New cloud disk name. Example: ARM-centOS_System disk
Portable	No	No	Bool	It indicates whether it is an elastic cloud disk. FALSE indicates a non-elastic cloud disk, and TRUE indicates an elastic cloud disk. Only non-elastic cloud disks can be converted into elastic cloud disks. Example: TRUE
DeleteWithInstance	No	No	Bool	It indicates whether the cloud disk, once successfully mounted to a cloud host, is destroyed along with the cloud host. TRUE indicates it will be destroyed along with the cloud server, and FALSE indicates it will not. It only supports pay-as-you-go cloud disk data disks. Example: FALSE
DiskType	No	No	String	When changing the cloud disk type, this parameter can be input in to indicate the target type. Value range: • CLOUD_PREMIUM: high-performance cloud disk • CLOUD_SSD: SSD cloud disk. Currently, batch type change is not supported, that is, when DiskType is input, DiskIds only supports entering one cloud disk ID; Changing the cloud disk type does not support changing other attributes at the same time. Example: CLOUD_SSD

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidDiskId.NotFound	The input DiskId does not exist.
InvalidDisk.NotSupported	The cloud disk does not support this operation.
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
ResourceNotFound.NotFound	The resource does not exist.

ModifyDiskExtraPerformance

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API (ModifyDiskExtraPerformance) is used to adjust the extra performance of cloud disks.

- Currently, it only supports ultra-fast SSD cloud disks (CLOUD_TSSD) and high performance SSD cloud disk (CLOUD_HSSD).

Default interface request rate limit: 20 times per second.

Interface update time: 2024-07-16 06:44:52.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: ModifyDiskExtraPerformance
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
ThroughputPerformance	Yes	No	Uint64	Additional cloud disk performance value to be purchased (in MB/s). Example: 300
DiskId	Yes	No	String	ID of the cloud disk that requires creating a snapshot. It can be queried through the DescribeDisks API. Example: disk-dritwhhm

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameter	Parameter error
ResourceInsufficient	Insufficient resources
UnsupportedOperation	Operation is not supported
InvalidDiskId.NotFound	The input DiskId does not exist.
InvalidDisk.Busy	The cloud disk is busy. Try again later.
InvalidDisk.Expire	The cloud disk has expired.
InvalidDisk.NotSupported	The cloud disk does not support this operation.
InvalidAccount.InsufficientBalance	Insufficient account balance

ResizeDisk

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API is used to scale out a cloud disk.

- It only supports scale-out of elastic cloud disks. The cloud disk type can be queried through the DescribeDisks API. For details, see the explanation of the Portable field in the output parameters. A cloud disk created along with CVM should be scaled out through the [ResizeInstanceDisks](/tcloud/api/Compute/CVM/APIs/CVM (cvm)/Version (2017-03-12)/Instance APIs/ResizeInstanceDisks) API.
- This API is asynchronous. When a response is returned successfully, the cloud disk is not immediately scaled out to the specified size. The status of the corresponding cloud disk can be queried through the DescribeDisks API. If the status of the cloud disk is EXPANDING, it indicates that the cloud disk is being scaled out. If the status changes to UNATTACHED, it indicates that the cloud disk is scaled out successfully.

Default interface request rate limit: 20 times per second.

Interface update time: 2025-03-29 00:43:48.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: ResizeDisk
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
DiskId	Yes	No	String	Cloud disk ID. It can be queried through the DescribeDisks API. Example: disk-7q2iy7gp
DiskSize	Yes	No	Uint64	Size of the cloud disk after scaling, in GB, which must be greater than the current cloud disk size. Value range: Common cloud disk: 10 GB to 4000 GB; High-performance cloud disk: 50 GB to 4000 GB; SSD cloud disk: 100 GB to 4000 GB. The step is 10 GB. Example: 100

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
InvalidDisk.NotSupported	The cloud disk does not support this operation.
InvalidDiskId.NotFound	The input DiskId does not exist.
InvalidDisk.Busy	The cloud disk is busy. Try again later.
InvalidDisk.Expire	The cloud disk has expired.
ResourceInsufficient.OverQuota	Insufficient quota

SwitchParameterCreateDisks

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API (SwitchParameterCreateDisks) is used to obtain order parameters for creating cloud disks.

Default interface request rate limit: 20 times per second.

Interface update time: 2024-07-16 06:58:53.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: SwitchParameterCreateDisks
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
DiskSize	No	No	Uint64	Cloud disk size, in GB. ▪ If SnapshotId is input, DiskSize is optional. In this case, the size of the new cloud disk will be the size of the snapshot. ▪ If SnapshotId and DiskSize are both input, then the cloud disk size must be greater than or equal to the snapshot size. ▪ The range of cloud disk sizes: Common cloud disk: 10 GB to 4000 GB; high performance cloud disk: 50 GB to 4000 GB; SSD cloud disk: 100 GB to 4000 GB. The step is 10 GB. Example: 10
DiskType	Yes	No	String	Hard disk media type. Value range: ▪ CLOUD_BASIC: common cloud disk ▪ CLOUD_PREMIUM: high-performance cloud disk ▪ CLOUD_SSD: SSD cloud disk. Example: CLOUD_BASIC
DiskCount	No	No	Uint64	Number of cloud disks purchased. There is a limit on the maximum number of cloud disks that can be created in a single request. For details, see the cloud disk use limits. The default value is 1. Example: 1
DiskChargeType	Yes	No	String	Payment mode. Currently, the only option is PREPAID. Example: PREPAID
DiskUsage	No	No	String	It specifies to create a system disk or data disk. Value range: ▪ SYSTEM_DISK: system disk ▪ DATA_DISK: data disk. Example: DATA_DISK
DiskChargePrepaid	Yes	No	DiskChargePrepaid	Prepaid mode, that is, the settings for the monthly subscription-related parameters. This parameter allows specifying the purchase duration of the monthly subscribed cloud disk, and whether to set auto-renewal, among other attributes. Example:View
Placement	Yes	No	Placement	The location of cloud disk. This parameter can specify the cloud disk availability zone and project. If the project is not specified, the cloud disk will be created under the default project. Example:View
SnapshotId	No	No	String	Snapshot ID. If this parameter is input, a cloud disk is created based on this snapshot. The snapshot type must be a data disk snapshot. Snapshots can be queried through the DescribeSnapshots API. For details, see the explanation of the output parameter DiskUsage. Example: snap-4zgc8cnk
DiskName	No	No	String	Display name of the cloud disk. If not specified, it defaults to "unnamed". The maximum length cannot exceed 60 bytes. Example: Untitled
Encrypt	No	No	String	Input this parameter to create an encrypted cloud disk. The value is fixed to ENCRYPT. Example: ENCRYPT
Tags	No	No	Array of Tag	Tag to which cloud disks are bound. Example:View
AutoSnapshotPolicyId	No	No	String	ID of the regular snapshot policy. When this parameter is input, the cloud disk will automatically be associated with the specified regular snapshot policy upon successful creation. Example: asp-mrsrn243
Shareable	No	No	Bool	Optional. The default value is False. When True is input, the cloud disk will be created as a shared cloud disk. Example: False
DiskStoragePoolGroup	No	No	String	Resource Pool Group Example: pool-ec
PlatformProjectId	No	No	String	Project ID Example: pr-42749a13

3. Output Parameters

Parameter name	Type	Description
DiskOrder	Array of DiskOrder	Billing order parameters. Example:View
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.

SwitchParameterResizeDisk

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API (SwitchParameterResizeDisk) is used to obtain the order parameters for scaling the cloud disk.

Default interface request rate limit: 20 times per second.

Interface update time: 2025-09-12 16:34:57.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: SwitchParameterResizeDisk
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
DiskId	Yes	No	String	Cloud disk ID. It can be queried through the DescribeDisks API. Example: disk-7q2iy7gp
DiskSize	Yes	No	Uint64	The disk size after scaling. It must be greater than the current value, with a maximum of 4000 GB and a step of 10 GB. Example: 100

3. Output Parameters

Parameter name	Type	Description
DiskOrder	Array of DiskOrder	Billing order parameters. Example:View
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
InvalidDisk.NotSupported	The cloud disk does not support this operation.
InvalidDiskId.NotFound	The input DiskId does not exist.

TerminateDisks

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API (TerminateDisks) is used to terminate cloud disks.

- Currently, only monthly subscribed cloud disks can be terminated.
- It supports batch operation. The upper limit for batch cloud disk requests is 50 at a time. If there are cloud disks in the batch not allowed for the operation, a specific error code will be returned for the request.

Default interface request rate limit: 20 times per second.

Interface update time: 2025-03-29 00:44:16.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: TerminateDisks
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
DiskIds	Yes	No	Array of String	List of cloud disks IDs to be returned. Example: ["disk-l7mxecdc"]

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
InvalidDisk.NotSupportRefund	The cloud disk cannot be returned.
InternalError.FailQueryResource	Failed to query the resource.
InvalidDisk.RepeatRefund	The cloud disk has been returned before and cannot be returned again.
InvalidDisk.Expire	The cloud disk has expired.
InsufficientRefundQuota	The number of returned cloud disks has reached the upper limit and no more cloud disks can be returned.
ResourceBusy	The resource is busy. Try again later.

Other APIs DescribeDedicatedClusterDiskStatistics

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API is used to obtain statistics on the cloud disk capacity of the CVM CDC cluster.

Default interface request rate limit: 20 times per second.

Interface update time: 2025-09-12 16:40:08.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DescribeDedicatedClusterDiskStatistics
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
DedicatedClusterId	Yes	No	String	CVM CDC Cluster ID Example: cluster-ay3cdplf

3. Output Parameters

Parameter name	Type	Description
DedicatedClusterDiskStatisticSet	Array of DedicatedClusterDiskStatistic	CVM CDC dedicated cluster cloud disk statistics. Example:View
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

There is no error code related to business logic for this interface, For other error codes, see public ErrorCode.

ModifyDisksSynchronizeFlag

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API is used to check whether the user has permission to operate the cloud block storage resource.

Default interface request rate limit: 20 times per second.

Interface update time: 2025-09-12 16:47:55.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: ModifyDisksSynchronizeFlag
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
DiskIds	Yes	No	Array of String	List of cloud disks to modify tags. Example: disk-qweasdzx
UpdateAction	Yes	No	String	modification action. Example: ADD
DiskFlag	Yes	No	String	Modified tag value. Example: SOURCE_SYNCHRONIZE_DISK

3. Output Parameters

Parameter name	Type	Description
ModifySuccessDiskIds	Array of String	Success List Example: disk-qweasdzx
ModifyFailDiskIds	Array of String	List of Failures Example: disk-qweasdzx
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

There is no error code related to business logic for this interface, For other error codes, see public ErrorCode.

Snapshot APIs
ApplySnapshot

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API is used to roll back a snapshot to the original cloud disk.

- It only supports rollback to the original cloud disk. For data disk snapshots, if you need to copy snapshot data to another cloud disk, use the [CreateDisks](../CBS APIs/CreateDisks) API to create an elastic cloud disk and then copy the snapshot data to the purchased cloud disk.
- Snapshots used for rollback must be in the NORMAL state. The snapshot status can be queried through the DescribeSnapshots API. For details, see the explanation of the SnapshotState field in the output parameters.
- For an elastic cloud disk, it must be in the unmounted state. The cloud disk mounting status can be queried through the [DescribeDisks](../CBS APIs/DescribeDisks) API. For details, see the explanation of the Attached field. For a non-elastic cloud disk purchased along with CVM, the CVM must be in the shutdown state. The CVM status can be queried through the [DescribeInstancesStatus](/tcloud/api/Compute/CVM/APIs/CVM (cvm)/Version (2017-03-12)/Instance APIs/DescribeInstancesStatus) API.

Default interface request rate limit: 20 times per second.

Interface update time: 2025-03-14 01:35:19.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: ApplySnapshot
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
SnapshotId	Yes	No	String	Snapshot ID. It can be queried through the DescribeSnapshots API. Example: snap-4zgc8cnk
DiskId	Yes	No	String	ID of the original cloud disk for the snapshot. It can be queried through the DescribeDisks API. Example: disk-7q2iy7gp
AutoStartInstance	No	No	Bool	Whether to start the instance after rollback is completed. Example: true
AutoStopInstance	No	No	Bool	Whether rollback with instance shut down is supported. Example: true

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
InvalidDisk.NotSupported	The cloud disk does not support this operation.
InvalidDiskId.NotFound	The input DiskId does not exist.
InvalidInstanceId.NotFound	The input InstanceId does not exist.
InvalidDisk.Busy	The cloud disk is busy. Try again later.
InvalidDisk.SnapshotCreating	The cloud disk is creating a snapshot. Try again later.
InvalidSnapshotId.NotFound	The input SnapshotId does not exist.
InvalidSnapshot.NotSupported	The snapshot does not support this operation.
InvalidParameter.DiskSizeNotMatch	The cloud disk size does not match the snapshot size.
ResourceBusy	The resource is busy. Try again later.
InvalidInstance.NotSupported	The Cloud Virtual Machine does not support mounting a cloud disk.

ApplySnapshotGroup

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API is used to roll back the snapshot group.

- i. You can roll back all or part of the disks in the snapshot group.
- ii. If the disks to be rolled back contain mounted disks, these disks must be mounted to the same instance, and the instance must be shut down before rollback can proceed.
- iii. Rollback is an asynchronous operation. A successful API return does not mean the rollback is successful. You can query the result of the asynchronous rollback task through the DescribeSnapshotGroups API.

Default interface request rate limit: 20 times per second.

Interface update time: 2021-12-04 04:04:15.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: ApplySnapshotGroup
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
SnapshotGroupId	Yes	No	String	ID of Snapshot Group Rolled Back Example: csnap-csoplz7u

Parameter name	Required	Allow NULL	Type	Description
ApplyDisks	Yes	No	Array of ApplyDisk	ID of Snapshot Rolled Back, Cloud Disk ID List Example:View

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
UnsupportedOperation.SnapshotGroupDiskAttachMultiInstance	The original disk of the snapshot in the snapshot group is not mounted to the same instance.

BindAutoSnapshotPolicy

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API (BindAutoSnapshotPolicy) is used to bind a cloud disk to a specified regular snapshot policy.

- A maximum of 10 regular snapshot policies can be created in each region, and each regular snapshot policy can be bound to up to 80 cloud disks.
- Regular snapshots will not be created for cloud disks bound to a regular snapshot policy that is in an unused state (that is, the elastic cloud disk is not mounted or the host of a non-elastic cloud disk is shut down).

Default interface request rate limit: 20 times per second.

Interface update time: 2025-03-29 00:39:19.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: BindAutoSnapshotPolicy
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AutoSnapshotPolicyId	Yes	No	String	Regular snapshot policy ID to be bound. Example: asp-6xtxct37
DiskIds	No	No	Array of String	ID List of cloud disks to be bound. Up to 80 cloud disks can be bound in a single request. Example: ["disk-7q2iy7gp"]
InstanceIds	No	No	Array of String	List of instances to be bound. Up to 100 instances can be bound at a time. Example: ["ins-locaa0iu"]

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidDisk.NotSupported	The cloud disk does not support this operation.

Error Code	Description
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.
InvalidParameterValue.LimitExceeded	The parameter value is out of range
InvalidAutoSnapshotPolicyId.NotFound	The input AutoSnapshotPolicyId does not exist.

CopyAutoSnapshotPolicyCrossAccount

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

Cross-account Replication of Automatic Snapshot Policy

Default interface request rate limit: 20 times per second.

Interface update time: 2025-08-28 01:27:46.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: CopyAutoSnapshotPolicyCrossAccount
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AutoSnapshotPolicyIds	Yes	No	Array of String	List of Regular Snapshot Policy IDs for Backup Example: ["asp-ljylgbyo"]
TargetAccountUin	Yes	No	String	Target Account UIN Example: 597568700

3. Output Parameters

Parameter name	Type	Description

Parameter name	Type	Description
AutoSnapshotPolicyIds	Array of String	ID List of Regular Snapshot Policies Already Backed Up Example: ["asp-ljylgbyo"]
TargetAccountUin	String	Target Account UIN Example: 597568700
Code	String	Returned Status Code Example: SUCCESS
MESSAGE	String	Return Message Example: SUCCESS
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
InvalidParameter	Parameter error
ResourceInsufficient	Insufficient resources

CopySnapshotCrossAccount

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

Snapshot Cross-account Replication

Default interface request rate limit: 20 times per second.

Interface update time: 2025-08-28 17:40:50.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: CopySnapshotCrossAccount
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
SnapshotId	Yes	No	String	Snapshot instance ID copied. Example: snap-4zgc8cnk
TargetAccountUin	Yes	No	String	Target Account UIN Example: 597568700
SnapshotName	No	No	String	Name of New Snapshot After Replication Example: new_snap_name

3. Output Parameters

Parameter name	Type	Description
SnapshotId	String	New Snapshot Instance ID After Replication. Example: snap-4zgc8cnk
TargetAccountUin	String	Target account UIN. Example: 597568700
Code	String	Success status Example: Success
Message	String	Additional Information Example: Copy task has existed
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
InvalidDisk.TypeError	Incorrect cloud disk type.
InvalidSnapshot.NotSupported	The snapshot does not support this operation.

CopySnapshotCrossRegions

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API is used to replicate a snapshot across regions.

- This API is asynchronous. When a request for cross-region replication is successfully issued, a new snapshot ID is returned, but the snapshot is not immediately replicated to the target region. The status of the new snapshot in the target region can be queried through the DescribeSnapshots API to determine whether the replication is completed. If the snapshot status is NORMAL, it indicates that the snapshot is replicated successfully.

Default interface request rate limit: 20 times per second.

Interface update time: 2025-05-06 18:49:18.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: CopySnapshotCrossRegions
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
SnapshotId	No	No	String	ID of the source snapshot that requires cross-region replication. It can be queried through the DescribeSnapshots API. Example: snap-ckgjwkqh
DestinationRegions	Yes	No	Array of String	Target region to which the snapshot will be replicated. The standard values for each region can be queried through the DescribeRegions API. The input region should support snapshots. Example: ["region1"]
SnapshotName	No	No	String	Name of the new copied snapshot. If this parameter is not input, the default value is "Copied snap-11112222 from region name". Example: snap-hfdadhusads

3. Output Parameters

Parameter name	Type	Description
SnapshotCopyResultSet	Array of SnapshotCopyResult	Result of snapshot cross-region replication. If the request is successfully issued, the new snapshot ID of the corresponding region will be returned; otherwise, Error is returned. Example:View
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidAccount.InsufficientBalance	Insufficient account balance
InsufficientSnapshotQuota	The snapshot quota is insufficient.
InvalidSnapshotId.NotFound	The input SnapshotId does not exist.
InvalidSnapshot.NotSupported	The snapshot does not support this operation.

CreateAutoSnapshotPolicy

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API is used to create a scheduled snapshot policy.

*A maximum of 10 scheduled snapshot policies can be created in each region.
*The number and capacity of snapshots that can be created in each region are restricted. For details, see the notes on the Schedule Snapshot Policy page of the CBS console.

Default interface request rate limit: 20 times per second.

Interface update time: 2025-03-28 01:10:14.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: CreateAutoSnapshotPolicy
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
Policy	Yes	No	Array of Policy	Execution policy of regular snapshots. Example:View
AutoSnapshotPolicyName	No	No	String	Name of the regular snapshot policy to be created. If this parameter is not input, it defaults to "unnamed". The maximum length cannot exceed 60 bytes. Example: backup_data_friday
IsActivated	No	No	Bool	It indicates whether the regular snapshot policy is activated. FALSE indicates not activated, and TRUE indicates activated. The default value is TRUE. Example: false
IsPermanent	No	No	Bool	It indicates whether the snapshot created by the regular snapshot policy is permanently retained. FALSE indicates not permanently retained, and TRUE indicates permanently retained. The default value is FALSE. Example: false
RetentionDays	No	No	Uint64	Retention days for snapshots created by the regular snapshot policy. The default value is 7 days. This parameter cannot conflict with the IsPermanent parameter. If the regular snapshot policy is set to retain permanently, this parameter must be set to 0. Example: 20
DryRun	No	No	Bool	It indicates whether to create a regular snapshot policy. TRUE indicates only to access the start time of the backup initially, without actually creating a regular snapshot policy. FALSE indicates creating a policy. The default value is FALSE. Example: false
RetentionMonths	No	No	Uint64	Number of months the snapshot created by this regular snapshot policy can be retained. This parameter cannot conflict with the IsPermanent/RetentionDays parameters. Example: 5
RetentionAmount	No	No	Uint64	Maximum number of snapshots that can be retained through this regular snapshot policy. If this number is exceeded, the earliest created snapshot will be automatically deleted. This parameter cannot conflict with the IsPermanent parameter. Example: 9
AdvancedRetentionPolicy	No	No	AdvancedRetentionPolicy	Advanced retention policy for regular snapshot. This parameter cannot conflict with the IsPermanent parameter. Example:View
IsCrossRegion	No	No	Bool	Whether to enable cross-region snapshot replication for this scheduled snapshot policy. Default value: 0, indicating that it is disabled. Example: true
DestinationRegions	No	No	Array of String	Target region for cross-region snapshot replication. Example: ["region1"]
CrossRegionIsPermanent	No	No	Bool	Whether cross-region snapshots created based on this scheduled snapshot policy are permanently retained. Default value: 0, indicating that such snapshots are not retained permanently. Example: true
CrossRegionRetentionDays	No	No	Uint64	Retention days for cross-region snapshots created by using this scheduled snapshot policy. The default value is 7 days. This parameter cannot conflict with the CrossRegionIsPermanent parameter. Example: 7

3. Output Parameters

Parameter name	Type	Description
AutoSnapshotPolicyId	String	ID of the created regular snapshot policy. Example: asp-1lebc9r3
NextTriggerTime	String	First backup start time. Example: 2018-08-08 16:00:00
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameterValue.LimitExceeded	The parameter value is out of range
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
AutoSnapshotPolicyOutOfQuota	The number of regular snapshot policies has reached the upper limit.

CreateSnapshot

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API is used to create a snapshot for a specified cloud disk.

- Only cloud disks with the snapshot capability can create snapshots. The snapshot capability of a cloud disk can be queried through the [DescribeDisks](../CBS APIs/DescribeDisks) API. For details, see the SnapshotAbility field.
- For the limit on the number of snapshots that can be created, see the product use limits.

Default interface request rate limit: 20 times per second.

Interface update time: 2025-03-29 00:40:39.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: CreateSnapshot
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
DiskId	Yes	No	String	ID of the cloud disk that requires creating a snapshot. It can be queried through the DescribeDisks API. Example: disk-7q2iy7gp
SnapshotName	No	No	String	Snapshot name. If this parameter is not input, the new snapshot name defaults to "unnamed". Example: Snapshot-for-io
ProjectId	No	No	Int64	Project ID of the snapshot. If this parameter is not input, it is consistent with the cloud disk project ID by default. Example: 0
Tags	No	No	Array of Tag	List of tags to which the snapshot is bound. If this parameter is not input, it is consistent with the cloud disk tag by default. Example:View
PlatformProjectId	No	No	String	Project ID in the internal environment. It is unrelated to ProjectId in the public cloud. If it is not specified, it remains consistent with the CBS project ID. Example: pr-42749a13
CreateInCluster	No	No	Bool	Whether to create the snapshot in the Outpost cluster of the cloud disk. Example: true

3. Output Parameters

Parameter name	Type	Description
SnapshotId	String	Created snapshot ID. Example: snap-3xvkxyxp
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidDisk.NotSupportSnapshot	The cloud disk does not have the snapshot capability.
InvalidDisk.SnapshotCreating	The cloud disk is creating a snapshot. Try again later.
InvalidDisk.TypeError	Incorrect cloud disk type.
InvalidDiskId.NotFound	The input DiskId does not exist.
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.
InternalError.ComponentError	Third-party Component Error
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
ResourceInUse.DiskMigrating	The cloud disk is migrating.
InsufficientSnapshotQuota	The snapshot quota is insufficient.
InvalidDisk.NotSupported	The cloud disk does not support this operation.
InvalidAccount.InsufficientBalance	Insufficient account balance
ResourceNotFound.NotFound	The resource does not exist.

CreateSnapshotGroup

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API is used to create snapshot groups.

- The disk list for the snapshot group creation must be mounted to the same instance. You can select all or some of the disks mounted to an instance to create a snapshot group.

Default interface request rate limit: 20 times per second.

Interface update time: 2025-03-29 00:41:04.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: CreateSnapshotGroup
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
DiskIds	Yes	No	Array of String	ID List of Cloud Disks for Which Snapshot Group Needing to Be Created Example: ["disk-6rag8jef"]
SnapshotGroupName	No	No	String	Snapshot Group Name Example: snp092101

3. Output Parameters

Parameter name	Type	Description
SnapshotGroupId	String	Snapshot Group ID Example: csnap-oavcw09g
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
InternalError.ComponentError	Third-party Component Error
InvalidDisk.NotSupported	The cloud disk does not support this operation.
InvalidParameter	Parameter error
InvalidDisk.SnapshotCreating	The cloud disk is creating a snapshot. Try again later.
InvalidDiskId.NotFound	The input DiskId does not exist.
InvalidAccount.InsufficientBalance	Insufficient account balance
InsufficientSnapshotQuota	The snapshot quota is insufficient.
ResourceInUse.DiskMigrating	The cloud disk is migrating.
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.
ResourceNotFound.NotFound	The resource does not exist.

DeleteAutoSnapshotPolicies

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API (DeleteAutoSnapshotPolicies) is used to delete regular snapshot policies.

- It supports batch operation. If there are multiple regular snapshot policies that cannot be deleted, the operation will not be executed, and a specific error code will be returned.

Default interface request rate limit: 20 times per second.

Interface update time: 2025-03-29 00:41:26.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DeleteAutoSnapshotPolicies
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AutoSnapshotPolicyIds	Yes	No	Array of String	List of regular snapshot policy IDs to delete. Example: ["asp-mrsrn243","asp-3lpl7ev3"]

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
InvalidParameterValue.LimitExceeded	The parameter value is out of range
InvalidAutoSnapshotPolicyId.NotFound	The input AutoSnapshotPolicyId does not exist.

DeleteSnapshotGroup

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API is used to delete a snapshot group. Only one snapshot group can be deleted per call.

- By default, all snapshots in the snapshot group will be deleted.
- If a snapshot in the snapshot group has an associated image, the deletion will fail and no snapshot will be deleted. If simultaneous deletion of images bound to the snapshot is required, the parameter DeleteBindImages=true can be specified.

Default interface request rate limit: 20 times per second.

Interface update time: 2021-12-04 03:48:05.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DeleteSnapshotGroup
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
SnapshotGroupId	Yes	No	String	Snapshot Group ID Example: csnap-gv9naz6b
DeleteBindImages	No	No	Bool	It indicates whether to delete the image associated with the snapshot group at the same time. The value false indicates not deleting the image bound to the snapshot group. In this case, if the snapshot group has bound images, the deletion will fail. The value true indicates deleting the image bound to the snapshot group at the same time. The default value is false. Example: true
OnlyDismiss	No	No	Bool	Dissolve the snapshot group.The value true indicates only dissolving the snapshot group without deleting the snapshots associated with the group. The value false indicates deleting the snapshots associated with the group as well. The default value is false. Example: true

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
UnsupportedOperation.SnapshotHasBindedImage	The snapshot is already associated with an image.
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
InvalidParameter	Parameter error

DeleteSnapshots

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API is used to delete a snapshot.

- The snapshot must be in the NORMAL state. The snapshot status can be queried through the DescribeSnapshots API. For details, see the explanation of the SnapshotState field in the output parameters.
- Batch operations are supported. If any one in multiple snapshots cannot be deleted, the operation will not be executed and a specific error code will be returned.

Default interface request rate limit: 20 times per second.

Interface update time: 2025-03-29 00:42:26.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DeleteSnapshots
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
SnapshotIds	Yes	No	Array of String	ID list of snapshots to be deleted. It can be queried through the DescribeSnapshots API. Example: ["snap-gybrif0z"]
DeleteBindImages	No	No	Bool	It indicates whether to delete the image associated with the snapshot at the same time. The value false indicates deleting the image bound to the snapshot. In this case, if the snapshot has bound images, the deletion will fail. The value true indicates deleting the image bound to the snapshot at the same time. The default value is false. Example: true

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
UnsupportedOperation.SnapshotHasBindedImage	The snapshot is already associated with an image.
InvalidSnapshot.NotSupported	The snapshot does not support this operation.
InvalidSnapshotId.NotFound	The input `SnapshotId` does not exist.
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.
ResourceUnavailable.SnapshotCreating	The snapshot is being created.

DescribeAutoSnapshotPolicies

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API (DescribeAutoSnapshotPolicies) is used to query regular snapshot policies.

- You can query the details of regular snapshot policies by their ID, name, or status. Different conditions are in an AND relationship. For more details, refer to the Filter .
- If the parameter is empty, a certain number (specified by Limit , default is 20) of regular snapshot policies for the current user is returned.

Default interface request rate limit: 20 times per second.

Interface update time: 2024-08-16 06:42:32.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DescribeAutoSnapshotPolicies
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AutoSnapshotPolicyIds	No	No	Array of String	List of regular snapshot policy IDs to query. Parameters AutoSnapshotPolicyIds and Filters cannot be specified at the same time. Example: ["asp-c9m1berf"]
Filters	No	No	Array of Filter	Filter criteria. Parameters AutoSnapshotPolicyIds and Filters cannot be specified at the same time. • auto-snapshot-policy-id - Array of String - Required: No - (Filter criteria) Filter by regular snapshot policy ID. The regular snapshot policy ID is in the format of asp-11112222 . • auto-snapshot-policy-state - Array of String - Required: No - (Filter criteria) Filter by the status of the regular snapshot policy. The regular snapshot policy ID is in the format of asp-11112222 . (NORMAL: Normal
Limit	No	No	Uint64	Number of returned values. The default value is 20, and the maximum value is 100. For more information about Limit , refer to the relevant sections in API [Overview]. Example: 20
Offset	No	No	Uint64	Offset. The default value is 0. For more information about Offset , refer to the relevant sections in the API [Overview]. Example: 0
Order	No	No	String	Output order of the regular snapshot list. Value range: • ASC: ascending • DESC: descending. Example: ASC
OrderField	No	No	String	Field based on which the regular snapshot list is sorted. Value range: • CREATETIME: Sort by the creation time of the regular snapshot Snapshots are sorted by the creation time by default. Example: CREATETIME
TryCreateDefaultAutoSnapshotPolicy	No	No	Bool	When this parameter is input, the system will try to create a default regular snapshot policy. If the user already has a regular snapshot policy, no policy is created. Example: true
InnerSearch	No	No	String	Fuzzy Search Parameter Example: asp-asgag224

3. Output Parameters

Parameter name	Type	Description
TotalCount	Uint64	Number of valid regular snapshot policies. Example: 0
AutoSnapshotPolicySet	Array of AutoSnapshotPolicy	List of regular snapshot policies. Example:View
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
InvalidFilter	The specified filter is not supported.
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.

DescribeCopyAutoSnapshotPolicy

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API is used to query the details of cross-account automatic snapshot policy replication.

Default interface request rate limit: 20 times per second.

Interface update time: 2025-08-28 17:44:30.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DescribeCopyAutoSnapshotPolicy
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
OrderField	No	No	String	Field based on which the regular snapshot list is sorted. CREATETIME: Sort by the creation time of the regular snapshot. Snapshots are sorted the by creation time by default. Example: CREATE_TIME
CopyAction	Yes	No	String	Automatic snapshot policy replication action. REMOTE: Backup to another account; LOCAL: Backup to my account Example: REMOTE
AutoSnapshotPolicyIds	No	No	Array of String	List of regular snapshot policy IDs to query. Parameters AutoSnapshotPolicyIds and Filters cannot be specified at the same time. Example: ["asp-ljylgbyo"]
Limit	No	No	Uint64	Number of returned values. The default value is 20, and the maximum value is 100. Example: 20
Filters	No	No	Array of Filter	Filter criteria. Parameters AutoSnapshotPolicyIds and Filters cannot be specified at the same time. • auto-snapshot-policy-id - Array of String - Required: No - (Filter criteria) Filter by regular snapshot policy ID. The regular snapshot policy ID is in the format of asp-11112222 . • auto-snapshot-policy-state - Array of String - Required: No - (Filter criteria) Filter by the status of the regular snapshot policy. The regular snapshot policy ID is in the format of asp-11112222 . (NORMAL: Normal
Offset	No	No	Uint64	Offset. The default value is 0. Example: 0
Order	No	No	String	Output order of the regular snapshot list. ASC indicates ascending and DESC indicates descending. Example: ASC

3. Output Parameters

Parameter name	Type	Description
TotalCount	Uint64	Number of Active Regular Snapshot Policies Example: 10
CopyAutoSnapshotPolicySet	Array of AutoSnapshotPolicy	List of Regular Snapshot Policies for Cross-account Replication Example:View
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
InvalidFilter	The specified filter is not supported.
ResourceBusy	The resource is busy. Try again later.
InvalidParameter	Parameter error

DescribeDiskAssociatedAutoSnapshotPolicy

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API (DescribeDiskAssociatedAutoSnapshotPolicy) is used to query the regular snapshot policy to which the cloud disk is bound.

Default interface request rate limit: 20 times per second.

Interface update time: 2018-04-19 09:44:23.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DescribeDiskAssociatedAutoSnapshotPolicy
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
DiskId	Yes	No	String	ID of the cloud disk to be queried. Example: disk-pe1vkgim

3. Output Parameters

Parameter name	Type	Description
TotalCount	Uint64	Number of regular snapshots bound to cloud disks. Example: 0
AutoSnapshotPolicySet	Array of AutoSnapshotPolicy	List of regular snapshots to which cloud disks are bound. Example:View

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.
MissingParameter	Missing parameter. The request does not include the mandatory parameter.

DescribeSnapshotGroups

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API is used to query the snapshot group list, returning snapshot groups that are in the states of Normal, Creating, or Rolling Back.

Default interface request rate limit: 20 times per second.

Interface update time: 2025-08-28 18:03:43.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DescribeSnapshotGroups
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
Filters	No	No	Array of Filter	Filter criteria. • snapshot-group-id - Array of String - Required: No - (Filter criteria) Filter by snapshot group ID • snapshot-group-state - Array of String - Required: No - (Filter criteria) Filter by snapshot group state. (NORMAL: Normal
Offset	No	No	Uint64	Offset. The default value is 0. Example: 0
Limit	No	No	Uint64	Number of returned values. The default value is 20, and the maximum value is 100. Example: 20
Order	No	No	String	Results sorting order. DESC indicates descending, and ASC indicates ascending. Example: ASC
InnerSearch	No	No	String	Fuzzy search parameter. Example: snap-j1ytqcd7

3. Output Parameters

Parameter name	Type	Description
TotalCount	Uint64	Total Number of Qualified Items Example: 100

Parameter name	Type	Description
SnapshotGroupSet	Array of SnapshotGroup	Details of Snapshot Group List Example:View
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameter	Parameter error

DescribeSnapshotGroupsDeniedActions

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API (DescribeSnapshotGroupsDeniedActions) is used to query the operation mask of the snapshot group.

- Based on the snapshot group IDs provided by the input parameter SnapshotGroupIds, the operation mask of the corresponding snapshot group is returned.

Default interface request rate limit: 20 times per second.

Interface update time: 2025-08-28 18:09:44.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DescribeSnapshotGroupsDeniedActions
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
SnapshotGroupIds	Yes	No	Array of String	Snapshot Group ID List Example: ["csnap-csoplz7u"]

3. Output Parameters

Parameter name	Type	Description
SnapshotGroupDeniedActionSet	Array of SnapshotGroupDeniedAction	Operation Mask Set for Querying Snapshot Groups Example:View
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
InvalidParameter	Parameter error

DescribeSnapshotSharePermission

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API (DescribeSnapshotSharePermission) is used to query the sharing information of snapshots.

Default interface request rate limit: 20 times per second.

Interface update time: 2021-04-30 07:27:23.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DescribeSnapshotSharePermission
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
SnapshotId	Yes	No	String	ID of the snapshot to query, which can be accessed through DescribeSnapshots. Example: snap-j1ytqcd7

3. Output Parameters

Parameter name	Type	Description
SharePermissionSet	Array of SharePermission	Collection of Snapshot Sharing Information Example:View
TotalCount	Uint64	Total Example: 0
MaxShareAccount	Uint64	Maximum Shareable Quantity Example: 50
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
MissingParameter	Missing parameter. The request does not include the mandatory parameter.

DescribeSnapshots

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API (DescribeSnapshots) is used to query details of snapshots.

- Filter can be performed by snapshot ID, and ID and type of the cloud disk that created the snapshot. These conditions are in an AND relationship. For detailed filter information, refer to the Filter .
- If parameters are empty, a certain number (specified by Limit , default is 20) of the snapshot list for the current user is returned.

Default interface request rate limit: 20 times per second.

Interface update time: 2024-08-14 08:44:35.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DescribeSnapshots
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
SnapshotIds	No	No	Array of String	List of snapshot IDs to query. Parameters SnapshotIds and Filters cannot be specified at the same time. Example: snap-j1ytqcd7
Filters	No	No	Array of Filter	Filter criteria. The parameters SnapshotIds and Filters cannot be specified at the same time. • snapshot-id - Array of String - Required: No - (Filter criteria) Filter by snapshot ID in the format like snap-11112222 . • snapshot-name - Array of String - Required: No - (Filter criteria) Filter by snapshot name. • snapshot-state - Array of String - Required: No - (Filter criteria) Filter by snapshot status. (NORMAL: normal
Offset	No	No	Uint64	Offset. The default value is 0. Example: 0
Limit	No	No	Uint64	Number of returned values. The default value is 20, and the maximum value is 100. Example: 20
Order	No	No	String	Output cloud disk list sorting order. Value range: • ASC: ascending • DESC: descending. Example: DESC
OrderField	No	No	String	Field based on which the snapshot list is sorted. Value range: • CREATE_TIME: Sort by the creation time of the snapshot Snapshots are sorted by creation time by default. Example: CREATE_TIME
InnerSearch	No	No	String	Fuzzy search parameter. Example: snap-j1ytqcd7

3. Output Parameters

Parameter name	Type	Description
TotalCount	Uint64	Number of snapshots. Example: 20
SnapshotSet	Array of Snapshot	List of snapshots details. Example:View
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
InvalidFilter	The specified filter is not supported.

DescribeSnapshotsDeniedActions

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API (DescribeSnapshotsDeniedActions) is used to query the operation masks of snapshots.

- Based on the snapshot ID provided by the input parameter SnapshotIds, the operation mask of the corresponding snapshot is returned.

Default interface request rate limit: 20 times per second.

Interface update time: 2022-01-07 11:49:00.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DescribeSnapshotsDeniedActions
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
SnapshotIds	Yes	No	Array of String	Snapshot ID List Example: snap-j1ytqcd7

3. Output Parameters

Parameter name	Type	Description
SnapshotDeniedActionSet	Array of SnapshotDeniedAction	Operation Mask Set for Querying Snapshots Example:View
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
InvalidParameter	Parameter error

GetSnapOverview

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API is used to obtain snapshot overview information.

Default interface request rate limit: 20 times per second.

Interface update time: 2025-08-28 18:33:11.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: GetSnapOverview
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface

3. Output Parameters

Parameter name	Type	Description
TotalSize	Float	Total Size of User Snapshots Example: 4.28
RealTradeSize	Float	Total Size of User Snapshots (for Billing) Example: 4.28

Parameter name	Type	Description
FreeQuota	Float	Snapshot Free Quota Example: 0.0
TotalNums	Int64	Total Number of Snapshots Example: 2
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.

ModifyAutoSnapshotPolicyAttribute

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API (ModifyAutoSnapshotPolicyAttribute) is used to modify the attributes of regular snapshot policies.

- You can call this API to modify the execution policy, name, and whether to activate, among other attributes of the regular snapshot policy.
- When modifying the retention days, ensure there is no conflict with the permanent retention attribute; otherwise, the whole operation fails and a specific error code is returned.

Default interface request rate limit: 20 times per second.

Interface update time: 2025-03-17 19:43:46.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: ModifyAutoSnapshotPolicyAttribute
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AutoSnapshotPolicyId	Yes	No	String	Regular snapshot policy ID. Example: asp-nqu08k2l
Policy	No	No	Array of Policy	Execution policy of regular snapshots. Example:View
AutoSnapshotPolicyName	No	No	String	Name of the regular snapshot policy to be created. If this parameter is not input, it defaults to "unnamed". The maximum length cannot exceed 60 bytes. Example: data_disk_auto_snapshot
IsActivated	No	No	Bool	It indicates whether the regular snapshot policy is activated. FALSE indicates not activated, and TRUE indicates activated. The default value is TRUE. Example: 1
IsPermanent	No	No	Bool	It indicates whether the snapshot created by the regular snapshot policy is permanently retained. FALSE indicates not permanently retained, and TRUE indicates permanently retained. The default value is FALSE. Example: TRUE
RetentionDays	No	No	Uint64	Retention days for snapshots created by the regular snapshot policy. This parameter cannot conflict with the IsPermanent parameter. That is, if the regular snapshot policy is set to retain permanently, RetentionDays must be set to 0. Example: 6
RetentionMonths	No	No	Uint64	Number of months the snapshot created by this regular snapshot policy can be retained. This parameter cannot conflict with the IsPermanent/RetentionDays parameters. Example: 6
RetentionAmount	No	No	Uint64	Maximum number of snapshots that can be retained through this regular snapshot policy. If this number is exceeded, the earliest created snapshot will be automatically deleted. This parameter cannot conflict with the IsPermanent parameter. Example: 6
AdvancedRetentionPolicy	No	No	AdvancedRetentionPolicy	Advanced retention policy for regular snapshot. This parameter cannot conflict with the IsPermanent parameter. Example:View
IsCrossRegion	No	No	Bool	Whether to enable cross-region snapshot replication. The default value is 0, indicating that it is disabled. Example: true
UnsetAdvancedRetentionPolicy	No	No	Bool	Whether to delete the advanced snapshot policy. Example: true

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
InvalidSnapshot.NotSupported	The snapshot does not support this operation.
InvalidAutoSnapshotPolicyId.NotFound	The input AutoSnapshotPolicyId does not exist.

ModifySnapshotAttribute

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API is used to modify the attributes of a specified snapshot.

*Currently, only the snapshot name can be modified, and a non-permanent snapshot can be changed to a permanent one.
*The snapshot name is used only for snapshot management by users. The cloud platform does not use this name for submitting tickets or performing snapshot management operations.

Default interface request rate limit: 20 times per second.

Interface update time: 2024-07-16 07:44:51.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: ModifySnapshotAttribute
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
SnapshotId	Yes	No	String	Snapshot ID. It can be queried through the DescribeSnapshots API. Example: snap-gybrif0z
SnapshotName	No	No	String	New snapshot name. The maximum length is 60 characters. Example: snap_201711301143
IsPermanent	No	No	Bool	Retention period of the snapshot. FALSE indicates non-permanent, and TRUE indicates permanent. It only supports changing a non-permanent snapshot to a permanent one. Example: true
ProjectId	No	No	Int64	New snapshot project ID. Example: 1002189

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.
InvalidSnapshotId.NotFound	The input SnapshotId does not exist.
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
InvalidSnapshot.NotSupported	The snapshot does not support this operation.

ModifySnapshotsSharePermission

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API (ModifySnapshotsSharePermission) is used to modify snapshot sharing information.
After a snapshot is shared, the shared account can create cloud disk with that snapshot.

- Each snapshot can be shared with up to 50 accounts.
- Shared snapshots cannot have their name or description changed and can only be used for creating cloud disks.
- It supports sharing only within the same region of the recipient account.
- It only supports sharing data disk snapshots.

Default interface request rate limit: 20 times per second.

Interface update time: 2024-07-16 07:34:08.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: ModifySnapshotsSharePermission
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AccountIds	Yes	No	Array of String	ID list of accounts that receive a shared snapshot. Example: ["100001122000"]
Permission	Yes	No	String	Operation, including SHARE and CANCEL. SHARE indicates sharing, and CANCEL indicates canceling share. Example: SHARE
SnapshotIds	Yes	No	Array of String	Snapshot ID. It can be queried through the DescribeSnapshots API. Example: ["snap-4zgc8cnk"]

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.

UnbindAutoSnapshotPolicy

1. Interface Description

Interface request domain name: cbs.api3.convergecloud.com.

This API (UnbindAutoSnapshotPolicy) is used to unbind the regular snapshot policy to which the cloud disk is bound. * It supports batch operation. Multiple cloud disks can be unbound from the same regular snapshot policy at a time. * If the input cloud disk is not bound to the current regular snapshot policy, the API will be automatically skipped and only the cloud disk bound to the current regular snapshot policy is unbound.

Default interface request rate limit: 20 times per second.

Interface update time: 2025-03-29 00:44:39.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: UnbindAutoSnapshotPolicy
Version	Yes	No	String	Public parameter, The value of this interface: 2017-03-12
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
DiskIds	No	No	Array of String	ID list of cloud disks from which the regular snapshot policy is unbound. Example: ["disk-6rag8jef"]
AutoSnapshotPolicyId	Yes	No	String	Regular snapshot policy ID to unbind. Example: asp-mrsrn243
InstanceIds	No	No	Array of String	List of instance IDs to unbind from the regular snapshot policy. Example: ["ins-locaa0iu"]

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
InvalidAutoSnapshotPolicyId.NotFound	The input AutoSnapshotPolicyId does not exist.

Data structure

SharePermission

Snapshot Sharing Information Set

Referenced by the following interfaces: DescribeSnapshotSharePermission

name	required	NULL allowed	type	description
CreatedTime	yes	no	Datetime	Time of Snapshot Sharing Example: 2023-02-26 01:21:03
AccountId	yes	no	String	Shared Account ID Example: 100022975249

Tag

Tag.

Referenced by the following interfaces: CreateDisks, CreateSnapshot, DescribeAutoSnapshotPolicies, DescribeCopyAutoSnapshotPolicy, DescribeDiskAssociatedAutoSnapshotPolicy, DescribeDiskAssociatedSnapshots, DescribeDisks, DescribeSnapshotGroups, DescribeSnapshots, DescribeSnapshotsInternal, SwitchParameterCreateDisks

name	required	NULL allowed	type	description
Key	yes	no	String	Tag key. Example: TagKey
Value	yes	no	String	Tag value. Example: TagValue

AdvancedRetentionPolicy

Advanced Snapshot Retention Policy

Referenced by the following interfaces: CreateAutoSnapshotPolicy, DescribeAutoSnapshotPolicies, DescribeCopyAutoSnapshotPolicy, DescribeDiskAssociatedAutoSnapshotPolicy, ModifyAutoSnapshotPolicyAttribute

name	required	NULL allowed	type	description
Days	no	yes	Uint64	The latest snapshot of each day within Days is retained. Example: 6
Weeks	no	yes	Uint64	The latest snapshot of each week within Weeks is retained. Example: 8
Months	no	yes	Uint64	The latest snapshot of each month within Months is retained. Example: 3
Years	no	yes	Uint64	The latest snapshot of each year within Years is retained. Example: 1

Policy

It describes the execution strategy of regular snapshots.

Referenced by the following interfaces: CreateAutoSnapshotPolicy, DescribeAutoSnapshotPolicies, DescribeCopyAutoSnapshotPolicy, DescribeDiskAssociatedAutoSnapshotPolicy, ModifyAutoSnapshotPolicyAttribute

name	required	NULL allowed	type	description
DayOfWeek	no	no	Array of Uint64	Selects a day from Monday to Sunday to create snapshots. The value range is from 0 to 6. 0 indicates triggered on Monday, and so on. Example: [0,1,2]
Hour	yes	no	Array of Uint64	It specifies the trigger time for the regular snapshot policy in the unit of hours. The value range is from 0 to 23. A total of 24 options from 00:00 to 23:00 are available, with 1 indicating 01:00, and so on. Example: [11,12,13]
IntervalDays	no	no	Uint64	Specifies the interval days for creating automatic snapshots. Value range: [1, 365]. For example, if it is set to 5, creating a automatic snapshot will be triggered every 5 days. Note: When you choose daily backups, theoretically, the time for the first backup is on the day when the backup policy is created. If the time of policy creation on that day is later than the specified backup time, the first backup occurs on the second backup cycle. Example: 50
DayOfMonth	no	no	Array of Uint64	Specified days in each month to trigger snapshot creation. Example: [0]

ResizeGoodsDetail

Product renewal details. It is only used for the SwitchParameterResizeDisks API.

Referenced by the following interfaces: SwitchParameterModifyDiskAttributes

name	required	NULL allowed	type	description
Pid	yes	no	Uint64	Product ID. Example: 8000
ResourceId	yes	no	String	Cloud disk ID. Example: disk-7q2ly7gp
CurDeadline	yes	no	String	Current expiration time of the cloud disk. Example: 2023-03-21 12:00:00
ProductInfo	yes	no	Array of ProductInfo	Product information description. Example:View
NewConfig	yes	no	DiskResizeConfig	Cloud disk configuration after scaling. Example:View
OldConfig	yes	no	DiskResizeConfig	Configuration before cloud disk scaling. Example:View

AttachDetail

It describes the number of mounted and mountable data disks for an instance.

Referenced by the following interfaces: DescribeInstancesDiskNum

name	required	NULL allowed	type	description
InstanceId	yes	yes	String	Instance ID. Example: ins-jw0vit58
AttachedDiskCount	yes	yes	Uint64	Number of data disks mounted to the instance. Example: 1
MaxAttachCount	yes	yes	Uint64	Maximum number of data disks to which an instance can mount. Example: 20

SnapshotDeniedAction

Operation Mask of Single Snapshot

Referenced by the following interfaces: DescribeSnapshotsDeniedActions

name	required	NULL allowed	type	description
SnapshotId	yes	no	String	Based on Snapshot ID Example: snap-hdbc13ad
DeniedActions	yes	no	Array of DeniedAction	Specific Operation Mask Information for Snapshot Example:View

AutoSnapshotPolicy

It describes the details of the regular snapshot policy.

Referenced by the following interfaces: DescribeAutoSnapshotPolicies, DescribeCopyAutoSnapshotPolicy, DescribeDiskAssociatedAutoSnapshotPolicy

name	required	NULL allowed	type	description
AutoSnapshotPolicyId	no	no	String	Regular snapshot policy ID. Example: asp-nz3ybdwj
AutoSnapshotPolicyName	no	no	String	Name of the regular snapshot policy. Example: cloud_demo
AutoSnapshotPolicyState	no	no	String	Status of the regular snapshot policy. Value range: • NORMAL: normal • ISOLATED: isolated. Example: NORMAL
IsActivated	no	no	Bool	It indicates whether the regular snapshot policy is activated. Example: true
IsPermanent	no	no	Bool	It indicates whether snapshots created by this regular snapshot policy are permanently retained. Example: false
RetentionDays	no	no	Uint64	It indicates the retention days for snapshots created by the regular snapshot policy. Example: 1
CreateTime	no	no	Datetime	Creation time of the regular snapshot policy. Example: 2022-01-11 10:45:51
NextTriggerTime	no	no	Datetime	Next trigger time of regular snapshots. Example: 2022-05-31 11:00:00
Policy	no	no	Array of Policy	Execution policy of regular snapshots. Example:View
DiskIdSet	no	no	Array of String	It lists IDs of cloud disks that have been bound to the current regular snapshot policy. Example: ["disk-pe1vkgim"]
IsCopyToRemote	yes	yes	Uint64	Cross-account Snapshot Replication or Not Example: 0
CopyFromAccountUin	yes	yes	String	Source Account UIN Example: 100004603241
CopyToAccountUin	yes	yes	String	Destination Account UIN Example: 100004603246
InstanceIdSet	no	no	Array of String	List of instances bound to the current regular snapshot policy. Example: ["ins-og3j7lur"]
RetentionMonths	no	no	Uint64	Maximum Retention Months for Snapshots Created by This Regular Snapshot Policy Example: 1
RetentionAmount	no	no	Uint64	Maximum Number of Retained Snapshots Created by This Regular Snapshot Policy Example: 5
AdvancedRetentionPolicy	no	yes	AdvancedRetentionPolicy	Advanced Snapshot Retention Policy Example:View
IsCrossRegion	no	yes	Bool	Whether cross-region snapshot replication is enabled in the regular snapshot policy. Example: false
Tags	no	yes	Array of Tag	Tag. Example:View
CrossRegionRetentionDays	no	yes	Int64	Retention days for cross-region snapshots created by using this scheduled snapshot policy. The default value is 7 days. This parameter cannot conflict with the CrossRegionIsPermanent parameter. Example: 7
CrossRegionIsPermanent	no	yes	Bool	Whether cross-region snapshots created by using this scheduled snapshot policy are retained permanently. The default value is 0, indicating the snapshots are not retained permanently. Example: false

ProductInfo

It describes the product information output by the order parameter conversion API.

Referenced by the following interfaces: SwitchParameterCreateDisks, SwitchParameterModifyDiskAttributes, SwitchParameterRenewDisks, SwitchParameterResizeDisk

name	required	NULL allowed	type	description
Name	yes	no	String	Attribute name. Example: CBS instance alias

name	required	NULL allowed	type	description
Value	yes	no	String	Attribute value. Example: Unnamed

DiskOrder

Order parameters for request billing.

Referenced by the following interfaces: SwitchParameterCreateDisks, SwitchParameterRenewDisks, SwitchParameterResizeDisk

name	required	NULL allowed	type	description
GoodsCategoryId	no	no	Uint64	Product category. Example: 100000
GoodsNum	no	no	Uint64	Number of products. Example: 10
ProjectId	no	no	Uint64	Project ID. Example: 1002189
GoodsDetail	no	no	GoodsDetail	Product details. Example:View
RegionId	no	no	Uint64	Region ID. Example: 10
ZoneId	no	no	Uint64	Availability zone ID. Example: 100001
PayMode	no	no	Uint64	Payment mode. Example: 2
Type	no	no	String	Cloud disk type. Example: CLOUD_TSSD

SnapOverview

User snapshot information overview. It includes the total number and total capacity of created snapshots.

Referenced by the following interfaces: DescribeUserDiskResources

name	required	NULL allowed	type	description
SnapshotCapacityUsedTotal	yes	no	Float	Total snapshot capacity created by the user. Example: 0.0
SnapshotNumberUsedTotal	yes	no	Int64	Total number of snapshots created by the user. Example: 24

DiskDeniedAction

Operation Mask of Single Cloud Disk

Referenced by the following interfaces: DescribeDisksDeniedActions

name	required	NULL allowed	type	description
DiskId	yes	no	String	It indicates the cloud disk ID, such as "disk-7rf24jvb". Example: disk-pe1vkgim
DeniedActions	yes	no	Array of DeniedAction	Specific Cloud Disk Operation Mask Information Example:View

Filter

It describes the key-value pair filter, used for conditional filter query.

Referenced by the following interfaces: DescribeAutoSnapshotPolicies, DescribeCopyAutoSnapshotPolicy, DescribeDisks, DescribeSnapshotGroups, DescribeSnapshots, DescribeSnapshotsInternal

name	required	NULL allowed	type	description
Name	yes	no	String	Filter key name. Example: snapshot-id
Values	yes	no	Array of String	One or more filter values. Example: snap-9asdgcsq

SnapshotCopyResult

It describes the results of snapshot cross-region replication.

Referenced by the following interfaces: CopySnapshotCrossRegions

name	required	NULL allowed	type	description
DestinationRegion	yes	no	String	Target region for cross-region replication. Example: "ap-region1"
SnapshotId	yes	no	String	New snapshot ID copied to the target region. Example: snap-hdbc13ad
Code	no	no	String	Returned Status Code Example: Success
Message	no	no	String	Return Message Example: Success

ApplyDisk

This parameter is used as the input for the snapshot group rollback API, indicating the cloud disk to be rolled back, and the snapshot list.

Referenced by the following interfaces: ApplySnapshotGroup

name	required	NULL allowed	type	description
SnapshotId	yes	no	String	Snapshot ID Example: snap-czneir4gl
DiskId	yes	no	String	Original Cloud Disk ID for Snapshot Example: disk-1ly0wu8c

BlockStorage

It describes the details of the block storage.

Referenced by the following interfaces: DescribeBlockStorages

name	required	NULL allowed	type	description
DiskId	yes	no	String	Block storage ID. Example: disk-pe1vkgim
DiskName	yes	no	String	Block storage name. Example: CBS instance name-zzl
DiskSize	yes	no	Uint64	Block storage size (in GB). Example: 50
DiskType	yes	no	String	Block Storage Medium Type. Value range: • CLOUD_BASIC • CLOUD_PREMIUM • CLOUD_SSD • CLOUD_ENHANCEDSSD • LOCAL_BASIC • LOCAL_SSD Example: CLOUD_TSSD
DiskUsage	yes	no	String	Block storage type. Value range: • SYSTEM_DISK: system disk • DATA_DISK: data disk. Example: DATA_DISK
InstanceId	yes	yes	String	ID of the instance to which block storage is mounted. Example: ins-9oxguxa
InstanceName	yes	no	String	Name of the instance to which block storage is mounted. Example: thins-kailwxe

GoodsDetail

Product details.

Referenced by the following interfaces: SwitchParameterCreateDisks, SwitchParameterRenewDisks, SwitchParameterResizeDisk

name	required	NULL allowed	type	description
TimeSpan	no	no	Uint64	Duration for purchasing or renewing a cloud disk. Example: 1
Pid	no	no	Uint64	ProductID. Example: 99903
DiskType	no	no	String	It indicates the cloud disk type. Options are system disk or data disk. Example: CLOUD_TSSD
ProductInfo	no	no	Array of ProductInfo	Product information description. Example:View
MediumType	no	no	String	Cloud disk media type. Example: CLOUD_HSSD
TimeUnit	no	no	String	Unit of duration "TimeSpan". Example: m
CbsSize	no	no	Uint64	Cloud disk size. Example: 50

Filters

It describes the key-value pair filter, used for conditional filter query.

Referenced by the following interfaces: DescribeDiskStoragePoolGroups

name	required	NULL allowed	type	description
Name	yes	no	String	Filter key name. Example: snapshot-group-id
Values	yes	no	Array of String	One or more filter values. Example: csnap-c44200bd

ResourcesDetail

It describes the details of the user resource.

Referenced by the following interfaces: DescribeUserDiskResources

name	required	NULL allowed	type	description
Region	yes	no	String	Resource region. Example: region1
SnapshotNumberTotal	yes	no	Uint64	Snapshot user quantity quota. Example: 99
SnapshotNumberUsed	yes	no	Uint64	Number of used snapshots. Example: 29

name	required	NULL allowed	type	description
SnapshotCapacityTotal	yes	no	Uint64	Snapshot capacity quota. Example: 13
SnapshotCapacityUsed	yes	no	Uint64	Used snapshot capacity (in GB). Example: 425
DiskNumber	yes	no	Int64	Number of cloud disks. Example: 14
NewDiskFlag	yes	no	Bool	It indicates whether the current region includes newly purchased cloud disks. Example: true
SnapshotSupportCrossCopy	yes	no	Bool	It indicates whether the current region supports cross-region replication Example: false
DiskNumberIsolated	yes	no	Uint64	Number of cloud disks in the recycle bin. Example: 0
ExpiredNumber	yes	no	Int64	Number of expired cloud disks. Example: 0
DiskNumberExpireIn7days	yes	no	Int64	Number of cloud disks that will expire within 7 days. Example: 0
DiskNumberUnderDataProtection	yes	no	Int64	Number of protected cloud disks. Example: 0

CopyFromBackupDetail

Snapshot replicated from a backup, obtaining the information of the corresponding backup.

Referenced by the following interfaces: DescribeDiskAssociatedSnapshots, DescribeSnapshots, DescribeSnapshots3Internal

name	required	NULL allowed	type	description
BackupGroupId	yes	yes	String	Backup group ID. Example: cbackup-ccy6hfmj
CreateTime	yes	yes	Datetime	Backup creation time. Example: 2024-08-09 15:03:44
BackupDiskId	yes	yes	String	Backup disk ID. Example: disk-pe1vkgim
BackupId	yes	yes	String	Backup ID, which is the unique identifier of a backup. Example: backup-c9agada4

PoolCustomAttributes

It describes the custom attributes of the user resource pool.

Referenced by the following interfaces: DescribeDiskStoragePoolGroups

name	required	NULL allowed	type	description
CpuArchitecture	no	yes	Array of String	CPU architecture. Example: aarch64
CpuModel	no	yes	Array of String	CPU model. Example: Virtual

DiskChargePrepaid

It describes the instance billing mode.

Referenced by the following interfaces: CreateDisks, InquiryPriceCreateDisks, InquiryPriceCreateSnapshots, InquiryPriceRenewDisks, RenewDisk, SwitchParameterCreateDisks, SwitchParameterRenewDisks

name	required	NULL allowed	type	description
Period	yes	no	Uint64	Duration of cloud disk purchase. The default unit is month, with options 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 24, and 36. Example: 1
RenewFlag	no	no	String	Auto-renewal flag. Value range: • NOTIFY_AND_AUTO_RENEW: Notify that the account is about to expire and automatically renew the account. • NOTIFY_AND_MANUAL_RENEW: Notify that the account is about to expire but do not automatically renew the account. • DISABLE_NOTIFY_AND_MANUAL_RENEW: Neither notify that the account is about to expire nor automatically renew the account. Default value: NOTIFY_AND_MANUAL_RENEW. Notify that the account is about to expire but do not automatically renew the account. Example: NOTIFY_AND_AUTO_RENEW
CurInstanceDeadline	no	no	Datetime	When it is necessary to align the expiration time of the cloud disk with the mounted child machine, this parameter can be input. This parameter indicates the current expiration time of the child machine. At this time, if Period is input, it indicates the duration that the child machine needs to renew, and the cloud disk will automatically renew according to the expiration time after the child machine is renewed. Example: 2018-03-30 20:1:03

DedicatedClusterDiskStatistic

Cloud disk statistical information of CDC exclusive clusters.

Referenced by the following interfaces: DescribeDedicatedClusterDiskStatistics

name	required	NULL allowed	type	description
DiskType	yes	no	String	Hard disk media type. Example: CLOUD_SSD
TotalDiskSize	yes	no	Uint64	Total cloud disk capacity. Example: 1000
UsedDiskSize	yes	no	Uint64	Used cloud disk capacity. Example: 100
AvailableDiskSize	yes	no	Uint64	Available cloud disk capacity. Example: 900

DiskStoragePoolGroupInfo

Storage Pool Group of Cloud Disk

Referenced by the following interfaces: DescribeDiskStoragePoolGroups

name	required	NULL allowed	type	description
DiskStoragePoolGroupName	yes	no	String	Name of Storage Resource Group Example: ec
PoolBoundType	yes	no	String	Binding type Example: OPTIONAL
DiskStoragePoolGroup	yes	no	String	Storage Resource Pool Group Example: pool-ec
Zone	yes	no	String	Availability Zone Example: az1
DiskType	yes	no	String	Cloud disk type Example: CLOUD_TSSD
Available	yes	no	Bool	Sold Out or Not Example: true
Message	yes	no	String	Sell-out Information Example: ""
PoolCustomAttributes	no	yes	PoolCustomAttributes	It describes the CPU information of the resource pool. Example:View

Image

Image Associated with Snapshot

Referenced by the following interfaces: DescribeDiskAssociatedSnapshots, DescribeSnapshotGroups, DescribeSnapshots, DescribeSnapshots3Internal

name	required	NULL allowed	type	description
ImageId	yes	yes	String	Image instance ID. Example: img-m0w653yh
ImageName	yes	yes	String	Image name. Example: image-1

Placement

It describes the abstract location of the instance, including its availability zone and the project it belongs to.

Referenced by the following interfaces: CreateDisks, DescribeDiskAssociatedSnapshots, DescribeDisks, DescribeSnapshots, DescribeSnapshots3Internal, InquiryPriceCreateSnapshots, SwitchParameterCreateDisks

name	required	NULL allowed	type	description
Zone	yes	no	String	ID of the availability zone to which the instance belongs. This parameter can also be obtained from the Zone field returned by calling the DescribeZones API. Example: ap-region1-2
ProjectId	no	no	Uint64	ID of the project to which the instance belongs. This parameter can be obtained from the projectId field returned by calling the DescribeProjects API. If this parameter is not specified, the default project is used. Example: 0
CdcId	no	yes	String	ID of the exclusive cluster to which the instance belongs. As an input parameter, it indicates operating the resources of the specified CdcId exclusive cluster, which can be empty. As an output parameter, it indicates the ID of the exclusive cluster to which the resource belongs, which can be empty. Example: cdc-98asazed
CageId	no	yes	String	Cage ID. As an input parameter, it indicates operating the resources of a specified cage ID, which can be empty. As an output parameter, it indicates the cage ID to which the resource belongs, which can be empty. Example: cage-axqr92da
CdcName	no	yes	String	Name of the exclusive cluster. As an input parameter, ignore this parameter. As an output parameter, it indicates the name of the exclusive cluster the cloud disk belongs to, which can be empty. Example: TesetCdcName
ProjectName	no	yes	String	Project Name Example: Default Project
DedicatedClusterId	no	yes	String	ID of the CVM exclusive cluster. As an input parameter, it indicates operating on the resources of the specified exclusive cluster. As an output parameter, it indicates the ID of the exclusive cluster to which the resource belongs, which can be empty. Example: cluster-ay3cdptf
ClusterGroupId	no	yes	String	Cluster group ID. Example: ""

Disk

It describes the details of the cloud disk.

Referenced by the following interfaces: DescribeDisks

name	required	NULL allowed	type	description
DiskId	no	no	String	Cloud disk ID. Example: disk-pe1vkgim
DiskUsage	no	no	String	Cloud disk type. Value range: • SYSTEM_DISK: system disk • DATA_DISK: data disk. Example: SYSTEM_DISK
DiskChargeType	no	no	String	Payment MODE. Value range: • PREPAID: prepaid, that is, monthly subscription • POSTPAID_BY_HOUR: postpaid, that is, pay-as-you-go. Example: PREPAID
Portable	no	no	Bool	It indicates whether it is an elastic cloud disk. The value false indicates non-elastic cloud disk, and the value true indicates elastic cloud disk. Example: true
Placement	no	no	Placement	Location of a cloud disk. Example:View

name	required	NULL allowed	type	description
SnapshotAbility	no	no	Bool	It indicates whether the cloud disk has the capability to create a snapshot. Value range: • false: Do not have this capacity • true: Have this capacity. Example: true
DiskName	no	no	String	Cloud disk name. Example: unname
DiskSize	no	no	Uint64	Cloud disk size. Example: 50
DiskState	no	no	String	Cloud disk status. Value range: • UNATTACHED: Unmounted • ATTACHING: Mounting • ATTACHED: Mounted • DETACHING: Unmounting • EXPANDING: Expanding • ROLLBACKING: Rolling back. Example: UNATTACHED
DiskType	no	no	String	It indicates the type of cloud disk media. Value range: • CLOUD_BASIC: Indicates a basic cloud disk. • CLOUD_PREMIUM: Indicates a high-performance cloud disk. • CLOUD_SSD: Indicates an SSD cloud disk. Example: CLOUD_SSD
Attached	no	no	Bool	It indicates whether the cloud disk is mounted to the cloud host. Value range: • false: Not mounted • true: Mounted. Example: false
InstanceId	no	no	String	ID of the cloud host to which the cloud disk is mounted. Example: ins-9oxguxa
CreateTime	no	no	String	Creation time of cloud disks. Example: 2021-01-01 12:00:00
DeadlineTime	no	no	String	Expiration time of cloud disks. Example: 2021-02-01 12:00:00
Rollbacking	no	no	Bool	It indicates whether the cloud disk is in snapshot rollback status. Value range: • false: It is not in snapshot rollback status • true: It is in snapshot rollback status. Example: false
RollbackPercent	no	no	Uint64	Progress of cloud disk snapshot rollback. Example: 0
Encrypt	no	no	Bool	It indicates whether the cloud disk is encrypted. Value range: • false: Indicates an unencrypted disk. • true: Indicates an encrypted disk. Example: false
AutoRenewFlagError	no	no	Bool	The cloud disk has been mounted to the child machine, and both the child machine and the cloud disk are on a monthly subscription basis. • true: The auto renewal flag is set to the child machine but not to the cloud disk. • false: The cloud disk's auto-renewal flag is normal. Example: true
RenewFlag	no	no	String	Auto renewal flag. Value range: • NOTIFY_AND_AUTO_RENEW: Notify that the account is about to expire and automatically renew the account. • NOTIFY_AND_MANUAL_RENEW: Notify that the account is about to expire but do not automatically renew the account. • DISABLE_NOTIFY_AND_MANUAL_RENEW: Neither notify that the account is about to expire nor automatically renew the account. Example: NOTIFY_AND_AUTO_RENEW
DeadlineError	no	no	Bool	This field is valid only when the cloud disk is mounted to the instance, and both the instance and the cloud disk are on a monthly subscription basis. • true: The expiration time of the cloud disk is earlier than that of the instance. • false: The expiration time of the cloud disk is later than that of the instance. Example: true
IsReturnable	no	yes	Bool	It determines if a prepaid cloud disk supports proactive return. • true: It supports proactive return. • false: It does not support proactive return. Example: true
ReturnFailCode	no	yes	Uint64	For prepaid cloud disks that do not support proactive returns, this parameter indicates the specific reason for the lack of support. Value range: • 1: The cloud disk has been returned. • 2: The cloud disk has expired. • 3: The cloud disk does not support return. • 8: The maximum number that can be returned is exceeded. Example: 10
AutoSnapshotPolicyIds	no	yes	Array of String	Regular snapshot ID associated with the cloud disk. This parameter is only returned when the DescribeDisks API is called and ReturnBindAutoSnapshotPolicy is set to TRUE. Example: ["asp-ljyhj6yo"]
Tags	yes	yes	Array of Tag	It indicates the tag bound to the cloud disk. If the cloud disk is not bound to any tag, the value is empty. Example:View
DeleteWithInstance	yes	yes	Bool	It indicates whether the cloud disk is destroyed along with the mounted instance. • true: The cloud disk is destroyed along with the instance. Only hourly postpaid cloud disks are supported. • false: The cloud disk is not destroyed when the instance is destroyed. Example: false
DifferDaysOfDeadline	yes	yes	Int64	Days until disk expiration (only valid for prepaid disks). Example: 20
Migrating	yes	yes	Bool	It indicates whether the cloud disk currently is undergoing a type change. Value range: • false: Indicates the cloud disk is not undergoing a type change. • true: Indicates the cloud disk has initiated a type change and is in the process of migrating. Example: false
MigratePercent	yes	yes	Uint64	It indicates the migration progress of cloud disk type change. The value range is 0 to 100. Example: 0
Shareable	yes	yes	Bool	Is the cloud disk a shared cloud disk. Example: false
InstanceIdList	yes	yes	Array of String	For non-shared cloud disks, this parameter is an empty array. For shared cloud disks, it indicates the CVM instance ID to which the cloud disk is currently mounted. Example: ["ins-9oxguxa"]
SnapshotCount	yes	yes	Int64	Total number of snapshots owned by the cloud disk. Example: 0
SnapshotSize	yes	yes	Uint64	Total snapshot capacity of the cloud disk, in MB. Example: 0
PlatformProjectId	yes	yes	String	Project ID Example: pr-42749a13
BackupDisk	yes	yes	Bool	When a cloud disk is destroyed due to overdue or expiration, the flag indicates whether to use snapshot backup data. The value true indicates creating a snapshot for data backup on destruction. The value false indicates direct destruction without data backup. Example: true
InstanceType	yes	yes	String	Model of Child Machine to Which Disk Mounted Example: CVM
DiskStoragePoolGroup	yes	yes	String	Resource Pool to Which Disk Belongs Example: pool-ec
LatestOperationRequestId	yes	yes	String	It indicates the API RequestId corresponding to the last executed asynchronous task. Example: 597205a3-68c8-4f8a-bd3d-91cf0fa5a46c
ThroughputPerformance	yes	yes	Uint64	Additional cloud disk performance value (in MB/s). Example: 0
LatestOperation	yes	yes	String	Last executed asynchronous task operation. Example: ApplySnapshot
LastAttachInsId	yes	yes	String	ID of Last Cloud Virtual Machine to Which Cloud Disk Mounted Example: ins-9oxguxa
LatestOperationState	yes	yes	String	Result of the last executed asynchronous task. Example: SUCCESS
DeleteSnapshot	yes	yes	Int64	Delete associated non-permanently retained snapshots when destroying the cloud disk. 0 indicates non-permanent snapshots are not deleted with cloud disk destruction, 1 indicates non-permanent snapshots are deleted with cloud disk destruction, and the default value is 0. Whether a snapshot is permanently retained is determined by the IsPermanent field in the snapshot details returned by the DescribeSnapshots API. The value true indicates permanent snapshot, and the value false indicates non-permanent snapshot. Example: 0
ErrorPrompt	yes	yes	String	Error Message Returned to User upon Failure Example: "SUCCESS"
EncryptAlgorithm	yes	yes	String	Encryption algorithm type. SM4: National Standard; AES256: Non-National Standard Example: SM4
DiskBackupQuota	no	yes	Uint64	Quota of cloud disk backup points. It indicates the maximum number of backup points that can be retained for a cloud disk. Example: 0
DiskBackupCount	no	yes	Uint64	Number of cloud disk backup points used. Example: 0
BurstPerformance	no	yes	Bool	Whether performance burst is enabled for the cloud disk. Example: false

DiskFeature

Features supported by a cloud disk product.

Referenced by the following interfaces: DescribeDiskSupportFeatures

name	required	NULL allowed	type	description
TradeType	yes	yes	String	Billing system used by the cloud disk. Value range: • bmppro: billing system • bmp: metering system Example: bmp

DiskOverview

User cloud disk information overview. It includes the total number of cloud disks, the number of expired cloud disks, and the number of cloud disks expiring within 7 days.

Referenced by the following interfaces: DescribeUserDiskResources

name	required	NULL allowed	type	description
DiskNumberTotal	yes	no	Int64	Total number of cloud disks. Example: 10
ExpiredNumberTotal	yes	no	Int64	Total number of expired cloud disks. Example: 5
DiskNumberExpireIn7days	yes	no	Int64	Total Number of Cloud Disks Expiring Within 7 days Example: 2
DiskNumberUnderDataProtection	yes	no	Int64	Number of protected disks. Example: 1

DiskResizeConfig

Cloud disk scaling configuration. It is used for the output of the SwitchParameterResizeDisk API.

Referenced by the following interfaces: SwitchParameterModifyDiskAttributes

name	required	NULL allowed	type	description
Pid	yes	no	Uint64	Product ID. Example: 8000
CbsSize	yes	no	Uint64	Cloud disk size. Example: 50
MediumType	yes	no	String	Cloud disk media type. Example: CbsPremium

Price

It describes the price of the cloud disk.

Referenced by the following interfaces: InquirePriceModifyDiskExtraPerformance, InquiryPriceCreateDisks, InquiryPriceCreateSnapshots, InquiryPriceModifyDiskAttributes, InquiryPriceRenewDisks, InquiryPriceResizeDisk

name	required	NULL allowed	type	description

name	required	NULL allowed	type	description
OriginalPrice	no	yes	Float	Original price of pre-payment (in USD). Example: 1.0
DiscountPrice	no	yes	Float	Prepaid discounted price (in USD). Example: 1.0
UnitPriceDiscount	no	yes	Float	Discounted unit price of postpaid cloud disks (in USD). Example: 1.0
UnitPrice	no	yes	Float	Original unit price of postpaid cloud disks (in USD). Example: 1.0
UnitPriceHigh	no	yes	String	Original Unit Price of High-precision Postpaid Cloud Disk (in USD) Example: "1.00"
OriginalPriceHigh	no	yes	String	Original price of high-precision prepaid cloud disks (in USD). Example: "1.00"
DiscountPriceHigh	no	yes	String	Discounted Price of High-precision Prepaid Cloud Disk (in USD) Example: "1.00"
UnitPriceDiscountHigh	no	yes	String	Discounted Unit Price of High-precision Postpaid Cloud Disk (in USD) Example: "1.00"
ChargeUnit	no	yes	String	It indicates the billing cell of the postpaid cloud disk. Value range: • HOUR: Indicates the billing cell of the postpaid cloud disk is calculated on an hourly basis. Example: "h"
DetailPrices	no	yes	Array of DetailPrice	Details list of billable items. Example:View

DetailPrice

Price of a prepaid or postpaid cloud disk.

Referenced by the following interfaces: InquirePriceModifyDiskExtraPerformance, InquiryPriceCreateDisks, InquiryPriceCreateSnapshots, InquiryPriceModifyDiskAttributes, InquiryPriceRenewDisks, InquiryPriceResizeDisk

name	required	NULL allowed	type	description
PriceTitle	no	yes	String	Name of a billable item. Example: DiskSpace
PriceName	no	yes	String	Name of the billable item, which is displayed in the user console. Example: CBS space
OriginalPrice	no	yes	Float	Original price of a prepaid cloud disk, in USD. Example: 100.0
DiscountPrice	no	yes	Float	Discounted price of a prepaid cloud disk, in USD. Example: 100.0
UnitPrice	no	yes	Float	Original unit price of postpaid cloud disks (in USD). Example: 10.0
UnitPriceDiscount	no	yes	Float	Discounted unit price of postpaid cloud disks (in USD). Example: 10.0
ChargeUnit	no	yes	String	Billing unit for a postpaid cloud disk. Valid value: HOUR: indicates that the billing unit of a postpaid cloud disk is hour. Example: HOUR
OriginalPriceHigh	no	yes	String	Original price of a high-precision prepaid cloud disk, in USD. Example: 100
DiscountPriceHigh	no	yes	String	Discounted price of a high-precision prepaid cloud disk, in USD. Example: 100
UnitPriceHigh	no	yes	String	Original unit price of a high-precision postpaid cloud disk, in USD. Example: 100
UnitPriceDiscountHigh	no	yes	String	Discounted unit price of a high-precision postpaid cloud disk, in USD. Example: 100

Snapshot

It describes the details of the snapshot.

Referenced by the following interfaces: DescribeDiskAssociatedSnapshots, DescribeSnapshots, DescribeSnapshotsInternal

name	required	NULL allowed	type	description
SnapshotId	no	no	String	Snapshot ID. Example: snap-4agc8cnk
Placement	no	no	Placement	Location of the snapshot. Example:View
DiskUsage	no	no	String	Type of the cloud disk that created this snapshot. Value range: • SYSTEM_DISK: system disk • DATA_DISK: data disk. Example: SYSTEM_DISK
DiskId	no	no	String	ID of the cloud disk that created this snapshot. Example: disk-pe1vkgim
DiskSize	no	no	Uint64	Size of the cloud disk that created this snapshot. Example: 10
SnapshotState	no	no	String	Snapshot status. Value range: • NORMAL: normal • CREATING: creating • ROLLBACKING: rolling back • COPYING_FROM_REMOTE: copying snapshots cross regions. Example: NORMAL
SnapshotName	no	no	String	Snapshot name, which is a custom snapshot alias. This field can be modified by calling the ModifySnapshotAttribute API. Example: snap-1
Percent	no	no	Uint64	Snapshot creation progress percentage. This field is always 100 after the snapshot is created successfully. Example: 100
CreateTime	no	no	String	Snapshot creation time. Example: 2019-07-15 00:00:00
DeadlineTime	no	no	String	Snapshot expiration time. If the snapshot is permanently retained, this field is empty. Example: 2021-11-24 16:00:00
Encrypt	no	no	Bool	Snapshot created by encrypted disks. Value range: • true: This snapshot is created by an encrypted disk. • false: This snapshot is not created by an encrypted disk. Example: false
IsPermanent	no	no	Bool	It indicates whether the snapshot is a permanent snapshot. Value range: • true: permanent snapshot • false: non-permanent snapshot. Example: true
CopyingToRegions	yes	yes	Array of String	It indicates the destination region of the snapshot currently under cross region replication. The default value is []. Example: ["region1"]
CopyFromRemote	yes	yes	Bool	It indicates whether the snapshot is a snapshot for cross region replication. Value range: • true: Indicates a snapshot for cross region replication. • false: Snapshot of this region. Example: false
Images	yes	yes	Array of Image	List of images associated with snapshots. Example:View
ImageCount	yes	yes	Uint64	Number of images associated with the snapshot. Example: 0
SnapshotType	yes	yes	String	Snapshot type. Currently, the type PRIVATE_SNAPSHOT or SHARED_SNAPSHOT is supported. Example: PRIVATE_SNAPSHOT
ShareReference	yes	yes	Uint64	Number of Snapshots Currently Shared Example: 0
Tags	no	yes	Array of Tag	List of snapshot tags. Example:View
PlatformProjectId	no	yes	String	Internal project ID of the snapshot. Example: pr-42749a13
EncryptAlgorithm	no	yes	String	Encryption algorithm. Example: AES256
SavePlaceName	no	yes	String	Snapshot storage location. Example: ttttst
SnapFormat	no	yes	String	Snapshot format type. Valid values: ("RAW", "QCOW2"). Example: RAW
SnapSize	no	yes	Uint64	Snapshot size Example: 50
CopyFromBackupDetail	no	yes	CopyFromBackupDetail	Snapshot replicated from a backup, obtaining the information of the corresponding backup. Example:View

ResizeDiskOrder

Order parameters for request billing (scaling).

Referenced by the following interfaces: SwitchParameterModifyDiskAttributes

name	required	NULL allowed	type	description
GoodsCategoryId	yes	no	Uint64	Product category. Example: 100000
GoodsNum	yes	no	Uint64	Number of products. Example: 1
ProjectId	yes	no	Uint64	Project ID. Example: 0
GoodsDetail	yes	no	ResizeGoodsDetail	Product details. Example:View
RegionId	yes	no	Uint64	Region ID. Example: 10
ZoneId	yes	no	Uint64	Availability zone ID. Example: 100001
PayMode	yes	no	Uint64	Payment mode. Example: 1
Type	yes	yes	String	Cloud disk type. Example: cbs
SubProductCode	yes	yes	String	Sub-product ID. Example: sp_cbs_snr

SnapshotGroup

It describes the snapshot group details.

Referenced by the following interfaces: DescribeSnapshotGroups

name	required	NULL allowed	type	description
SnapshotGroupId	yes	no	String	Snapshot Group ID Example: csnap-csiqd1z7u
SnapshotGroupType	yes	yes	String	Snapshot group type. NORMAL: common snapshot group, non-consistency snapshot Example: NORMAL
ContainRootSnapshot	yes	yes	Bool	It indicates whether the snapshot group includes a system disk snapshot. Example: true
SnapshotIdSet	yes	no	Array of String	ID List of Snapshots in Snapshot Group Example: ["snap-31gfd8na"]

name	required	NULL allowed	type	description
SnapshotGroupState	yes	yes	String	Snapshot group status. ▸ NORMAL: Normal ▸ CREATING: Creating ▸ DELETED: Deleted ▸ FAILED: Creation failed ▸ DISMISS: Dismissed ▸ ROLLBACKING: Rolling back Example: NORMAL
Percent	yes	yes	Uint64	Snapshot Group Creation Progress Example: 100
CreateTime	yes	yes	String	Creation Time Example: 2022-05-25T12:14:53.243893+08:00
ModifyTime	yes	yes	String	Last Modification Time Example: 2022-05-25T12:14:53.243893+08:00
Images	yes	no	Array of Image	List of Images Associated with Snapshot Group Example:View
SnapshotGroupName	yes	yes	String	Snapshot Group Name Example: snp092101
ImageCount	yes	yes	Uint64	Number of Images Associated with Snapshot Group Example: 0
LatestOperationState	yes	yes	String	Result of Last Executed Asynchronous Task Example: SUCCESS
LatestOperationRequestId	yes	yes	String	It indicates the API RequestId corresponding to the last executed asynchronous task. Example: 597205a3-68c8-4f8a-bd3d-91cf0fa5a46c
LatestOperation	yes	yes	String	Last Executed Asynchronous Task Operation Example: ApplySnapshot
OperationFailedResourceIdSet	yes	yes	Array of String	List of Resources for Which Last Asynchronous Task Operation Failed Example: snap-dwl3bti9l
OperationSuccessResourceIdSet	yes	yes	Array of String	List of Successfully Operated Resources in Last Asynchronous Operation Example: snap-pfsifv2l3
ErrorPrompt	yes	yes	String	Prompt Code of Last Asynchronous Operation Result Example: SUCCESS
IsPermanent	yes	yes	Bool	Whether the snapshots copied across regions are retained permanently. The default value is True, indicating the snapshots are retained permanently. Currently, this parameter is only valid for snapshots automatically initiated by ASP. Snapshots initiated by users are retained permanently by default. Example: true
DeadlineTime	yes	yes	Datetime_iso	Deadline Example: 2022-05-25T12:14:53.243893+08:00
Tags	yes	no	Array of Tag	Tag Value Example:View

SnapshotGroupDeniedAction

Operation Mask of Single Snapshot Group

Referenced by the following interfaces: DescribeSnapshotGroupsDeniedActions

name	required	NULL allowed	type	description
SnapshotGroupId	yes	no	String	Snapshot Group ID Example: csnap-csoptz7u
DeniedActions	yes	no	Array of DeniedAction	Specific Operation Mask Information for Snapshot Example:View

DiskConfig

Cloud disk configuration.

Referenced by the following interfaces: DescribeDiskConfigQuota

name	required	NULL allowed	type	description
Available	no	no	Bool	Whether the configuration is available. Example: false
DiskType	no	no	String	Cloud disk media type. Value range: ▸ CLOUD_BASIC: standard cloud disk ▸ CLOUD_PREMIUM: high-performance cloud disk ▸ CLOUD_SSD: SSD cloud disk Example: CLOUD_BASIC
DiskUsage	no	no	String	Cloud disk type. Value range: ▸ SYSTEM_DISK: system disk ▸ DATA_DISK: data disk. Example: DATA_DISK
DiskChargeType	no	no	String	Payment mode. Value range: ▸ PREPAID: Indicates prepaid, that is, monthly subscriptions ▸ POSTPAID_BY_HOUR: Indicates postpaid, that is, pay-as-you-go. Example: POSTPAID_BY_HOUR
MaxDiskSize	no	no	Uint64	Maximum configurable cloud disk size. Example: 32000
MinDiskSize	no	no	Uint64	Minimum configurable cloud disk size. Example: 10
Zone	no	no	String	Availability zone. Example: ap-region1-2
DeviceClass	no	yes	String	Instance model. Example: VSELF_2
InstanceFamily	no	yes	String	Instance model series. For details, see the instance types. Example: N1
StepSize	yes	yes	Uint64	Step length Example: 10
ExtraPerformanceRange	no	yes	Array of Int64	Additional performance range. Example: 0
RecommendDiskSize	no	yes	Uint64	Recommended size for the current type of cloud disk, in GB. Example: 100

DiskConfigSet

Cloud Disk Configuration

Referenced by the following interfaces: DescribeZoneDiskConfigInfos

name	required	NULL allowed	type	description
DiskUsage	yes	no	String	System Disk or Data Disk Example: DATA_DISK
DiskType	yes	no	String	CLOUD_BASIC, CLOUD_PREMIUM, CLOUD_SSD Example: CLOUD_SSD
DiskSize	yes	no	Int64	Cloud Disk Size Example: 100
SnapshotId	no	no	String	Snapshot ID Example: snap-7vh165a5

DeniedAction

API That Cannot Be Executed by Cloud Disk

Referenced by the following interfaces: DescribeDisksDeniedActions, DescribeSnapshotGroupsDeniedActions, DescribeSnapshotsDeniedActions

name	required	NULL allowed	type	description
Action	yes	no	String	Name of the inoperable API. For example, when you query the cloud disk list "DescribeDisks"; for multipurpose APIs, the response is in the form of "API name.input parameter", such as "ModifyDiskAttributes.DiskType". Example: DetachDisks
Message	yes	no	String	Reason Why the API Cannot Be Operated Example: Only mounted elastic data disks support uninstallation.
Code	yes	no	String	Error Code for API Operation Failure Example: ResourceUnavailable.NotSupported

ErrorCode

Feature Description

If there is an Error field in the response, it means that the API call failed. For example:

```
{
    "Response": {
        "Error": {
            "Code": "AuthFailure.SignatureFailure",
            "Message": "The provided credentials could not be validated. Please check your signature is cor
rect."
        },
        "RequestId": "ed93f3cb-f35e-473f-b9f3-0d451b8b79c6"
    }
}
```

Code in Error indicates the error code, and Message indicates the specific information of the error.

Error Code List

Common Error Codes

Error codes	Description
AuthFailure.InvalidSecretId	Invalid key (not a Cloud API key type).
AuthFailure.MFAFailure	MFA failed.
AuthFailure.SecretIdNotFound	Key does not exist. Check if the key has been deleted or disabled in the console, and if not, check if the key is correctly entered. Note that whitespaces should not exist before or after the key.
AuthFailure.SignatureExpire	Signature expired. Timestamp and server time cannot differ by more than five minutes. Please ensure your current local time matches the standard time.
AuthFailure.SignatureFailure	Invalid signature. Signature calculation error. Please ensure you've followed the signature calculation process described in the Signature API documentation.
AuthFailure.TokenFailure	Token error.

Error codes	Description
AuthFailure.UnauthorizedOperation	The request does not have CAM authorization.
DryRunOperation	DryRun Operation. It means that the request would have succeeded, but the DryRun parameter was used.
FailedOperation	Operation failed.
InternalError	Internal error.
InvalidAction	The API does not exist.
InvalidParameter	Incorrect parameter.
InvalidParameterValue	Invalid parameter value.
LimitExceeded	Quota limit exceeded.
MissingParameter	A parameter is missing.
NoSuchVersion	The API version does not exist.
RequestLimitExceeded	The number of requests exceeds the frequency limit.
ResourceInUse	Resource is in use.
ResourceInsufficient	Insufficient resource.
ResourceNotFound	The resource does not exist.
ResourceUnavailable	Resource is unavailable.
UnauthorizedOperation	Unauthorized operation.
UnknownParameter	Unknown parameter.
UnsupportedOperation	Unsupported operation.
UnsupportedProtocol	HTTPS request method error. Only GET and POST requests are supported.
UnsupportedRegion	API does not support the requested region.

Service Error Codes

Error codes	Description
InvalidDisk.Attached	The cloud disk has been mounted.

Error codes	Description
InvalidDisk.NotPortable	Non-elastic cloud disks are not supported.
InvalidAutoSnapshotPolicyId.NotFound	The input AutoSnapshotPolicyId does not exist.
InvalidDisk.Busy	The cloud disk is busy. Try again later.
ResourceInUse.DiskMigrating	The cloud disk is migrating.
TradeDealConflict	Order conflict.
InternalError.ResourceOpFailed	The operation on the resource failed. For specific error information, view the Message field in the error description. Try again later or contact customer service for help.
InvalidDiskId.NotFound	The input DiskId does not exist.
InsufficientRefundQuota	The number of returned cloud disks has reached the upper limit and no more cloud disks can be returned.
ResourceBusy	The resource is busy. Try again later.
InvalidFilter	The specified filter is not supported.
ResourceInsufficient	Insufficient resources
InvalidAccount.InsufficientBalance	Insufficient account balance
InvalidDisk.SnapshotCreating	The cloud disk is creating a snapshot. Try again later.
UnsupportedOperation.SnapshotHasBindedImage	The snapshot is already associated with an image.
AutoSnapshotPolicyOutOfQuota	The number of regular snapshot policies has reached the upper limit.
UnsupportedOperation	Operation is not supported

Error codes	Description
InvalidDisk.Expire	The cloud disk has expired.
InvalidInstance.NotSupported	The Cloud Virtual Machine does not support mounting a cloud disk.
InternalError.FailQueryResource	Failed to query the resource.
ResourceUnavailable.TypeError	Disk type error.
InvalidSnapshot.NotSupported	The snapshot does not support this operation.
InvalidParameter.DiskSizeNotMatch	The cloud disk size does not match the snapshot size.
InternalError.ComponentError	Third-party Component Error
ResourceInsufficient.OverQuota	Insufficient quota
UnsupportedOperation.SnapshotGroupDiskAttachMultiInstance	The original disk of the snapshot in the snapshot group is not mounted to the same instance.
InvalidDisk.NotSupported	The cloud disk does not support this operation.
InvalidParameter	Parameter error
InvalidParameter.DiskConfigNotSupported	The current region does not support the cloud disk of the current configuration.
InvalidParameterValue	Invalid parameter value. The format of the parameter value is incorrect or the parameter value is not supported.
MissingParameter	Missing parameter. The request does not include the mandatory parameter.
ResourceNotFound.NotFound	The resource does not exist.
InvalidDisk.TypeError	Incorrect cloud disk type.
InvalidInstanceId.NotFound	The input InstanceId does not exist.

Error codes	Description
ZoneNotMatch	The cloud disk and the instance are not in the same availability zone.
InvalidDisk.RepeatRefund	The cloud disk has been returned before and cannot be returned again.
InsufficientSnapshotQuota	The snapshot quota is insufficient.
ResourceUnavailable.SnapshotCreating	The snapshot is being created.
InvalidParameterValue.LimitExceeded	The parameter value is out of range
InvalidSnapshotId.NotFound	The input SnapshotId does not exist.
InvalidDisk.NotSupportSnapshot	The cloud disk does not have the snapshot capability.
InvalidDisk.NotSupportRefund	The cloud disk cannot be returned.