

Cloud Block Storage

Product Introduction

Product Documentation



Copyright Notice

©2013-2025 Tencent Cloud. All rights reserved.

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

Trademark Notice



All trademarks associated with Tencent Cloud and its services are owned by the Tencent corporate group, including its parent, subsidiaries and affiliated companies, as the case may be. Trademarks of third parties referred to in this document are owned by their respective proprietors.

Service Statement

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

Contents

Product Introduction

- Overview

- Product Strengths

- Use Cases

- Regions and AZs

Cloud Disk Instances

- Cloud Disk Types

- Cloud Disk Status

- Cloud Disk Encryption

- Enhanced SSD Performance

- Cloud Disk Data Backup Point

Cloud Disk Snapshots

- Snapshot Overview

- Snapshot Principles

Access Management

- An Overview of Access's Management

- Authorized resource types

- Authorization Policy Syntax

Related Products

- Predecessor

Product Introduction

Overview

Last updated : 2023-12-20 16:46:37

What is Tencent Cloud Block Storage?

Cloud Block Storage (CBS) is a highly available, highly reliable, low-cost, and customizable block storage device. It can be used as an independent and scalable disk for CVM, providing efficient and reliable [storage](#) devices. CBS provides long-term storage at the data block level. It is typically used as the primary storage device for data that requires frequent and fine-grained updates (such as file system and database), featuring high availability, reliability, and performance. CBS uses a three-copy distributed mechanism to back up your data on different physical machines to avoid data loss caused by a single point of failure, improving data reliability.

You can easily purchase, adjust, and manage your cloud disk devices through the console, and build a file system to create storage space larger than that of a single cloud disk. Cloud disks can be classified according to different lifecycles as follows:

Lifecycle of non-elastic cloud disk completely follows that of the CVM. It is purchased with the CVM and used as a system disk. It does not support mounting and unmounting.

Lifecycle of elastic cloud disk is independent of that of the CVM. It can be purchased separately and manually mounted to the CVM. It can also be purchased with the CVM and automatically mounted to the CVM as a data disk. Elastic cloud disk supports mounting to and unmounting from the CVM in the same availability zone at any time. You can mount multiple elastic cloud disks to the same CVM, or unmount a cloud disk from CVM A and then mount it to CVM B.

Tencent Cloud places quota limits on users' cloud disks. For more information, see [Use Limits](#).

Typical Use Cases

You can purchase and mount one or more cloud disks to meet the requirements for storage capacity when the disk space on your CVM is insufficient.

When purchasing the CVM, you do not need additional storage space. When you have storage requirements, you can expand the CVM storage capacity by purchasing a cloud disk.

When there is a data exchange request between multiple CVMs, you can unmount cloud disks (data disks) and remount them to other CVMs.

You can expand the maximum storage capacity of a single cloud disk by purchasing multiple cloud disks and configuring a Logical Volume Manager (LVM).

You can expand the maximum I/O capacity of a single cloud disk by purchasing multiple cloud disks and configuring a Redundant Array of Independent Disks (RAID) policy.

Features

Tencent Cloud provides a variety of long-term storage devices, allowing the user to select an appropriate type of cloud disk to store files, build databases, etc.

Tencent Cloud CBS comes with the following disk types: Premium Cloud Disk, SSD, Enhanced SSD, and ulTra SSD.

Elastic mounting and unmounting: all types of elastic cloud disks support elastic mounting and unmounting. You can mount multiple cloud disks on a CVM to build a file system with high capacity.

Elastic expansion: a single disk supports a maximum capacity of 32 TB. You can scale up the disk at any time.

Snapshot backup: you can create and roll back a snapshot to back up data. You can also create a cloud disk from the snapshot to accelerate your business deployment.

Product Strengths

Last updated : 2023-12-20 17:02:06

Reliable

Cloud disks use a three-copy distribution mechanism, and the system ensures that data is written in three copies before returning a response of "written successfully". If any of the copies fail, the backend data replication mechanism can quickly replicate a new copy by using methods such as data migration, ensuring the availability of three data copies at all times to provide you with a secure storage service. Data is stored across racks, and reliability reaches up to 99.9999999%.

Elastic

You can configure the storage capacity freely, and expand the capacity on demand without business interruption. The maximum capacity of cloud disks is 32 TB. A single CVM can attach up to 20 elastic cloud disks as data disks, to efficiently handle TB/PB-level big data processing scenarios and meet the requirements of large-capacity file systems for big data, data warehouse, log processing and other businesses.

High Performance

Premium Cloud Disk uses a cache mechanism to satisfy the normal business requirements of users. SSD uses NVMe standards and can provide 26,000 random read-write Input/Output Operations per Second (IOPS) on a single disk. This is suitable for scenarios with high requirements for I/O capability.

Easy-to-Use

You can easily manage and use your cloud disks through simple operations such as creation, attaching, detaching, and deletion, reducing both costs and business deployment time.

Snapshot Backup

You can create a snapshot of the cloud disk to back up data at any time, and you can also quickly create a cloud disk by using a snapshot file, achieving rapid business deployment.

Use Cases

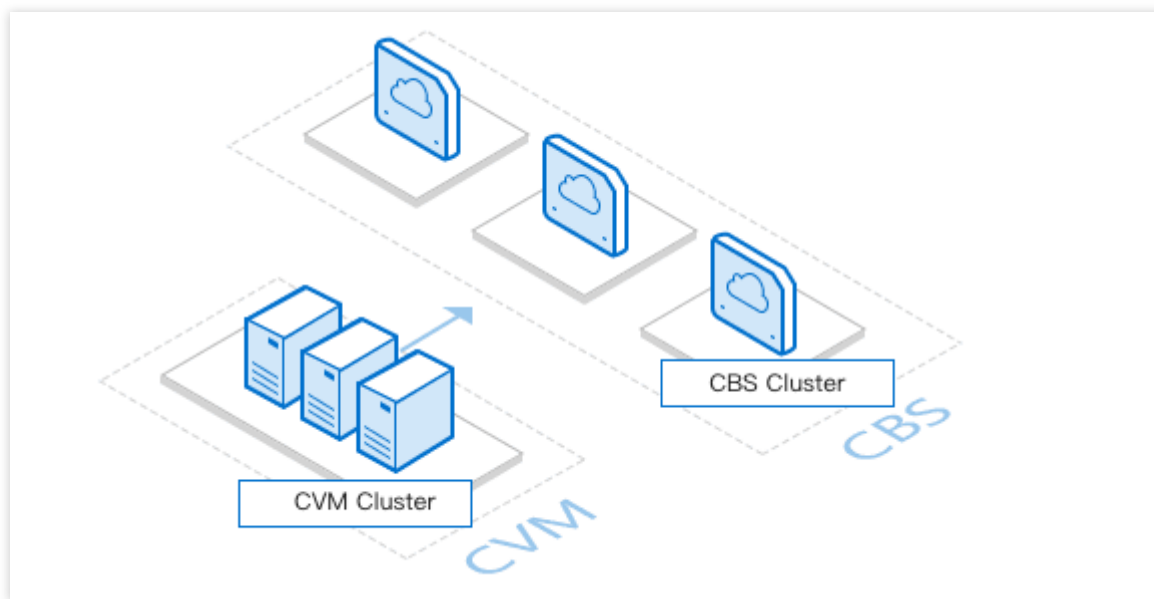
Last updated : 2023-12-20 17:28:04

Typical Use Cases

Delocalization

Data storage with high performance and high reliability: Cloud Block Storage (CBS) efficiently supports hot migration of CVMs, preventing business interruption caused by physical failures. CBS uses the three-copy data redundancy mechanism to back up your data and snapshots, and recover data within seconds, making it suitable for high-load and mission-critical systems.

Elastic expansion: cloud disks can be freely mounted and unmounted in the same availability zone without shutting down or restarting the CVM. Cloud disk capacity can be elastically configured, and expanded on demand.

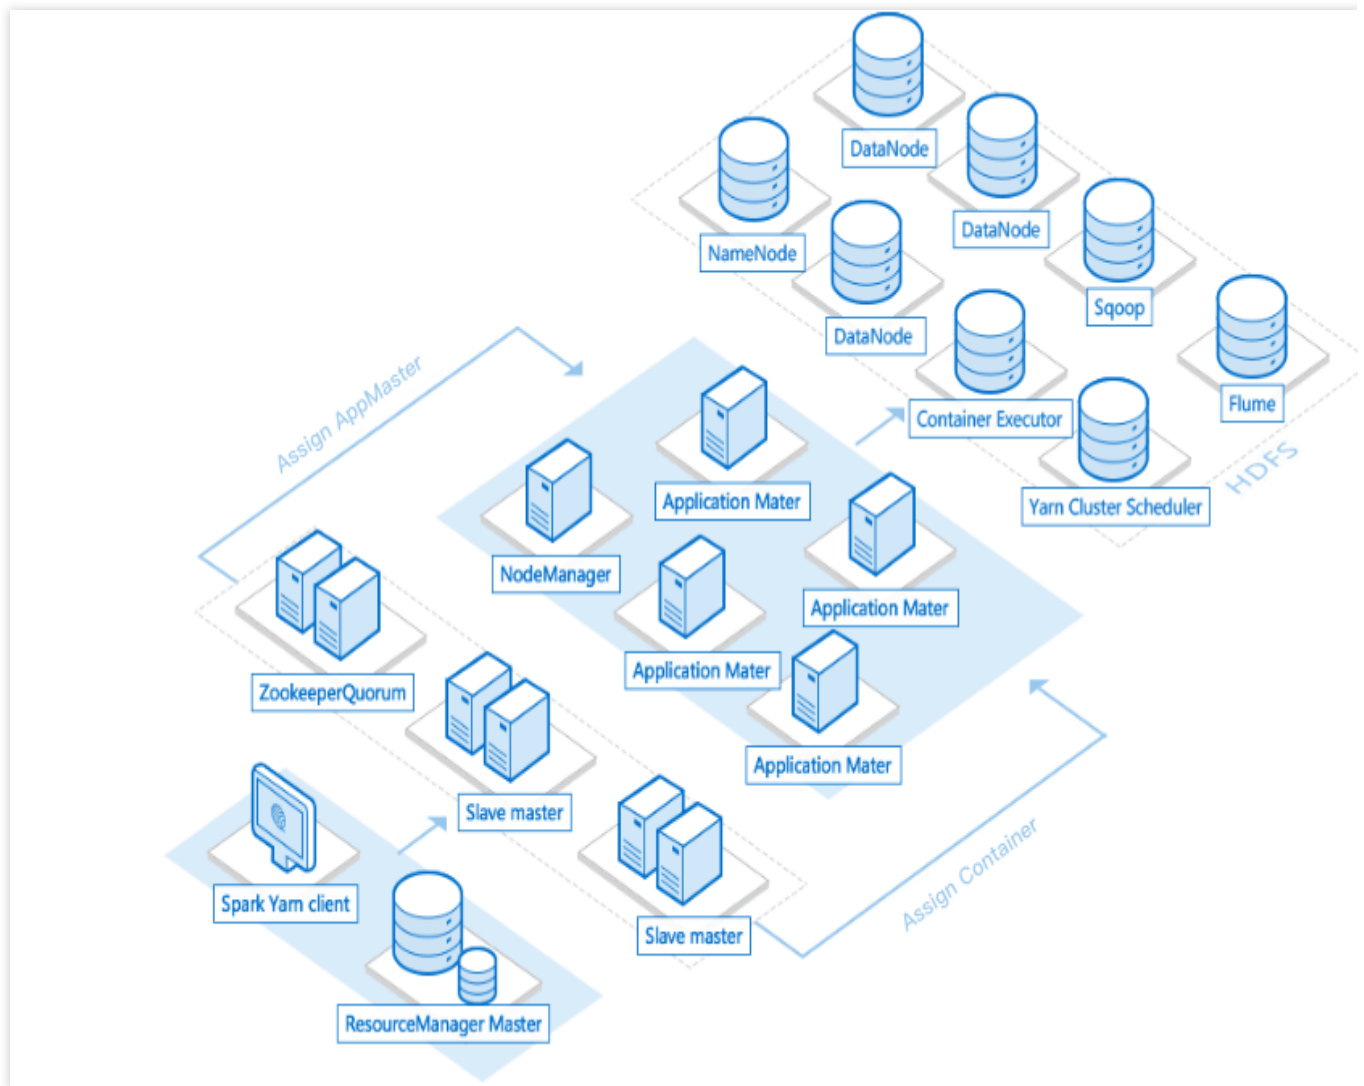


Massive data analysis

In a typical Spark-HDFS offline data analytics framework, RDD reads/writes and shuffle writes of disks by are all sequential IO, except for the random shuffle read I/O. The sequential I/O accounts for 95%. CBS features an excellent multi-thread concurrent throughput performance, enabling efficient offline data processing at the terabyte and petabyte levels for Hadoop-Mapreduce, HDFS, and Spark.

With multi-disk concurrency, a single HDFS cluster can achieve a throughput up to 1 GB/s.

CBS supports big data applications such as data analytics, data mining, and business intelligence for companies like Xiaohongshu, Giant Interactive Group, Ele, Yoho!BUY, and wepiao.com.



Deployment environment: 5 CVM servers (with 12-Core 40 GB RAM), simulating offline data analysis of 1.5TB data volume.

Test performance:

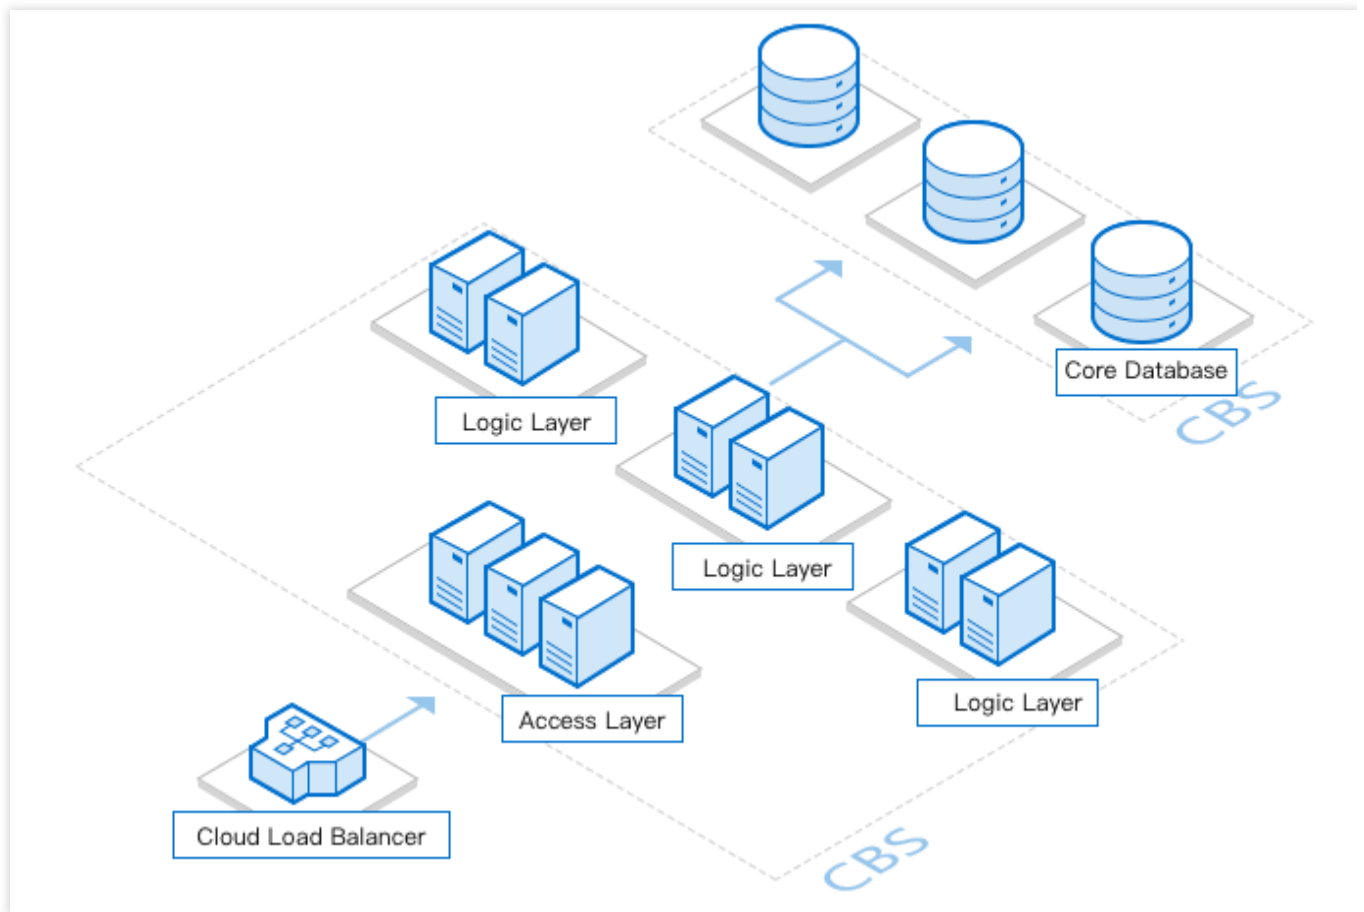
Each CVM has mounted one 1 TB HDD cloud disk. Five HDD cloud disks provide a read speed of 500 MB/s, allowing data to be read to memory in 50 minutes.

Each CVM has mounted one 1 TB SSD, allowing data to be read to memory in 25 minutes.

Core database

SSD is ideal for scenarios with high requirements for IO performance and data reliability. It is particularly suitable for medium and large relational database applications like PostgreSQL, MySQL, Oracle, and SQL Server; for IO-intensive core business systems with high data reliability requirements; and for medium and large development and testing environments with high data reliability requirements.

SSD offers both data reliability and high performance. It constantly provides reliable support for companies such as Heroes Evolved, Wendao, Yoho!BUY, weipiao.com, Xiaohongshu, etc.

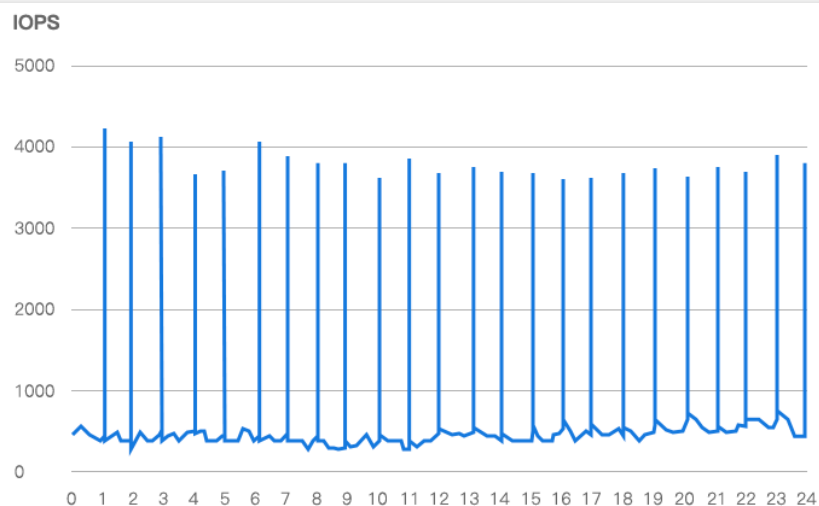


Deployment environment: 4 CVM servers (with 4-Core 8 GB RAM). Each has one 800 GB SSD cloud disk mounted, with MySQL version 5.5.42 deployed.

Test performance: simulate OLTP performance testing using sysbench, with a test set of 10 million records. In this test, TPS and QPS reach 1,616 and 29,000 respectively, meaning a single disk is sufficient to support 10 thousand concurrent transactions per second.

Typical business scenarios for I/O models

Hourly data persistence



Hourly data synchronization

Typical large-scale e-commerce business servers sync large amounts of data to the disk (asynchronous data placement) on the hour every day. Disk load is relatively high at that moment, with a read-write ratio of approximately 3:7.

Peak IOPS varies between 3,000–5,000, peak throughput varies between 10–120MB/s, and duration varies between 1–5 minutes. At other times, the load is relatively low. We recommend you use high-performance cloud disks.

Typical Users:




Ele.me



Xiaohongshu

High load OLTP services



High-load OLTP services

CBS hybrid cloud disks guarantee stable disk IO performance 99.5% of the time, avoiding unexpected fluctuations that impact user experience. This service meets the online transaction processing (OLTP) needs of the finance industry by working day and night to fulfill high-load IO read-write requirements.

Typical Users:




webank



Futu Securities

Periodic ultra-high load



Periodic ultra-high load

Users include top Chinese game developers such as 37Games, Wendao Games, and Heroes Evolved. Peak access is between 7–10 PM and 3–5 AM.

During peak hours, the application program is transmitted to the file system. The IO requests transmitted to the block device are random small amounts (either 4K, 8K or 16K), with peak demand and throughput exceeding 20,000 IOPS and 200MBps respectively. If disk performance can not meet this demand, problems such as `io await` or `%uitl` will occur, causing game players to experience stalling or disconnection. We recommend you use SSD hard disk for this scenario.

Typical Users:



Wendao Games




Feng Le Games



Heroes Evolved

Continuous sequential read-write

Bandwidth (MB/s)



Continuous high sequential read-write

A typical Spark-HDFS offline data analysis framework for reading and writing to the disk. RDD read/write and shuffle write use sequential IO. Only shuffle read IO is random IO. 95% of disk IO is sequential IO. Log collection and distribution are dependent on Kafka.

When Kafka writes a message, it writes to page cache, and the thread is asynchronously flushed. When it reads a message, the message is directly transferred from page cache to the socket and sent (the sendfile function implements zero copy). When no corresponding data is found in page cache, disk IO is generated, messages are loaded from the disk to the page cache, and then directly sent from the socket. All I/O is sequential.

We recommend using hybrid cloud disk, which provides reliable long-term throughput performance that can reach up to 130MB/s on a single disk, and up to 1GB/s if a software RAID solution is overlaid.

Typical Users: :



Wind Play



Netdragon

Regions and AZs

Last updated : 2024-10-28 16:43:42

Regions

A region is the physical location of an IDC. In Tencent Cloud, regions are fully isolated from each other, ensuring cross-region stability and fault tolerance. When purchasing Tencent Cloud services, we recommend selecting the region closest to your end users to minimize access latency and improve download speed.

Regions have the following characteristics:

The networks of different regions are fully isolated. Tencent Cloud services in different regions **cannot communicate via a private network by default**.

Cloud products in different regions can access the Internet through the [public network service](#). Cloud products in a VPC can also use Tencent Cloud's [Peering Connection](#) to communicate with each other by using Tencent Cloud's high-speed network to implement a connection that is faster and more stable than Internet access.

[Cloud Load Balancer](#) currently supports intra-region traffic forwarding by default. If you enable the [cross-region binding](#) feature, a CLB instance can be bound to CVM instances in another region.

Availability Zones

An availability zone (AZ) is a physical IDC of Tencent Cloud with independent power supply and network in the same region. It can ensure business stability, as failures (except for major disasters or power failures) in one AZ are isolated without affecting other AZs in the same region. By starting an instance in an independent AZ, users can protect their applications from being affected by a single point of failure.

Availability zones have the following characteristics:

For cloud products under the same Tencent Cloud account, products in the same region but different availability zones and under the same [VPC](#) can communicate with each other through the private network, and the [private network service](#) can be used for direct access.

Resources under different Tencent Cloud accounts are availability zone-specific, meaning they cannot communicate via a private network, even if they are in the same region.

China

Region	AZ
South China (Guangzhou)	Guangzhou Zone 1 (sold out)

ap-guangzhou	ap-guangzhou-1
	Guangzhou Zone 2 (sold out) ap-guangzhou-2
	Guangzhou Zone 3 ap-guangzhou-3
	Guangzhou Zone 4 ap-guangzhou-4
	Guangzhou Zone 6 ap-guangzhou-6
	Guangzhou Zone 7 ap-guangzhou-7
East China (Shanghai) ap-shanghai	Shanghai Zone 1 (sold out) ap-shanghai-1
	Shanghai Zone 2 ap-shanghai-2
	Shanghai Zone 3 ap-shanghai-3
	Shanghai Zone 4 ap-shanghai-4
	Shanghai Zone 5 ap-shanghai-5
	Shanghai Zone 8 ap-shanghai-8
East China (Nanjing) ap-nanjing	Nanjing Zone 1 ap-nanjing-1
	Nanjing Zone 2 ap-nanjing-2
	Nanjing Zone 3 ap-nanjing-3
North China (Beijing) ap-beijing	Beijing Zone 1 (sold out) ap-beijing-1
	Beijing Zone 2

	ap-beijing-2
	Beijing Zone 3 ap-beijing-3
	Beijing Zone 4 ap-beijing-4
	Beijing Zone 5 ap-beijing-5
	Beijing Zone 6 ap-beijing-6
	Beijing Zone 7 ap-beijing-7
	Beijing Zone 8 ap-beijing-8
Southwest China (Chengdu) ap-chengdu	Chengdu Zone 1 ap-chengdu-1
	Chengdu Zone 2 ap-chengdu-2
Southwest China (Chongqing) ap-chongqing	Chongqing Zone 1 ap-chongqing-1
Hong Kong/Macao/Taiwan (Hong Kong, China) ap-hongkong	Hong Kong Zone 1 (Nodes in Hong Kong, China can cover Hong Kong/Macao/Taiwan regions) (sold out) ap-hongkong-1
	Hong Kong Zone 2 (Nodes in Hong Kong, China can cover Hong Kong/Macao/Taiwan regions) ap-hongkong-2
	Hong Kong Zone 3 (Nodes in Hong Kong, China can cover Hong Kong/Macao/Taiwan regions) ap-hongkong-3

Note:

The product is in beta test for Jinan, Hangzhou, Fuzhou, Wuhan, Changsha, and Shijiazhuang regions. To try it out, submit a ticket for application.

Other Countries and Regions

Region	AZ
Southeast Asia (Singapore) ap-singapore	Singapore Zone 1 (Nodes in Singapore can cover Southeast Asia) ap-singapore-1
	Singapore Zone 2 (Nodes in Singapore can cover Southeast Asia) ap-singapore-2
	Singapore Zone 3 (Nodes in Singapore can cover Southeast Asia) ap-singapore-3
	Singapore Zone 4 (Nodes in Singapore can cover Southeast Asia) ap-singapore-4
Southeast Asia (Jakarta) ap-jakarta	Jakarta Zone 1 (Nodes in Jakarta can cover Southeast Asia) ap-jakarta-1
	Jakarta Zone 2 (Nodes in Jakarta can cover Southeast Asia) ap-jakarta-2
Northeast Asia (Seoul) ap-seoul	Seoul Zone 1 (Nodes in Seoul can cover Northeast Asia) ap-seoul-1
	Seoul Zone 2 (Nodes in Seoul can cover Northeast Asia) ap-seoul-2
Northeast Asia (Tokyo) ap-tokyo	Tokyo Zone 1 (Nodes in Tokyo can cover Northeast Asia) ap-tokyo-1
	Tokyo Zone 2 (Tokyo nodes can cover services in Northeast Asia) ap-tokyo-2
Southeast Asia (Bangkok) ap-bangkok	Bangkok Zone 1 (Nodes in Bangkok can cover Southeast Asia) ap-bangkok-1
	Bangkok Zone 2 (Nodes in Bangkok can cover Southeast Asia) ap-bangkok-2
South America (São Paulo) sa-saopaulo	Sao Paulo Zone 1 (Sao Paulo nodes cover services in South America) sa-saopaulo-1
Western US (Silicon Valley) na-siliconvalley	Silicon Valley Zone 1 (Nodes in Silicon Valley can cover Western US) na-siliconvalley-1
	Silicon Valley Zone 2 (Nodes in Silicon Valley can cover Western US)

	na-siliconvalley-2
Eastern US (Virginia) na-ashburn	Virginia Zone 1 (Nodes in Virginia can cover Eastern US) na-ashburn-1
	Virginia Zone 2 (Nodes in Virginia can cover Eastern US) na-ashburn-2
Europe (Frankfurt) eu-frankfurt	Frankfurt Zone 1 (Nodes in Frankfurt can cover Europe) eu-frankfurt-1
	Frankfurt Zone 2 (Nodes in Frankfurt can cover Europe) eu-frankfurt-2

Selection of Regions and AZs

When selecting a region and availability zone, take the following into consideration:

A cloud disk can only be attached to a CVM in the same availability zone.

The region of the CVM on which the cloud disk will be used, your location, and the location of your users. We recommend that you choose the region closest to your end users when purchasing Tencent Cloud services to minimize access latency and improve access speed.

The relationship between the CVM on which the cloud disk will be used and other Tencent Cloud services. When you select other Tencent Cloud services, we recommend you try to locate them all in the same availability zone of a single region to allow them to communicate with each other through the private network, reducing access latency and increasing access speed.

High availability and disaster recovery. Even if you have just one VPC, we still recommend that you deploy your businesses in different availability zones to prevent a single point of failure and enable cross-AZ disaster recovery.

There may be network latency among different AZs. We recommend you assess your business requirements and find the optimal balance between high availability and low latency.

Related Operations

Actions such as the use and viewing of cloud disks are region-specific. To easily migrate data and services to other regions or construct a cross-region disaster recovery system, you can copy snapshots to other regions. For more information, see [Cross-region Snapshot Replicating](#).

Cloud Disk Instances

Cloud Disk Types

Last updated : 2025-06-09 11:45:09

Cloud Block Storage (CBS) provides highly available, highly reliable, low-cost, and customizable network block device that can be used as a standalone and expandable disk for CVMs. CBS stores data at the data block level in a three-copy distributed mechanism to ensure data reliability. CBS is classified into five types: **Premium Disk**, **Balanced SSD**, **Cloud SSD**, **Enhanced SSD**, and **Tremendous SSD**. Each type has unique performance and characteristics, and the price varies, making CBS suitable for different use cases.

Notes

Currently, Enhanced SSD and Tremendous SSD are only available in certain availability zones. They will be supported in more availability zones.

The performance of Enhanced SSD is only guaranteed when it's attached to S5, M5, and SA2 models created after August 1, 2020, and all later generation models.

Tremendous SSD is exclusively available for purchase with certain instance types. It is not supported by other types of cloud servers.

Tremendous SSD cannot be used as the system disk.

Enhanced SSD and Tremendous SSD cannot be upgraded from other disk types.

Overview

Premium Disk

Premium Disk is a hybrid storage solution launched by Tencent Cloud, providing near-solid-state storage performance through the Cache mechanism. It is suitable for medium to small-scale use cases such as Web/App servers, business logic processing, and medium to small-scale website construction, all of which require high data reliability and moderate performance needs.

Balanced SSD

Balanced SSD is an entry-level all-flash block storage product. It's highly cost-effective and suitable for medium applications with high requirements for data reliability and standard requirements for performance, such as Web/App servers, business logical processing, KV services, as well as basic database services.

Cloud SSD

Cloud SSD, provided by Tencent Cloud, is an all-flash storage type based on NVMe SSD storage media, providing low latency, high random IOPS, and throughput I/O services. It is suitable for scenarios that demand high I/O

performance.

Enhanced SSD

Enhanced SSD, designed by Tencent Cloud based on a new-generation storage engine and leveraging all NVMe SSD storage media along with the latest network infrastructure, is a product type that employs a triple-replica distributed mechanism to provide low latency, high random IOPS, and high throughput I/O services. It is suitable for I/O-intensive scenarios such as large databases and NoSQL, which require low latency. Enhanced SSD supports independent configuration of performance and capacity, allowing you to configure the storage capacity and performance according to the actual needs of your business.

Tremendous SSD

Tremendous SSD is powered by Tencent Cloud's latest high-performance distributed storage engine, high-speed network infrastructure, and the latest storage hardware. It boasts long-term and stable performance with ultra low latency. It is suitable for I/O-intensive and throughput-intensive workloads that require ultra low latency, such as large databases (MySQL, HBase, Cassandra, etc.), key-value storage models (etcd, rocksdb, etc.), log search service (Elasticsearch, etc.), and real-time high-bandwidth businesses (video processing, live streaming, etc.). It performs well in key transaction workloads, core database services, large-scale OLTP services, video processing, and other scenarios. Uniquely, the performance and capacity of Tremendous SSD cloud disks can be independently adjusted to meet your requirements.

Performance metrics

The table below compares the performances of the five CBS services.

Performance Metric	Tremendous SSD	Enhanced SSD	Cloud SSD	Balanced SSD	Premium Disk
Max capacity of a single disk (GB)	32000	32000	32000	32000	32000
Max IOPS of a single disk	Up to 1,000,000 after stacking extra performance	Up to 100,000 after stacking extra performance	26,000	10,000	6,000
Random IOPS performance	Baseline Performance: Random IOPS = Min {4000 + Capacity (GiB) x 100, 50000}	Baseline Performance: Random IOPS = Min {1800 + Capacity (GiB) x 50, 50000}	Random IOPS = Min {1800 + Capacity (GiB) x 30, 26000}	Random IOPS = Min {1800 + Capacity (GiB) x 15, 10000}	Random IOPS = Min {1800 + Capacity (GiB) x 8, 6000}

	Extra Performance: Max IOPS = Min {Extra Performance Value x 128, 950000}	Extra Performance: Max IOPS = Min {Extra Performance Value x 128, 50000} For details, see Enhanced SSD Performance			
Max throughput of a single disk (MB/s)	Up to 4,000 MB/s with extra performance	Up to 1,000 MB/s with extra performance	260 MB/s	190 MB/s	150 MB/s
Throughput performance (MB/s)	Baseline Performance: Throughput = Min {120 + Capacity (GiB) x 0.5, 350} Extra Performance: Throughput = Min {Extra Performance Value x 1, 3650}	Baseline Performance: Throughput = Min {120 + Capacity (GiB) x 0.5, 350} Extra Performance: Throughput = Min {Extra Performance Value x 1, 650} For details, see Enhanced SSD Performance	Throughput = Min {120 + Capacity (GiB) x 0.2, 260}	Throughput = Min {100 + Capacity (GiB) x 0.2, 190}	Throughput = Min {100 + Capacity (GiB) x 0.15, 150}
Single-path random read/write latency	0.1 - 0.5 ms	0.3 - 1 ms	0.5 - 3 ms	0.5 - 3 ms	0.8 - 5 ms
Note	Tremendous SSD is exclusively available for purchase with certain instance types (enhanced storage standard S5se , standard	The performance of Enhanced SSD is only guaranteed when it's attached to S5, M5, and SA2 models and all later generation models.	N/A	N/A	N/A

	<p>SA3, standard S6, standard SA5, standard S8, and memory-optimized MA5). They cannot be used with other types of cloud servers, nor can they be purchased as standalone products.</p> <p>The max IOPS and max throughput performance of a single disk of Tremendous SSD are limited by the type of instance to which it is attached. Currently, the maximum performance of a single disk is reached exclusively when Tremendous SSD is attached to an S5se instance type. The metric information of the maximum achievable storage performance for different instance types are subject to continuous updates, and we encourage you to stay informed.</p>	<p>The max IOPS and the max throughput performance of a single disk of Enhanced SSD are limited by the type of instance to which it is attached. Currently, only when Enhanced SSD is attached to the latest generation instances (S6 and SA3), the maximum performance of a single disk can be achieved. Note that the metric information on the achievable maximum storage performance for different instance types will be continuously updated.</p>			
--	---	---	--	--	--

Note:

The main difference among cloud disks is the I/O performance.

The max IOPS performance is tested out in 4 KiB I/O cases, and the max throughput is tested out at 256 KiB I/O cases. For specific test methods, see [Measuring cloud disk performance](#).

Application Scenarios

Enhanced SSD is more suitable for latency-sensitive or I/O-intensive scenarios, including:

High performance and high data reliability: Suitable for high-load and core critical business systems. Features comprehensive data backup, snapshot, and second-level data recovery capabilities.

Medium and large databases: Support medium and large relational database applications that contain tables with millions of rows, such as MySQL, Oracle, SQL Server, and MongoDB.

Large NoSQL: Support NoSQL businesses such as HBase and Cassandra.

Elasticsearch: Support low-latency ES storage.

Video service: Suitable for applications with high requirements for storage bandwidth, such as audio/video encoding and decoding, live streaming and recording playback.

Big data analysis: Suitable for data analysis, data mining, business intelligence, and other fields. Provide distributed processing capabilities for data at TB and PB levels.

Tremendous SSD is more suitable for latency-sensitive scenarios that require ultra low latency, including:

Key-value (KV) storage: Support rocksdb, etcd, etc. The KV storage service generally writes data to disk in the serial I/O mode, which requires ultra low latency. The single thread latency determines the overall system performance.

Tremendous SSD guarantees the latency as low as tens of microseconds, making it fit for core business systems with high requirements for data reliability and availability.

Large databases: Support medium and large relational database applications that contain tables with millions of rows, such as MySQL, Oracle, SQL Server, and MongoDB.

Large NoSQL: Support NoSQL businesses such as HBase and Cassandra.

Elasticsearch: Support low-latency ES storage.

Video service: Suitable for applications with high requirements for storage bandwidth, such as audio/video encoding and decoding, live streaming and recording playback.

Core business systems: Suitable for I/O-intensive applications and other core business systems with high requirements for data reliability.

Big data analysis: Suitable for data analysis, data mining, business intelligence, and other fields. Provide distributed processing capabilities for data at TB and PB levels.

High performance and highway data reliability: Suitable for high-load and core critical business systems. Features comprehensive data backup, snapshot, and second-level data recovery capabilities.

Cloud SSD is applicable for applications with high and medium loads, including:

Medium databases: Medium and large relational database applications, such as MySQL.

Image processing: Support data analysis and storage businesses, such as image processing.

Balanced SSD is mainly used in the following data scenarios:

Medium applications with high requirements for data reliability and standard requirements for performance, such as Web/App servers, business logical processing, KV services, as well as basic database services.

Premium Disk is mainly suitable for the following data scenarios:

Small and medium databases and Web/App servers. Provide long-term and stable I/O performance.

Scenarios that require balanced storage capacity and performance, such as enterprise office services.

Core business testing and the front and back end debugging.

Billing Description

For pricing details of cloud disks, see [Price Overview](#).

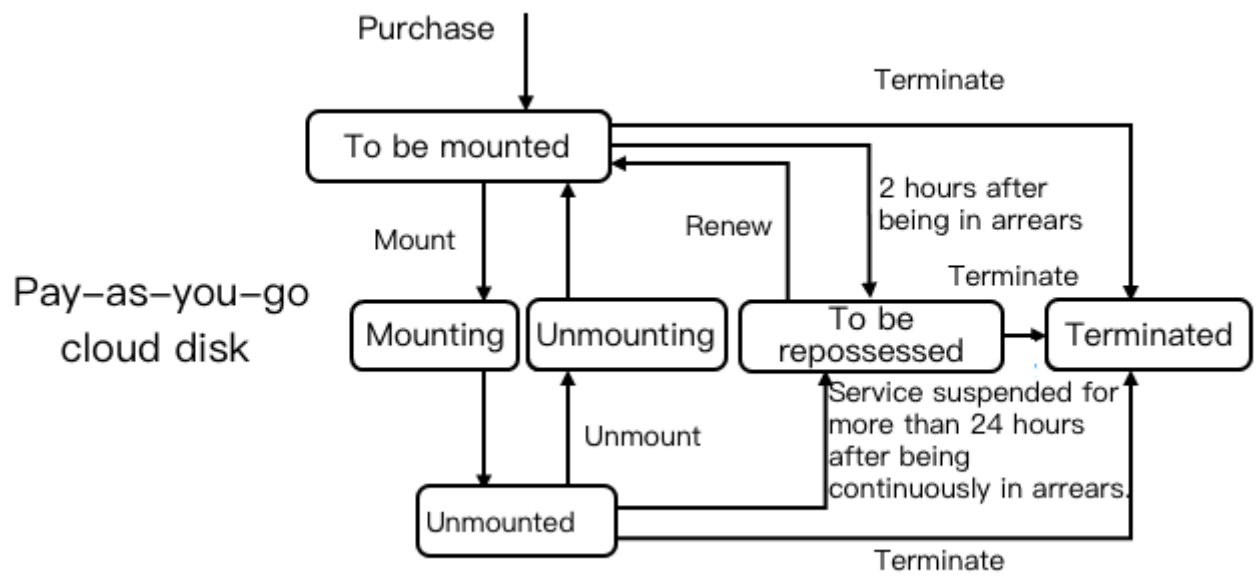
Cloud Disk Status

Last updated : 2023-12-21 14:21:23

Cloud disk has the following status:

Status	Attribute	Description
To be mounted	Stable status	The status after the cloud disk has been created and before it is mounted to a CVM.
Mounting	Interim status	When the cloud disk is being mounted, it enters the **mounting** status.
Mounted	Stable status	The status when the cloud disk has been mounted to a CVM in the same availability zone.
Unmounting	Interim status	When the cloud disk is being unmounted, it enters the **unmounting** status.
To be repossessed	Stable status	The status when a cloud disk that has not been renewed within the specified period after expiration, or a cloud disk with monthly subscription that has been manually terminated, is sent to the recycle bin after being suspended (the disk is unavailable, and can only store data) and forced to unmount.
Terminated	Stable status	The cloud disk is not renewed and retrieved before its storage time in the recycle bin expires, or the termination operation is completed. The original cloud disk no longer exists, and data has been completely erased.

The conversion relationship between cloud disk statuses is shown in the following figure:



Cloud Disk Encryption

Last updated : 2023-12-21 14:26:45

To encrypt the data stored in a cloud disk, you can enable the cloud disk encryption feature, provided by Tencent Cloud Key Management Service (KMS) instructed in [Product Overview](#).

Note:

This feature is currently in beta test. To use it, [contact us](#).

Key Management

Tencent Cloud encrypts data in your cloud disks using a data encryption key based on the standard AES-256 algorithm. When you use cloud disk encryption for the first time, the system automatically creates a customer master key (CMK) that allows you to use the cloud disk encryption feature in the corresponding region in the KMS. Only one CMK is automatically created and stored in the KMS, which is protected by strict physical and logical security controls. In each region, a unique 256-bit data key (DK) is used to encrypt the cloud disk. Snapshots created through encrypted cloud disks and encrypted cloud disks created through encrypted snapshots are all associated with this DK. The DK is protected by the key management infrastructure provided by KMS, which effectively blocks unauthorized access. The DK of a cloud disk is used only in the memory of the host where the instance resides, and is not stored in any persistent medium (including the cloud disk itself) in a plaintext form.

How it Works

When you configure your cloud disk as encrypted, the KMS encrypts the data and automatically decrypts it during the read operation. The encryption and decryption processes are performed on the host where the CVM instance resides, with minimal impact on the read and write performance of the cloud disk. To test the performance of cloud disks, refer to [Measuring Cloud Disk Performance](#).

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

Static data in the cloud disk;

Data transmitted between the cloud disk and instance (data in the operating system of the instance is not encrypted);

All snapshots created through encrypted cloud disks;

Limits

The cloud disk encryption feature is subject to the following limitations:

--	--

Limitation	Description
Cloud disk limitations	<p>All types of cloud disks can be encrypted, regardless of type of associated instance.</p> <p>Only cloud disks can be encrypted, not local disks.</p> <p>Only data disks can be encrypted, not system disks .</p> <p>An existing non-encrypted disk cannot be directly converted to an encrypted disk.</p> <p>An encrypted cloud disk cannot be converted to a non-encrypted cloud disk.</p> <p>To recognize the new capacity of an expanded encrypted cloud disk, you need to uninstall it and reattached it to the CVM.</p> <p>An encrypted cloud disk cannot be attached to an instance with local storage.</p>
Snapshots and images limitations	<p>A snapshot generated by an existing non-encrypted disk cannot be directly converted to an encrypted snapshot.</p> <p>An encrypted snapshot cannot be converted to a non-encrypted snapshot.</p> <p>An image with an encrypted snapshot cannot be shared.</p> <p>The encrypted snapshot and images created by it cannot be replicated across regions.</p>
Other Limitations	<p>The cloud disk encryption feature relies on the KMS in the same region. If you have no other operation requests, you do not need to perform additional operations in the KMS console.</p> <p>When you use the cloud disk encryption feature for the first time, you must activate KMS as instructed on the page. Otherwise, you cannot purchase the encrypted cloud disk.</p> <p>You can query the CMK created specifically by the system for cloud disk encryption in KMS console, but you cannot specify, delete, or change the CMK.</p>

Billing

Cloud disk encryption, CMK, and reads/writes of cloud disk data do not incur additional charges. When you manage the encrypted cloud disk either in the console or through an API, however, KMS is used as an API and your management operation will be counted as a KMS call in this region. You will be billed based on the number of KMS calls. For details, see [Billing Overview](#).

Management operations on an encrypted cloud disk include:

Create an encrypted cloud disk

Attach a cloud disk

Detach a cloud disk

Create a snapshot

Roll back a snapshot

Note:

Make sure you have sufficient account balance, otherwise the operation will fail.

Creating an encrypted cloud disk

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

Creating in the console

Creating from a snapshot

Creating Using an API

1. Log in to the [CBS console](#), select a region, and click **Create**.
2. In the **Purchase data disk** dialog box, select **Enable disk encryption**.

Note:

If you are using cloud disk encryption in this region for the first time, you need to get the authorization for KMS first.

3. Select the cloud disk configuration based on your actual needs and click **Ok**.

4. Once you have purchased the cloud disk, you can view encrypted cloud disks that have already been created on the [Cloud disk list](#) page.

The new encrypted cloud disk is in **to be attached** status, you can refer to [Attaching Cloud Disks](#) to attach the cloud disk to a CVM instance in the same availability zone.

You can select an encrypted snapshot to create a cloud disk. The cloud disk created in this way is encrypted automatically. For more information, see [Creating Cloud Disks Using Snapshots](#).

You can create an encrypted cloud disk by using the [CreateDisks API](#). The following two methods are supported.

Configure `Encrypt` as `true`.

Specify a `SnapshotId` for the encrypted snapshot.

Changing Data Encryption Status

To change the status of existing data in the cloud disk from non-encrypted to encrypted, we recommend you run the

`rsync` command in Linux system or the `robocopy` command in Windows system to copy the data from the non-encrypted disk to the new encrypted disk.

Enhanced SSD Performance

Last updated : 2023-12-21 14:32:54

Starting from January 2021, a single Enhanced SSD disk performance consists of **basic performance** and **extra performance**:

Basic performance varies linearly with cloud disk capacity and reaches the maximum value at the critical point.

If you have a higher performance requirement, you can enable and configure extra performance.

This document introduces basic performance and extra performance.

Overall Performance

For a single Enhanced SSD disk, its performance consists of two parts: **basic performance** and **extra performance**. The following table specifies the maximum performance, regardless of the performance proportion.

Performance Metric	Maximum Value
Random IOPS	100,000
Max throughput (MB/s)	1000

Basic Performance

For Enhanced SSD, its basic performance varies linearly with the disk capacity. The basic performance cannot be adjusted independently, which is calculated by the following formula.

Basic Performance Metric	Calculation Formula	Maximum Value
Random IOPS	$\min \{1800 + \text{capacity (GB)} \times 50, 50000\}$	50000
Max throughput (MB/s)	$\min \{120 + \text{capacity (GB)} \times 0.5, 350\}$	350

According to the formula,

If the disk capacity is 460 GB, the throughput reaches the maximum value and its basic performance is: random IOPS: 24,800; max throughput: 350 MB/s.

If the disk capacity is 964 GB, the random IOPS reaches the maximum value and its basic performance is: random IOPS: 50,000; max throughput: 350 MB/s.

Extra Performance

If you have a higher performance requirement, you can enable extra performance.

Prerequisites

The following requirements must be met to enable extra performance:

Currently, this feature is only available to **Enhanced SSD** and **ultra SSD** disks.

The extra performance can be configured only after either basic performance metric reaches its maximum value. In other words, the disk capacity must be greater than 460 GB according to the basic performance formula.

Extra performance formula

Refer to the following formula to configure the extra performance.

Extra Performance Metric	Calculation Formula	Maximum Value
Random IOPS	$\min \{\text{configured value} \times 128, 50,000\}$	50,000
Max throughput (MB/s)	$\min \{\text{configured value} \times 1, 650\}$	650

Extra performance price

For more information on extra performance prices, see [Price Overview](#).

Examples

Sample 1: An Enhanced SSD disk requires a capacity of 2,000 GB and throughput of 500 MB/s

This capacity configuration reaches the throughput limit of basic performance and needs to configure extra throughput performance: $(500-350)/1 = 150$. Therefore, you need to purchase an Enhanced SSD disk with a capacity of 2,000 GB and extra throughput performance of 150 MB/s.

Sample 2: An Enhanced SSD disk requires a capacity of 1,000 GB and IOPS of 50,000

This capacity configuration reaches the random IOPS limit of basic performance, and the required IOPS is met. Therefore, you only need to purchase an Enhanced SSD disk with a capacity of 1,000 GB.

Cloud Disk Data Backup Point

Last updated : 2025-01-15 16:46:30

Cloud Disk Data Backup Point is an automated data backup method provided by Tencent Cloud, offering fully managed data backup capabilities. A data backup point is a complete data copy of the cloud disk at the backup time, which can be used to restore historical versions of the data.

Unlike [Cloud Disk Snapshots](#), data backup points follow the lifecycle of the cloud disk and cannot exist independently. When the cloud disk is deleted, the data backup points are also automatically deleted and cannot be used to restore data.

Data Backup Point Feature

Historical Data Backup

After configuring the data backup point, Tencent Cloud will **automatically** retain a historical data backup of the cloud disk from a moment of the previous calendar day. If data anomalies occur due to viruses, intrusions, or misoperations, the automatically created historical data backup point can be used to restore the data to its state before the issue occurred.

Convert to long-term persistent snapshot

Important data backup points can be converted to snapshots for long-term backup that is independent of the CBS lifecycle.

Difference Between Data Backup Point and Snapshot

You Need to Know	Description
Cloud Disk Data Backup Point	<p>A data backup point is a historical data backup of the cloud disk automatically retained by Tencent Cloud. The Tencent Cloud platform is responsible for the automatic rotation and update of the backup.</p> <p>Advantage:</p> <p>No manual operation is required after configuration. Data backup and rotation are automatically managed by the cloud platform.</p> <p>The cost is only related to the size of the CBS and the maximum number of backup points retained. Prepaid is supported, and the cost can be estimated.</p> <p>Notes:</p> <p>The lifecycle completely follows that of the CBS. Backup data cannot be retained after the CBS is deleted.</p>

Cloud Disk Snapshots

Cloud Disk Snapshots are full data backups of CBS created manually or through automatic policies. Users need to configure their own backup generation and retention policies. For more information, see [Cloud Disk Snapshots](#). The lifecycle of snapshots is independent of the CBS, meaning backup data can be retained even if the CBS is deleted.

Notes:

Manual creation of snapshots or manual configuration of automatic policies is required, which has a certain operational threshold.

Cloud Disk Snapshots are billed on an hourly pay-as-you-go basis, which has some unpredictability.

Using Cloud Disk Data Backup Point

You can set a data backup point when purchasing a new cloud disk or set a data backup point for an existing cloud disk. For specific steps, please refer to the [Operation Guide](#).

The number of data backup points indicates the maximum number of data backups that can be retained. Tencent Cloud will **automatically** retain a historical data backup of the cloud disk with configured backup points at a moment in the **previous calendar day**. For example:

Case A: If you purchase a new cloud disk on May 2nd and set the data backup point to 1, then without any further action:

On May 3rd, the cloud disk will retain a historical data backup from any point in time on May 2nd.

On May 4th, the cloud disk will retain a historical data backup from any point in time on May 3rd, and the historical data backup from May 2nd will be automatically deleted.

This process will continue until the end of the cloud disk's lifecycle.

Case B: If you purchase a new cloud disk on May 2nd and set the data backup point to 3, then without any further action:

On May 3rd, the cloud disk will retain a historical data backup from any point in time on May 2nd.

On May 4th, the cloud disk will retain a historical data backup from any point in time on May 3rd and May 2nd.

On May 5th, the cloud disk will retain a historical data backup from any point in time on May 4th, May 3rd, and May 2nd.

On May 6th, the cloud disk will retain a historical data backup from any point in time on May 5th, May 4th, and May 3rd.

This process will continue until the end of the cloud disk's lifecycle.

Billing Instructions

Pricing information

The billing method for the data backup point quota is only related to the size of the CBS and is not affected by other factors. The billing model for the data backup point quota follows the CBS billing model. For specific prices, see the [Pricing List of Data Backup Point Quota](#).

Billing After the Adjustment

When you purchase data backup point quotas, increase or decrease the maximum number of retained data backup points, refer to the [Billing After the Adjustment](#) for specific billing rules and examples.

Notes

Data backup points are different from snapshots and are not billed according to the snapshot billing method. Data backup points will end with the lifecycle of the cloud disk and cannot retain data backups. For long-term retention, convert the data backup points to [Cloud Disk Snapshots](#).

Operation Guide

Operation Instructions	Reference Documentation
Set data backup point quota when purchasing a new CBS	Creating Cloud Disks
Increase, decrease/return the data backup point quota of existing CBS	Adjust Data Backup Point Quota
Viewing Cloud Disk Data Backup Point Information	Viewing Cloud Disk Information
Using Existing Data Backup Point to Restore Cloud Disk Data	Using Data Backup Point to Restore Cloud Disk Data
Convert Data Backup Point to Snapshot for Long-term Retention of Important Data Backup	Converting Data Backup Point into Snapshot
Deleting Data Backup Point	Delete Data Backup Point

Cloud Disk Snapshots

Snapshot Overview

Last updated : 2023-12-21 18:54:35

Overview

Real-time replica of online data

Snapshots are fully usable copies of cloud disks. When a problem occurs in a cloud disk where a snapshot has been created, you can use the snapshot to quickly restore the cloud disk to normal status. We recommend you create a snapshot for the cloud disk before making any major changes to your businesses, so that data can be quickly restored if the business changes fail.

Persistent backup at critical milestones

Snapshots can be used as persistent backups of business data to keep the business data at milestones.

Quick business deployment

Snapshots allow you to quickly clone multiple cloud disks for quick server deployment.

Applications

Snapshots provide a convenient and efficient data protection service, which can be used in the following business scenarios:

Daily data backup

You can use snapshots to regularly back up important business data to avoid data loss caused by incorrect operations, attacks, and viruses.

Quick data recovery

You can create snapshots before performing major operations, such as changing operating systems, upgrading applications, or migrating business data. If any problem occurs, you can use snapshots to restore the business data.

Application of multiple replicas of production data

You can create snapshots for the production data to provide near real-time data for applications such as data mining, report query, and development testing.

Quick environment deployment

You can create snapshots for a CVM instance to create a custom image, and use the custom image to create CVM instances for batch deployment with the same environment and less configuration time.

Billing

For more information about snapshot prices, see [Billing Overview](#) and [Price Overview](#).

Quota Limits

For more information about snapshot quota limits, see [Use Limits](#).

Snapshot Types

Manual snapshot

You can manually create a snapshot for a cloud disk at a certain point in time. This snapshot can be used to create more cloud disks with identical data, or to restore the cloud disk to the point in time when the snapshot was created. For more information, see [Creating Snapshots](#).

Scheduled snapshot

For businesses that are updated continuously, you can use scheduled snapshots to provide continuous data backups. To achieve continuous backups of cloud disk data over a certain time period, you only need to configure a backup policy and associate it with cloud disks, which significantly enhancing data security. For more information, see [Scheduled Snapshot](#).

Note

During snapshot creation, application data saved in the memory may not be persistently stored. As such, snapshots may not capture the latest and most complete cloud disk data. Refer to [Notes](#) to ensure snapshot data consistency.

Case Review

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

Customer A has never created a snapshot for the cloud disk. In May 2019, an operator performed a fio test on the cloud disk. The file system was corrupted. The data was damaged and could not be recovered.

Analysis: If customer A has created a snapshot for the cloud disk before testing, the snapshot can be used to roll back data and resume business immediately after the data damage.

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

Customer B created snapshots for multiple cloud disks, except those purchased after January 2019 for cost reasons. In June 2019, a cloud disk without snapshot protection had an unrecoverable data loss due to an accidental deletion of system-layer file data.

Analysis: If customer B has created scheduled snapshots for this cloud disk, the data can be recovered to the point in time when the last snapshot was created, thereby minimizing loss. After the incident, customer B created a snapshot for that cloud disk to enhance data protection.

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

Customer C created snapshots for all cloud disks. In May 2019, a startup exception occurred due to an incorrect operation.

Analysis: Customer C promptly restored data using the scheduled snapshot that was created two days ago, and the business remains stable.

These cases all involve data loss due to incorrect operations, but the results are different. By comparison, we can find that:

In situations where **a snapshot hasn't been created**, data is rarely recoverable when a server or cloud disk exception occurs, resulting in major loss.

In situations where **a snapshot has been created**, data can be recovered when a server or cloud disk exception occurs, minimizing loss.

We recommend regularly creating snapshots for businesses based on business types, enhancing data security and achieving low-cost, high-efficiency disaster recovery.

Others

For other questions, refer to [Snapshot FAQs](#).

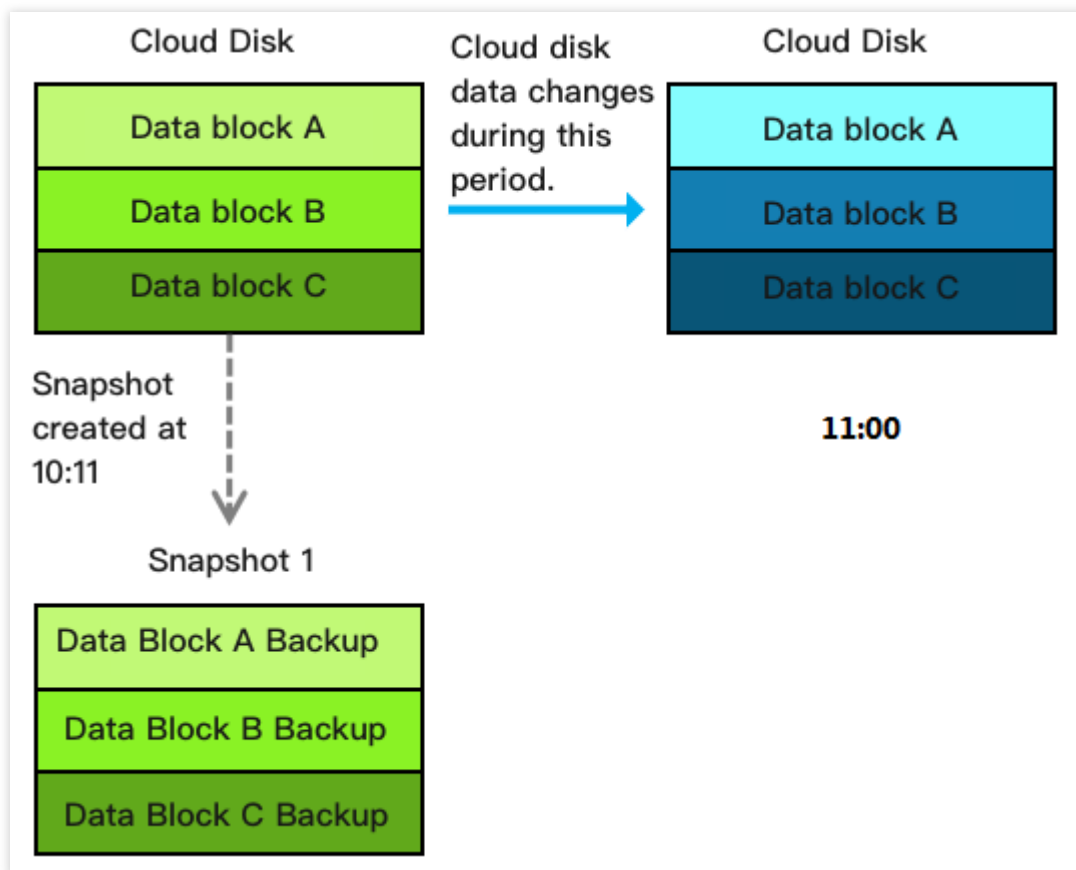
Snapshot Principles

Last updated : 2023-12-21 18:56:36

Snapshot and Source Disk

A snapshot is the data backup of a cloud disk at a certain point in time. Data writing and modification to the cloud disk do not affect snapshots that are already created. Users can use snapshots to back up cloud disk data at different points in time, which can be used for system recovery, disaster recovery, and cloud disk replication.

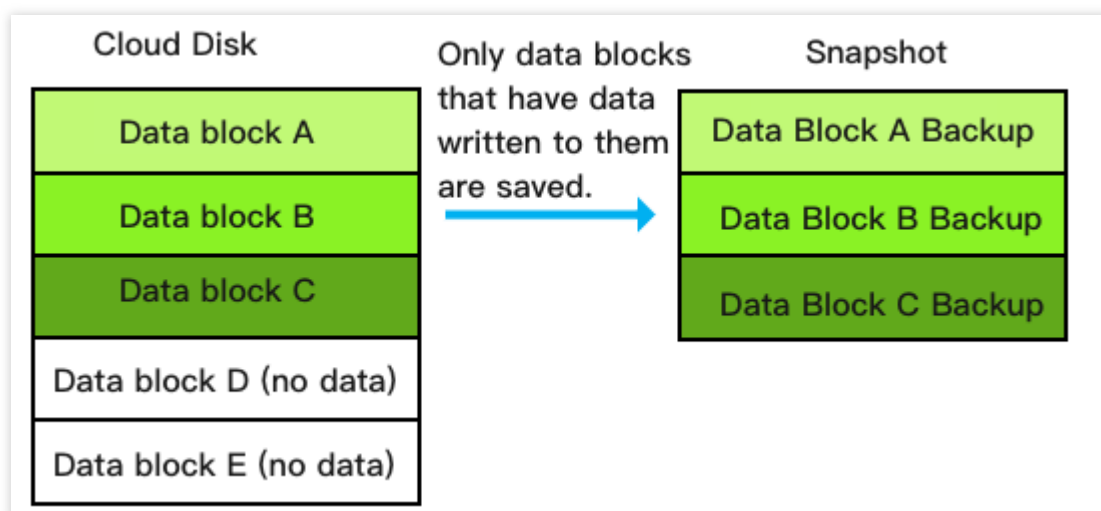
As shown in the following figure, Snapshot 1 retains data block information of the cloud disk at 10:00 (the snapshot creation time), regardless of any changes to the disk that occur after the snapshot is created.



Snapshot Size

A snapshot only saves data blocks in the cloud disk that have been written to or modified. Therefore, the size of the snapshot is smaller than the source cloud disk.

See below for more details:



Creating Incremental Snapshots

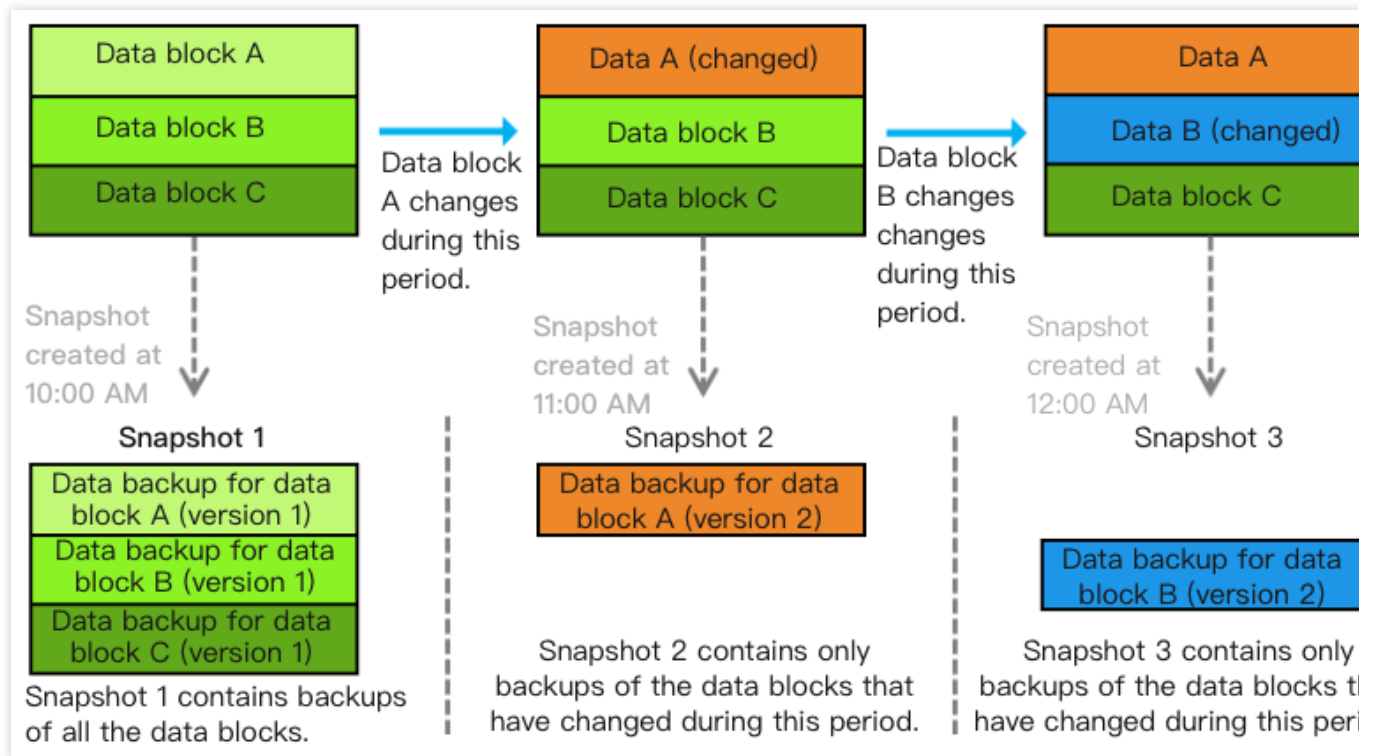
Tencent Cloud snapshots use an incremental snapshot mechanism. When you continuously create multiple snapshots of the same cloud disk, only the first snapshot is a full snapshot, and subsequent snapshots only contain data that has been modified relative to the previous snapshot (incremental snapshot). This can minimize the total storage capacity occupied when users continuously create snapshots, reducing user costs.

For example: Assume a cloud disk has three data blocks, A, B, and C. You make snapshots at 10:00, 11:00, and 12:00 respectively. Changes of data blocks on the disk between these points in time are shown in the following figure, and each snapshot should save the following data:

Snapshot 1 (initial snapshot): Contains data backups of all data blocks on the cloud disk at that time.

Snapshot 2: During the period, data block A on the cloud disk changes. Snapshot 2 only contains a backup of block A's newest data (usually called an incremental snapshot).

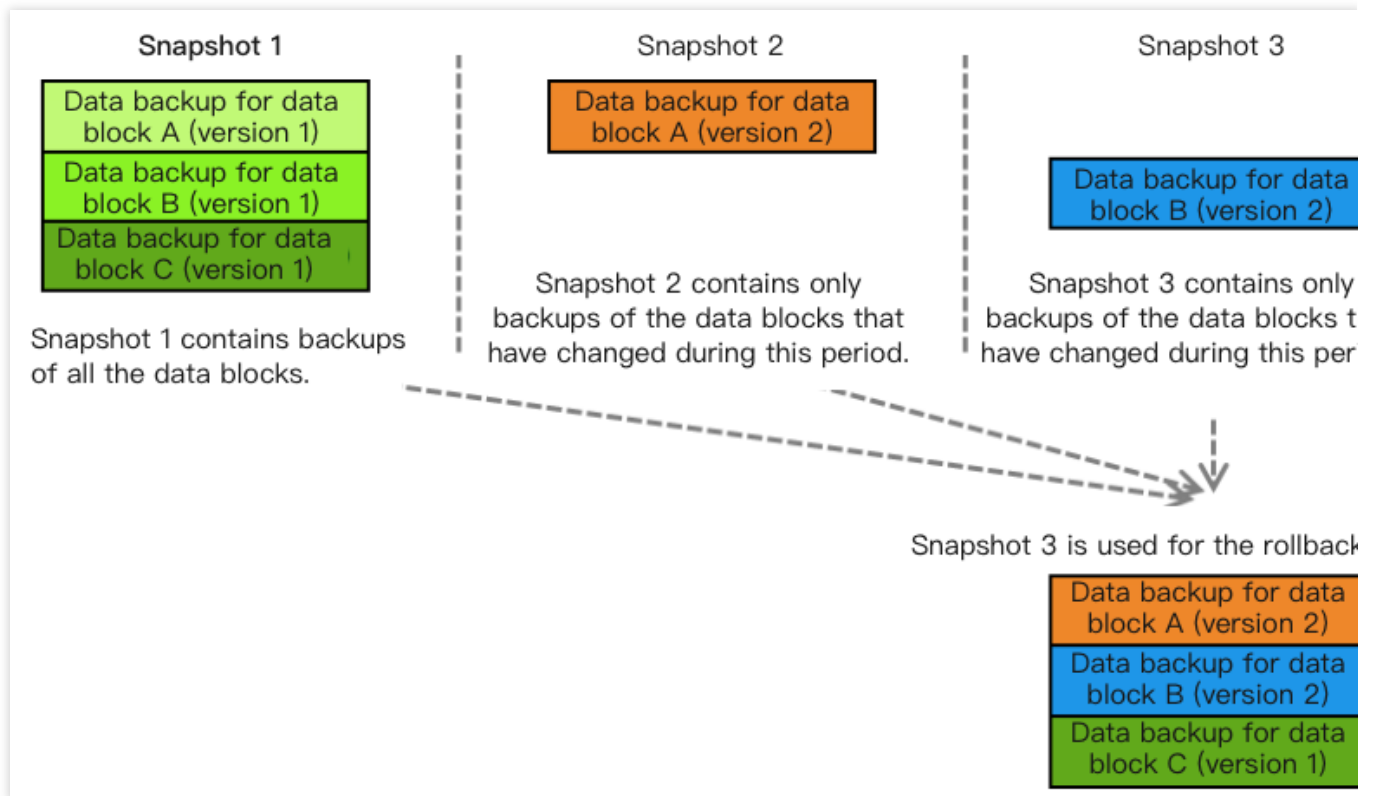
Snapshot 3: During the period, data block B on the cloud disk changes. Snapshot 3 only contains a backup of block B's newest data (usually called an incremental snapshot).



Rolling Back with Incremental Snapshots

Based on the previous example, when you use Snapshot 3 to perform data rollback, the system will merge the data in Snapshot 1, Snapshot 2, and Snapshot 3. If there is a data block in the same location, data in the newest snapshot will be taken. During final rollback, the merged data will be written to the cloud disk to be rolled back.

Incremental snapshot rollback process is as shown in the following figure:



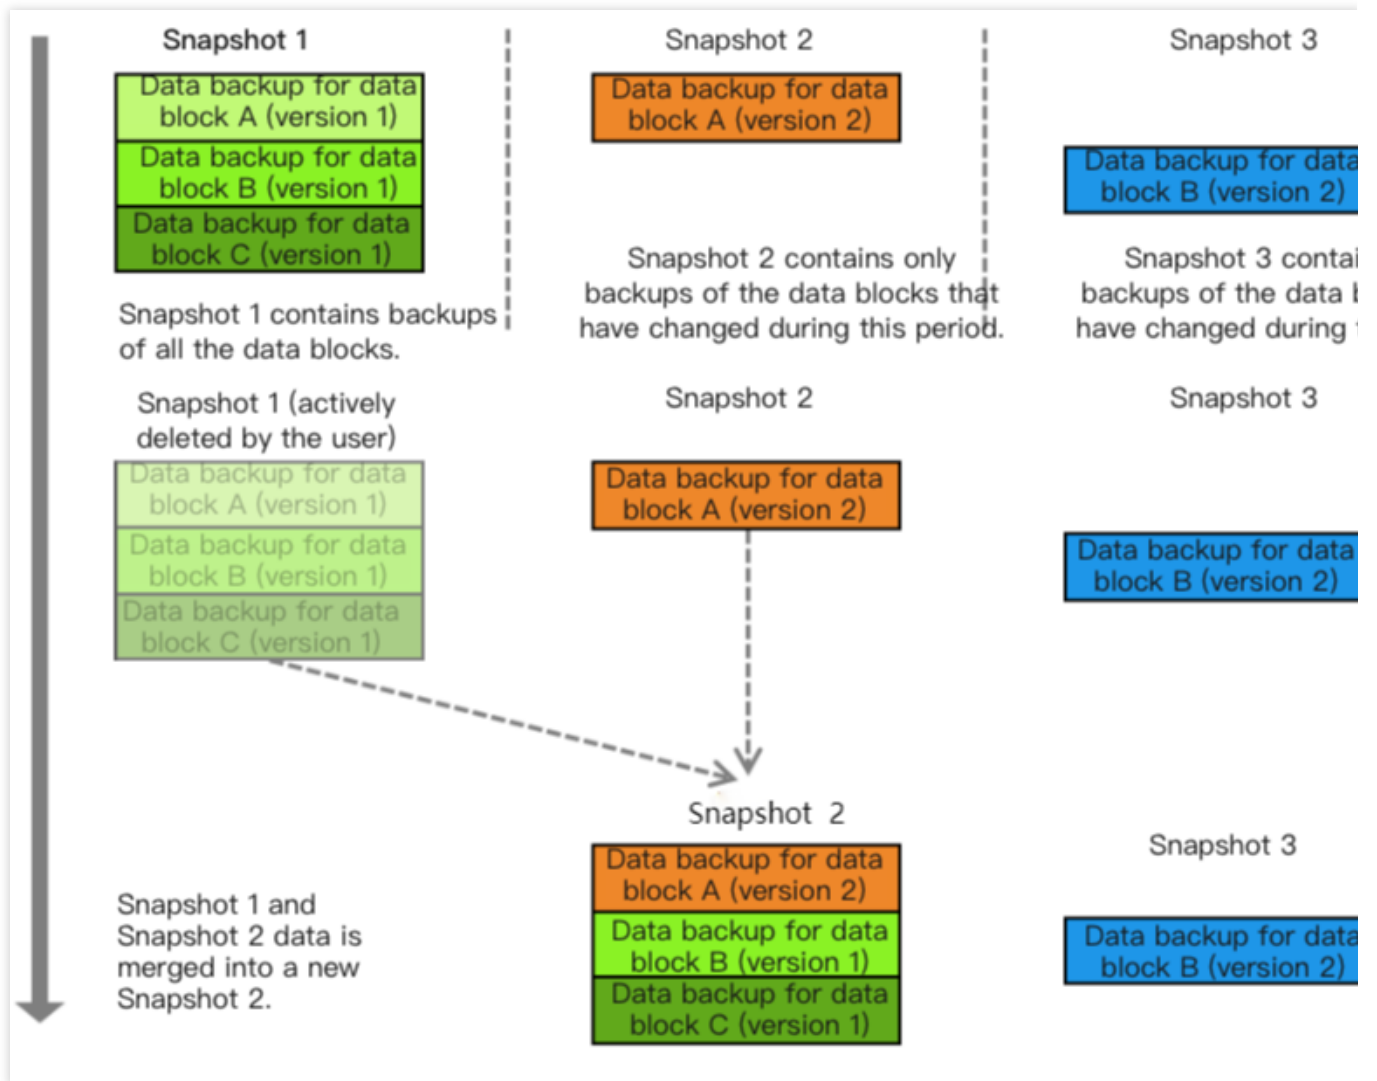
Deleting and Merging Incremental Snapshots

When a full snapshot (that is, the first snapshot) is being deleted, the system automatically merges the full snapshot with the next incremental snapshot.

When an incremental snapshot is being deleted, the system automatically merges the incremental snapshot with the next incremental snapshot. If there is no next incremental snapshot, it is directly deleted.

Based on the previous example, if you delete Snapshot 1, the system will merge Snapshot 1 and Snapshot 2, and will use Snapshot 2's data to overwrite Snapshot 1's data in the same location. After merging, Snapshot 2 is the new full snapshot.

Incremental snapshot deletion process is as shown in the following figure:



Access Management

An Overview of Access's Management

Last updated : 2023-12-21 18:59:21

If you are using multiple Tencent Cloud services such as CVM, CBS, VPC, and TencentDB that are managed by different users who share your Tencent Cloud account key, you may face the following problems.

Your key is shared by multiple users, leading to a high risk of disclosure.

You cannot control the access permissions of other users, which poses a security risk due to potential misoperations.

In this case, you can use sub-accounts to allow different users to manage different services to avoid these problems.

By default, a sub-account does not have the permission to use CVMs or CVM-related resources. Therefore, you need to create a policy to grant the required resources or permissions to the sub-account.

[Cloud Access Management \(CAM\)](#) is a set of web-based Tencent Cloud services that helps you securely manage and control access permissions to your Tencent Cloud resources. By using CAM, you can create, manage, and delete users (groups) and control who can use Tencent Cloud resources and which Tencent Cloud resources they can use through identity and policy management.

When using CAM, you can associate a policy with a user or a user group, which grants or denies them permission to use specified resources to perform specified tasks. For more information on CAM policy basics, see [Policy Syntax](#). For more information on the use of CAM policies, see [Policies](#).

If you do not need to manage the access permissions of sub-accounts to CBS resources, you can skip this section.

This will not affect your understanding and application of the remaining sections of this document.

Getting Started

A CAM policy must grant or deny the permission to one or more CBS operations. At the same time, it must specify the resources that can be operated on (which can be all resources or some resources for certain operations). A policy can also include the conditions set for the operations of the resources.

Task	Link
Learn the basic structure of a policy	Policy Syntax
Define operations in a policy	CBS Operations
Define resources in a policy	CBS Resource Paths
Restrict a policy by conditions	CBS Condition Keys
Learn the resource-level permissions supported by CBS	Resource-Level Permissions Supported by CBS
View console examples	Console Examples

Authorized resource types

Last updated : 2023-12-21 19:17:09

Resource-level permissions refer to the ability to specify which resources users are allowed to perform operations on. CBS supports resource-level permissions. That is, you can specify when users are allowed to perform some CBS operations that support resource-level permissions or which resources users are allowed to use.

The types of resources that can be authorized in Cloud Access Management (CAM) are as follows:

Resource Type	Resource Description Method in Authorization Policy
CBS APIs	<code>qcs::cvm:\$region::volume/*</code>

CBS APIs describe CBS API operations that currently support resource-level permissions as well as resources and condition keys supported by each operation. **When configuring the resource path**, you need to replace variable parameters such as `$region` and `$account` with your actual parameters. You can also use the `*` wildcard in the path. For more information, see [Console Example](#).

Caution:

CBS API operations not listed in the table do not support resource-level permissions. You can still authorize users to perform these operations, but the resource element of the policy statement must be specified as `*`.

CBS APIs

API Operation	Resource Path	Condition Key
Mount a cloud disk <code>AttachDisks</code>	<code>qcs::cvm:\$region:\$account:volume/*</code> <code>qcs::cvm:\$region:\$account:volume/\$diskId</code>	<code>cvm:region</code> <code>cvm:zone</code> <code>cvm:disk_type</code>
Create a cloud disk <code>CreateDisks</code>	<code>qcs::cvm:\$region:\$account:volume/*</code> <code>qcs::cvm:\$region:\$account:volume/\$diskId</code>	<code>cvm:region</code> <code>cvm:zone</code> <code>cvm:disk_type</code>
Query the list of cloud disk operation logs <code>DescribeDiskOperationLogs</code>	<code>qcs::cvm:\$region:\$account:volume/*</code> <code>qcs::cvm:\$region:\$account:volume/\$diskId</code>	<code>cvm:region</code> <code>cvm:zone</code> <code>cvm:disk_type</code>
Query the list of cloud disks <code>DescribeDisks</code>	<code>qcs::cvm:\$region:\$account:volume/*</code> <code>qcs::cvm:\$region:\$account:volume/\$diskId</code>	<code>cvm:region</code> <code>cvm:zone</code> <code>cvm:disk_type</code>
Unmount a cloud disk <code>DetachDisks</code>	<code>qcs::cvm:\$region:\$account:volume/*</code> <code>qcs::cvm:\$region:\$account:volume/\$diskId</code>	<code>cvm:region</code> <code>cvm:zone</code> <code>cvm:disk_type</code>

Modify the attributes of cloud disks <code>ModifyDiskAttributes</code>	<code>qcs::cvm:\$region:\$account:volume/*</code> <code>qcs::cvm:\$region:\$account:volume/\$diskId</code>	<code>cvm:region</code> <code>cvm:zone</code> <code>cvm:disk_type</code>
Modify the billing mode of a cloud disk <code>ModifyDisksChargeType</code>	<code>qcs::cvm:\$region:\$account:volume/*</code> <code>qcs::cvm:\$region:\$account:volume/\$diskId</code>	<code>cvm:region</code> <code>cvm:zone</code> <code>cvm:disk_type</code>
Modify the renewal flag of a cloud disk <code>ModifyDisksRenewFlag</code>	<code>qcs::cvm:\$region:\$account:volume/*</code> <code>qcs::cvm:\$region:\$account:volume/\$diskId</code>	<code>cvm:region</code> <code>cvm:zone</code> <code>cvm:disk_type</code>
Renew a cloud disk <code>RenewDisk</code>	<code>qcs::cvm:\$region:\$account:volume/*</code> <code>qcs::cvm:\$region:\$account:volume/\$diskId</code>	<code>cvm:region</code> <code>cvm:zone</code> <code>cvm:disk_type</code>
Expand the capacity of a cloud disk <code>ResizeDisk</code>	<code>qcs::cvm:\$region:\$account:volume/*</code> <code>qcs::cvm:\$region:\$account:volume/\$diskId</code>	<code>cvm:region</code> <code>cvm:zone</code> <code>cvm:disk_type</code>
Return a cloud disk <code>TerminateDisks</code>	<code>qcs::cvm:\$region:\$account:volume/*</code> <code>qcs::cvm:\$region:\$account:volume/\$diskId</code>	<code>cvm:region</code> <code>cvm:zone</code> <code>cvm:disk_type</code>

Authorization Policy Syntax

Last updated : 2023-12-21 19:21:50

Policy Syntax

CAM policy:

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

version: Required. It must be `"2.0"`.

statement: Detailed information of one or more permissions. Each permission is composed of a set of elements including effect, action, resource, and condition. A policy has only one statement element.

action: Required. It describes the operations to be allowed or denied, which can be APIs (described with the "name" prefix) or a feature set (a set of specific APIs described with the "permid" prefix).

resource: Required. It describes the specific data to be authorized in a six-segment format. Detailed resource definitions vary by product.

condition: Optional. It describes the conditions for the policy to take effect. A condition consists of an operator, operation key, and operation value. A condition value may contain information such as a time and IP address. Some services allow you to specify other information in conditions.

effect: Required. It describes the result returned by the statement, that is, whether the permission is allowed ("allow") or denied ("deny").

CBS Operations

In a CAM policy statement, you can specify any API operation from any service that supports CAM. For CBS, use the APIs prefixed with `name/cvm:`, for example, `name/cvm:CreateDisks` or `name/cvm:DescribeDisks`.

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 a wildcard to specify multiple operations. For example, you can specify all operations whose names begin with "Describe", as shown below.

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

To specify all operations in CVM, use the wildcard `*` as follows.

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

CBS Resource Paths

Every CAM policy statement contains the resources applicable to the policy itself. The general format of a resource path is shown below.

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

project_id: (Optional) Project information, which is only used to enable compatibility with earlier CAM logics.

service_type: Short name of a Tencent Cloud product, for example, "CVM".

region: Region information, for example, "bj".

account: Root account of a resource owner, for example, "uin/164256472".

resource: Specific resources of a product, for example, "volume/diskid1" or "volume/*".

You can specify a CBS resource in the statement, for example, "disk-abcdefg", as shown below.

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

You can also use the wildcard `*` to specify all CBS resources under an account, as shown below.

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

To specify all resources, or if an API operation does not support resource-level permission control, you can use the wildcard `*` in the resource element, as shown below.

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

To specify multiple resources in one statement, separate them with commas. In the following example, two resources are specified.

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

CBS Condition Keys(id:tiaojianmiyao)

In a policy statement, you can choose to specify the conditions for the policy to take effect. Each condition contains one or more key-value pairs. Condition keys are case-insensitive.

If you specify multiple conditions or keys in one condition, the condition is evaluated with the "AND" logical operator.

If you specify a key with multiple values in one condition, the condition is evaluated with the "OR" logical operator. The permission can be granted only after all conditions are met.

The following table describes the CBS condition keys that are used for specific services.

Condition Key	Reference Type	Key-Value Pair
cvm:region	String	cvm:region= region where region indicates a region (for example, "ap-guangzhou").
cvm:disk_type	String	cvm_disk_type= disk_type where disk_type indicates a disk type (for example, "CLOUD_PREMIUM").

Related Products

Last updated : 2023-12-21 19:24:32

Product Name	Relationship to CBS
Cloud Virtual Machine	Cloud Block Storage can act as system disks and data disks for CVMs, providing block storage devices that can be elastically scaled.
GPU Cloud Computing	SSD and Premium Cloud Disk can act as system disks and data disks for GPU Cloud Computing, providing block storage devices that can be easily scaled.
Tencent Cloud Observability Platform	You can use Tencent Cloud Observability Platform to view the performance indicators of CBS, including the disk read IOPS, disk read traffic, disk write IOPS, disk write traffic, disk I/O wait time, disk I/O service time, and disk I/O utilization ratio.
Tags	Tags are used to identify cloud resources. They can be used to categorize and search for CBS resources.

Predecessor

Last updated : 2023-12-21 19:25:26

HDD Cloud Disk

Note:

The HDD cloud disk is phased out currently. The following guide is only for reference when you are using the HDD cloud disk.

HDD cloud disks are the first generation of cloud disks provided by Tencent Cloud. It is suitable for business scenarios with low I/O loads where data is not accessed frequently. It uses magnetic storage, and a three-copy distributed mechanism to implement highly reliable data storage.

Performance metrics

Metrics	HDD Cloud Disk
Random IOPS	Offers the capacity options of 10 GB - 16000 GB, and hundreds of random IOPs choices.
Throughput (MB/s)	Supports the I/O throughputs up to 50 MB/s.
Latency	-

Use cases

For scenarios with infrequent requests, such as system logs, enterprise work files, data warehouses, small-sized blogs, and BBS, we recommend using HDD cloud disks to reduce costs.