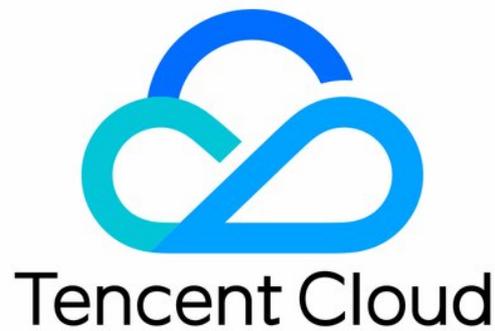


Cloud Block Storage Product Introduction



Copyright Notice

©2013–2023 Tencent Cloud. All rights reserved.

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

Trademark Notice



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

Service Notice

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

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

Contact Us

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

Contents

Product Introduction

Overview

Strengths

Scenarios

Regions and AZs

Cloud Disk Instances

Cloud Disk Types

Cloud Disk Status

Cloud Disk Encryption

Cloud Disk Data Backup Point

Enhanced SSD Performance

Cloud Disk Snapshots

Snapshot Overview

Cloud Access Management

CAM Overview

Authorizable Resource Type

Authorization Policy Syntax

Relevant Products

Predecessor

Product Introduction

Overview

Last updated: 2023-09-19 20:43:29

You can understand the concept, application scenarios, and functions of Tencent Cloud CBS through the following video:

[Watch video](#)

What is Tencent Cloud Block Storage?

Cloud Block Storage (CBS) is a highly available, highly reliable, low-cost, and customizable block storage device that can be used as an independent, scalable disk for cloud servers, providing efficient and reliable [storage](#) capabilities for cloud server instances.

Cloud Block Storage provides persistent storage capabilities at the data block level, typically serving as the primary storage device for data that requires frequent and granular updates, such as file systems and databases. It is characterized by its high availability, reliability, and performance.

Cloud Block Storage employs a three-copy distributed mechanism, backing up your data on different physical machines to avoid data loss caused by single-point failures, thereby enhancing data reliability.

You can easily purchase, adjust, and manage your CBS devices through the console, and create storage spaces larger than a single CBS by building file systems. Depending on their lifecycle, CBS can be categorized as:

- Non-elastic cloud disks

The lifecycle of a non-elastic cloud disk is completely tied to the cloud server. It is purchased alongside the cloud server and used as a **system disk**, with **no support for mounting or unmounting**.

- Elastic cloud disks

The lifecycle of an elastic cloud disk is independent of the cloud server. It can be purchased separately and then manually attached to a cloud server, or it can be purchased along with a cloud server and automatically attached to it, serving as a **data disk**. Elastic cloud disks support mounting or unmounting on a cloud server in the same availability zone at any time.

You can attach multiple elastic cloud disks to the same cloud server, or you can unmount an elastic cloud disk from cloud server A and then mount it to cloud server B.

Tencent Cloud imposes certain restrictions on the quota of CBS for users. For more details, please refer to [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.
- You don't need to purchase extra storage capacity while purchasing a CVM. You can purchase cloud disks later when it's necessary.
- When you need to transfer data from one CVM to another, just detach the cloud data disk from the source CVM and attach it to the target CVM.
- Use multiple cloud disks to form a Logical Volume Manager (LVM) to go beyond the physical limit of a single cloud disk.
- Use multiple cloud disks to form a Redundant Array of Independent Disks (RAID) configuration to go beyond the I/O performance limit of a single cloud disk.

Features

Tencent Cloud offers a variety of persistent storage devices. You have the flexibility to choose the type of hard drive and perform operations such as storing files and setting up databases on the hard drive.

- Five types of cloud disks are available for selection: Premium Cloud Disk, Balanced SSD, 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 cloud server to build a large-capacity file system.
- Elastic Expansion: You can scale up your cloud block storage at any time, with a single disk supporting up to 32TB.
- Snapshot Backup: Supports the creation of snapshots and snapshot rollbacks, enabling timely backup and recovery of critical data. It also supports the creation of hard drives using snapshots, facilitating rapid business deployment.

Strengths

Last updated: 2023-09-19 17:28:19

Reliable

Cloud Block Storage (CBS) employs a three-copy distributed mechanism, only returning a successful write response once the data has been written into all three copies. The background data replication mechanism can swiftly create a new copy through data migration if any copy fails, ensuring three copies are available at all times. This provides a secure and reliable data storage service. The data is stored across different racks, achieving a reliability of 99.9999999%.

Elasticity

You have the freedom to configure storage capacity and scale up as needed without interrupting your operations. The maximum capacity of a cloud disk is 32TB, and a single cloud server can mount up to 20 elastic cloud disks as data disks. This effectively addresses big data processing scenarios involving TB/PB level data, meets the requirements of large-capacity file systems, and is suitable for big data, data warehousing, log processing, and other scenarios.

High Performance

The high-performance cloud disk uses a Cache mechanism to meet regular business needs. The SSD cloud disk adopts the NVMe standard SSD, with a single disk providing 26,000 random read/write IOPS (Input/Output Operations per Second), making it suitable for scenarios with extreme I/O requirements.

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 back up data for your cloud disk at any time by creating a snapshot, and you can also rapidly deploy your business by creating a cloud disk using the snapshot file.

Scenarios

Last updated: 2023-09-21 15:32:44

Typical application scenarios for CBS

Delocalization

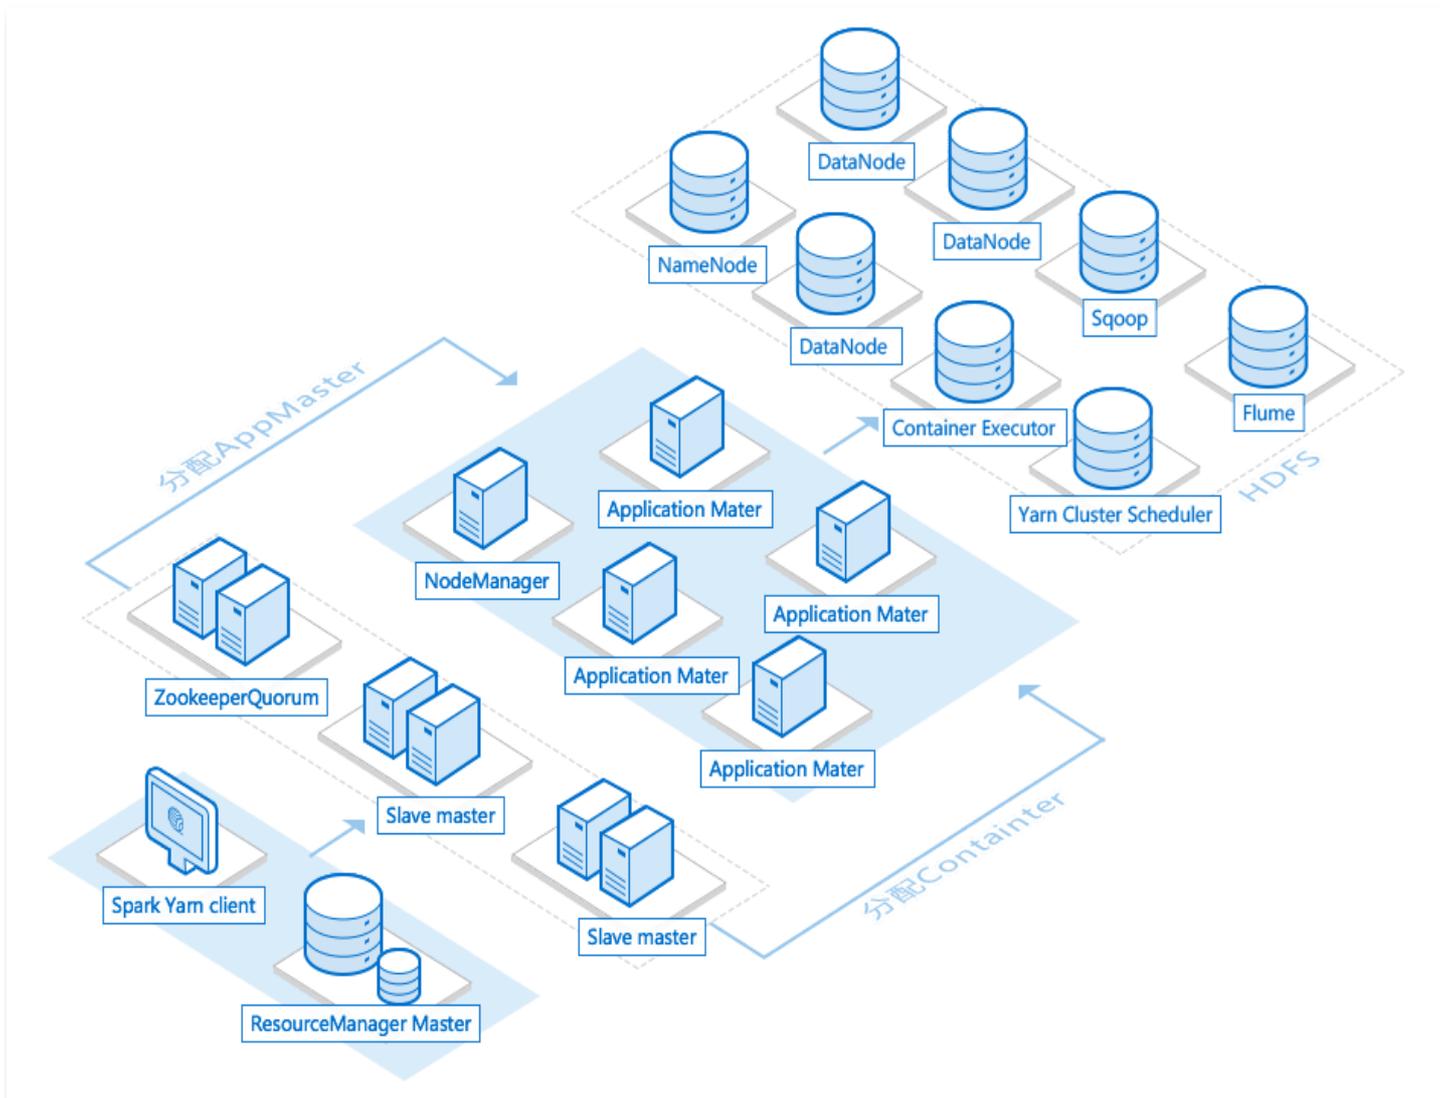
- **High-performance, highly reliable data storage:** CBS efficiently supports hot migration of cloud servers, preemptively avoiding business interruptions caused by physical failures. CBS provides three copies of data redundancy, complete data backup, snapshot, and second-level data recovery capabilities. CBS is suitable for high-load, core critical business systems.
- **Elastic scaling:** CBS can be freely attached and detached within the same availability zone without shutting down/restarting the cloud server. The capacity of CBS can be elastically configured and scaled as needed.

Massive data analysis

The typical Spark-HDFS offline data analysis framework performs sequential I/O for RDD read/write and shuffle write, while only shuffle read I/O is random I/O, with sequential I/O accounting for up to 95%. CBS's multi-threaded concurrent throughput performance is excellent, efficiently supporting Hadoop-Mapreduce, HDFS, Spark, and TB/PB level data offline processing capabilities.

In multi-disk concurrent scenarios, the throughput performance of a single HDFS cluster can reach 1GB/s.

Large enterprises such as Xiaohongshu, Giant Network, Ele.me, Yoho!BUY, Weipiao, and others have extensively carried out data analysis, mining, business intelligence, and other big data practices on CBS.



Deployment Environment: Five cloud servers with 12 cores and 40GB RAM each, simulating offline data analysis of 1.5TB data volume.

Performance Testing:

- Each CVM has one 1TB HDD cloud disk mounted. Five HDD cloud disks provide a read speed of 500MB/s, allowing data to be read to memory in 50 minutes.
- Each CVM has one 1TB SSD cloud disk mounted, allowing data to be read to memory in 25 minutes.

Core databases

SSD Cloud Disk is suitable for scenarios that require high I/O performance and high data reliability. It is particularly suitable for medium to large relational database applications such as PostgreSQL, MySQL, Oracle, SQL Server, core business systems with high I/O intensity and high data reliability requirements, and medium to large development and testing environments with high data reliability requirements.

Balancing data reliability and high-performance, SSD Cloud Disk has provided reliable

support for large enterprises such as Blade & Soul, Ask Tao, Yoho!BUY, WePiao, and Xiaohongshu.

Deployment Environment: Four cloud servers each with 4 cores and 8GB RAM, each mounted with an 800GB SSD CBS, running MySQL version 5.5.42.

Performance Test: Using sysbench to simulate OLTP performance test, with a test dataset of 10 million records, TPS can reach 1616, and QPS can reach 29,000. A single disk is sufficient to support tens of thousands of online transactions per second.

Regions and AZs

Last updated: 2023-09-19 17:37:06

Regions

A region refers to the geographical location of a physical data center. Tencent Cloud regions are completely isolated from each other, ensuring maximum stability and fault tolerance between different regions. To minimize access latency and enhance download speed, it is recommended to select the region closest to your customers.

The following are the characteristics of a region:

- The networks of different regions are entirely isolated from each other, and by default, Tencent Cloud services in different regions **cannot communicate with each other over the private network**.
- Tencent Cloud services across different regions can access the Internet via [public network services](#). Services within a private network can also communicate through Tencent Cloud's [Cloud Connect Network](#) via Tencent Cloud's high-speed interconnection network, providing a more stable and faster connection than Internet access.
- [Cloud Load Balancer](#) currently supports traffic forwarding within the same region by default, binding to CVMs in the local region. If the [Cross-Region Binding Load Balancer](#) feature is enabled, it can support binding CVMs across different regions to the load balancer.

Special Note for Shenzhen/Shanghai Finance Zones:

These zones are tailor-made for compliance with regulatory requirements in the finance industry, featuring high security and isolation. Currently, they provide services such as Cloud Virtual Machines, Cloud Block Storage, finance databases, Redis storage, and face recognition. Verified financial industry customers can submit a ticket to apply for use of these zones. For more information, please refer to [Introduction to Finance Cloud](#).

Availability Zones

An availability zone (AZ) is a physical data center within the same region of Tencent Cloud, with independent power supply and network. The objective is to ensure that failures are isolated between AZs (except for major disasters or power failures), preventing failure propagation and ensuring continuous online services for users. By launching instances within an independent AZ, users can protect their applications from being affected by a single point of failure.

The following are the characteristics of an availability zone:

- Under the same Tencent Cloud account, services in different availability zones but within the same [VPC \(Virtual Private Cloud\)](#) in the same region can communicate with each other

over the private network and can be accessed directly via [private network services](#) .

- 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

Regions	Availability Zones
South China (Guangzhou) ap-guangzhou	Guangzhou Zone 1 (Sold out) 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
South China (Shenzhen Finance) ap-shenzhen-fsi	Shenzhen Finance Zone 1 (available exclusively for financial institutions and enterprises through online consultation application) ap-shenzhen-fsi-1
	Shenzhen Finance Zone II (exclusively available to financial institutions and enterprises via online consultation for application) ap-shenzhen-fsi-2
	Shenzhen Finance Zone 3 (exclusively available to financial institutions and enterprises, which can apply for access via online consultation) ap-shenzhen-fsi-3
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 (Shanghai Finance) ap-shanghai-fsi	East China (Shanghai Finance) – Available exclusively for financial institutions and enterprises through online consultation application. ap-shanghai-fsi-1
	Shanghai Finance Zone 2 (Exclusive to financial institutions and enterprises, activation upon online consultation only) ap-shanghai-fsi-2
	ap-shanghai-fsi-2 (Exclusive to financial institutions and enterprises, activation upon online consultation only) ap-shanghai-fsi-3
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
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 and Taiwan, China (Hong Kong) ap-hongkong	Hong Kong Zone 1 (Nodes in Hong Kong, China can cover services in Hong Kong/Macao/Taiwan regions) (Sold out) ap-hongkong-1
	Hong Kong Zone 2 (Hong Kong nodes cover services in the China regions of Hong Kong, Macao, and Taiwan) ap-hongkong-2
	Hong Kong Zone 3 (Hong Kong nodes cover services in the China regions of Hong Kong, Macao, and Taiwan) ap-hongkong-3

Note

The product is currently in beta testing for the regions of Jinan, Hangzhou, Fuzhou, Wuhan, Changsha, and Shijiazhuang. If you wish to use it, please contact your account manager or submit a ticket to request access.

Other Countries and Regions

Regions	Availability Zones
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 (Jakarta nodes cover services in Southeast Asia) ap-jakarta-2
Northeast Asia (Seoul) ap-seoul	Seoul Zone 1 (Seoul nodes cover services in Northeast Asia) ap-seoul-1
	Seoul Zone 2 (Seoul nodes cover services in Northeast Asia) ap-seoul-2
Northeast Asia (Tokyo) ap-tokyo	Tokyo Zone 1 (Tokyo nodes cover services in Northeast Asia) ap-tokyo-1
	Tokyo Zone 2 (Tokyo node AZs cover services in Northeast Asia) ap-tokyo-2
South Asia (Mumbai) ap-mumbai	Mumbai Zone 1 (Mumbai nodes cover services in South Asia) ap-mumbai-1
	Mumbai Zone 2 (Mumbai nodes cover services in South Asia) ap-mumbai-2
Southeast Asia (Bangkok) ap-bangkok	Bangkok Zone 1 (Bangkok nodes cover services in Southeast Asia) ap-bangkok-1
	Bangkok Zone 2 (Bangkok nodes cover services in

	Southeast Asia) ap-bangkok-2
North America (Toronto) na-toronto	Toronto Zone 1 (Toronto nodes cover services in North America) na-toronto-1
South America (Sao Paulo) sa-saopaulo	São Paulo Zone 1 (São Paulo nodes cover services in South America) sa-saopaulo-1
US West (Silicon Valley) na-siliconvalley	Silicon Valley Zone 1 (Silicon Valley nodes cover services in West US) na-siliconvalley-1
	Silicon Valley Zone 2 (Silicon Valley nodes cover services in West US) na-siliconvalley-2
US East (Virginia) na-ashburn	Virginia Zone 1 (Virginia nodes cover services in East US) na-ashburn-1
	Virginia Zone 2 (Virginia nodes cover services in East US) na-ashburn-2
Europe (Frankfurt) eu-frankfurt	Frankfurt Zone 1 (Frankfurt nodes cover services in Europe) eu-frankfurt-1
	Frankfurt Zone 2 (Frankfurt nodes cover services in Europe) eu-frankfurt-2

Selection of Regions and AZs

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

- Restrictions on the attachment of cloud disks. Cloud disks can only be attached to a CVM in the same availability zone.
- Consider the region where the CVM that requires a cloud disk is located, as well as your location and the geographical location of your target users. When purchasing cloud services, it is recommended to select the region closest to your customers to minimize access latency and enhance access speed.
- The relationship between CVMs that require cloud disks and other Tencent Cloud

services. It is recommended that the selected services are all located in the same region and availability zone, enabling them to communicate with each other over the private network, thereby reducing access latency and enhancing access speed.

- Considerations for high availability and disaster recovery: Even in scenarios with only one VPC, it is recommended to deploy services across different availability zones at the very least. This ensures fault isolation between availability zones and facilitates cross-zone disaster recovery.
- There may be network latency among different availability zones. We recommend that you assess your business requirements and find the optimal balance between high availability and low latency.

Related Actions

Your actions such as using and viewing cloud disks are distinguished by region attributes. If you need to conveniently migrate data and services to other regions, or build a cross-region disaster recovery system, you can replicate snapshots to other regions. For more information, please refer to [Cross-Region Snapshot Replication](#).

Cloud Disk Instances

Cloud Disk Types

Last updated: 2023-09-19 17:40:10

Cloud Block Storage is a highly available, highly reliable, low-cost, and customizable network block device that can be used as a standalone scalable disk for cloud servers. It provides data storage at the data block level and employs a 3-copy distributed mechanism to ensure data reliability for cloud servers. Cloud Block Storage offers five types of disks: **premium cloud disk**, **balanced SSD cloud disk**, **SSD cloud disk**, **enhanced SSD cloud disk**, and **ultra-fast SSD cloud disk**. Each disk type has distinct performance, features, and pricing, allowing you to choose based on your application requirements.

Supports and Limits

- Currently, Enhanced SSD and ulTra SSD are only available in certain availability zones. They will be supported in more availability zones.
- Enhanced SSD is only supported when attached to S5, M5, SA2, and later generation models created after August 1, 2020. **Performance cannot be guaranteed when mounted on older generation instances.**
- **ulTra SSD can only be purchased together with Standard Storage Optimized S5se instances. Other types of cloud servers do not support ulTra SSD.**
- Enhanced SSD and ulTra SSD cannot be used as the system disk.
- Enhanced SSD and ulTra SSD cannot be encrypted.
- Enhanced SSD and ulTra SSD cannot be upgraded from other disk types.

Overview

- **Premium Cloud Disk**
Tencent Cloud Premium Cloud Disk is a hybrid storage type that employs the Cache mechanism to provide high-performance storage capabilities similar to solid-state storage while ensuring data reliability through a three-copy distributed mechanism. Premium Cloud Disk is suitable for small and medium-sized applications with high data reliability requirements and moderate performance demands, such as Web/App servers, business logic processing, and small to medium-sized website construction.
- **Balanced SSD**
Balanced SSD is an entry-level all-flash block storage product provided by Tencent Cloud. 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.

- **SSD Cloud Disk**

SSD Cloud Disk is an all-flash storage type provided by Tencent Cloud based on NVMe SSD storage media. It uses a three-copy distributed mechanism, offering low-latency, high random IOPS, and high-throughput I/O capabilities, with data security up to 99.9999999%. SSD Cloud Disk is suitable for scenarios with high I/O performance requirements.

- **Enhanced SSD**

Enhanced SSD is designed by Tencent Cloud based on the latest generation storage engine, utilizing full NVMe SSD storage media and the most recent network infrastructure. It employs a 3-copy distributed mechanism, providing low-latency, high random IOPS, high-throughput I/O capabilities, and up to 99.9999999% data security. Enhanced SSD is suitable for I/O-intensive scenarios with high latency requirements, such as large databases and NoSQL. Uniquely, Enhanced SSD supports separate configuration of performance and capacity, allowing you to tailor storage capacity and performance according to your business needs.

- **ulTra SSD**

ulTra SSD is a high-performance distributed storage engine developed by Tencent Cloud, combined with high-speed network infrastructure and the latest generation of storage hardware, providing long-term stable and ultra-low latency reliable performance. It is ideal for I/O-intensive and throughput-intensive workloads requiring extremely low latency, such as large MySQL, HBase, and Cassandra databases, etcd and rocksdb key-value storage, Elasticsearch log retrieval, video processing, and real-time high-bandwidth live streaming. It excels in critical transaction workloads, core database operations, large OLTP tasks, and video processing scenarios. ulTra SSD uniquely supports separate configuration of performance and capacity, allowing you to customize storage capacity and performance based on your business needs.

Performance Metrics

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

Performance Metrics	ulTra SSD	Enhanced SSD	SSD cloud disk	Balanced SSD	Premium Cloud Disk
Max size (GB)	32000	32000	32000	32000	32000
Max	Up to	Up to 100000 after	26000	10000	6000

IOPS of a single disk	1,000,000 after stacking extra performance	stacking extra performance			
Random IOPS Performance Calculation Formula	<p>Basic Performance: Random IOPS = $\text{Min}\{4000 + \text{Capacity (GiB)} \times 100, 50000\}$</p> <p>Additional Performance: Max IOPS = $\text{min}\{\text{Additional Performance Value} \times 128, 950000\}$</p>	<p>Basic Performance: Random IOPS = $\text{Min}\{1800 + \text{Capacity (GiB)} \times 50, 50000\}$</p> <p>Additional Performance: Maximum IOPS = $\text{Min}\{\text{Additional Performance Value} \times 128, 50000\}$</p> <p>For more details, please refer to Performance Description of Enhanced SSD.</p>	Random IOPS = $\text{Min}\{1800 + \text{Capacity (GiB)} \times 30, 26000\}$	Random IOPS = $\text{Min}\{1800 + \text{Capacity (GiB)} \times 15, 10000\}$	Random IOPS = $\text{Min}\{1800 + \text{Capacity (GiB)} \times 8, 6000\}$
Max throughput (MB/s)	Up to 4000MB/s with extra performance	Up to 1,000 MB/s with extra performance	260MB/s	190MB/s	150MB/s
Throughput performance (MB/s)	<p>Baseline Performance: Throughput = $\text{Min}\{120 + \text{Capacity (GiB)} \times 0.5, 350\}$</p> <p>Additional Performance: Throughput = $\text{Min}\{\text{Additional Performance Value} \times 1, 650\}$</p>	<p>Baseline Performance: Throughput = $\text{Min}\{120 + \text{Capacity (GiB)} \times 0.5, 350\}$</p> <p>Additional Performance: Throughput = $\text{Min}\{\text{Additional Performance Value} \times 1, 650\}$</p> <p>For more details, please refer to Performance Description of Enhanced SSD.</p>	Throughput = $\text{Min}\{120 + \text{Capacity (GiB)} \times 0.2, 260\}$	Throughput = $\text{Min}\{100 + \text{Capacity (GiB)} \times 0.2, 190\}$	Throughput = $\text{Min}\{100 + \text{Capacity (GiB)} \times 0.15, 150\}$

	e Value x 1, 3650}				
One-Way Random Read/Write Latency	0.1ms – 0.5ms	0.2ms – 1ms	0.5ms – 3ms	0.5ms – 3ms	0.8ms – 5ms
Note	ulTra SSD can only be purchased together with Standard Storage Optimized S5se instances. It cannot be used with other types of cloud servers, nor can it be purchased separately.	<ul style="list-style-type: none"> The performance of Enhanced SSD is only guaranteed when it's attached to S5, M5, and SA2 models and all later generation models. The maximum IOPS and throughput performance of Enhanced SSD is limited by the instance type it is attached to. Currently, only the latest generation S6 and SA3 instances can achieve the maximum performance per disk. The maximum storage performance metrics for different 	–	–	–

		instance types will be continuously updated, so please stay tuned.		
--	--	--	--	--

Note

- The main difference among cloud disks is the I/O performance.
- For performance metrics, the maximum IOPS can be tested with a 4KiB IO size, and the maximum throughput can be tested with a 256KiB IO size. For specific testing methods, please refer to [Measuring performance of cloud disks](#).

Scenarios

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

- High performance and high data reliability: Suitable for high-load, mission-critical business systems. It offers triple data redundancy and comprehensive data backup, snapshot, and near-instant data recovery capabilities.
- Mid and large-sized databases: support mid and large-sized relational database applications that contain tables with millions of rows, such as MySQL, Oracle, SQL Server, and MongoDB.
- Large-sized NoSQL: satisfy storage performance requirements of 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.

ultra SSD is suitable for latency-sensitive scenarios with extreme latency requirements, such as:

- Key-Value Storage: Such as RocksDB, container etcd, etc., key-value storage often writes to disk in a serial I/O mode, with very strict requirements for the speed and latency of each I/O operation. The overall performance of the system is determined by the latency of a single path. ultra SSD provides latency as low as tens of microseconds, making it highly suitable for core business systems that require 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-sized NoSQL: satisfy storage performance requirements of 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 high data reliability: Suitable for high-load, mission-critical business systems. It offers triple data redundancy and comprehensive data backup, snapshot, and near-instant data recovery capabilities.

SSD cloud disks are more suitable for medium to large workloads, such as:

- Medium databases: Medium and large relational database applications, such as MySQL.
- Image processing: Support data analysis and storage businesses, such as image processing.

Common use cases for Balanced SSD Cloud Disks include:

Suitable for medium-sized applications with high data reliability requirements and moderate performance demands, such as Web/App servers, business logic processing, KV services, and basic database services.

Primary use cases for premium cloud disks include 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.
- Satisfy I/O requirements for core business testing and developing joint testing environments.

Billing Description

For detailed cloud disk costs, please see [Cloud Disk Pricing Overview](#).

Cloud Disk Status

Last updated: 2023-09-26 11:31:29

Cloud disk has the following status:

Status	Properties	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.
Detaching	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.

Cloud Disk Encryption

Last updated: 2023-09-19 17:43:02

Should your business necessitate the encryption of data stored on cloud disks due to security or compliance requirements, you can enable the cloud disk encryption feature. Utilizing the infrastructure provided by [Tencent Cloud Key Management Service \(KMS\)](#), you can effectively safeguard the privacy of your data.

Note

The cloud disk encryption feature is currently in the trial phase. If required, please apply through [Online Support](#).

Key Management

Tencent Cloud employs the industry-standard AES-256 algorithm, using data keys to encrypt your cloud disk data. Upon first use of the encrypted cloud disk, the system automatically creates a Customer Master Key (CMK) in the corresponding region of KMS specifically for cloud disk encryption. There is only one automatically created key, which is stored in the key management service that is stringently protected by physical and logical security controls. Each encrypted cloud disk in a region uses a unique 256-bit data key (DK) for encryption in the corresponding region. The snapshots created from the encrypted cloud disk, as well as the encrypted cloud disks created using encrypted snapshots, are all associated with this key. This key is protected by the key management infrastructure provided by KMS, effectively preventing unauthorized access. The data key (DK) of the cloud disk is used only in the memory of the host to which the instance belongs and will not be stored in plaintext on any persistent medium, including the cloud disk itself.

How does it work

When you set the cloud disk as an encrypted cloud disk, KMS encrypts the data and automatically decrypts it when the data is read. The encryption and decryption process runs on the host where the cloud server instance is located, having virtually no impact on the read-write performance of the cloud disk. You can refer to [Measuring Performance of Cloud Disks](#) for cloud disk performance testing.

After creating an encrypted cloud disk and attaching it to an instance, the system will encrypt 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;

Usage Limits

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

Item	Note
Cloud disk limitations	<ul style="list-style-type: none">• Cloud disk encryption supports all cloud disk types and instance types.• Only cloud disks can be encrypted, not local disks.• Only data disks can be encrypted, not system disks .• An existing non-encrypted cloud disk cannot be directly converted into an encrypted cloud disk.• An encrypted cloud disk cannot be converted to a non-encrypted cloud disk.• To recognize the new capacity of an expanded encrypted cloud disk, you need to uninstall it and reattached it to the CVM.• An encrypted cloud disk cannot be attached to an instance with local storage.
Snapshots and images limitations	<ul style="list-style-type: none">• A snapshot generated by an existing non-encrypted disk cannot be directly converted to an encrypted snapshot.• An encrypted snapshot cannot be converted to a non-encrypted snapshot.• An image with an encrypted snapshot cannot be shared.• The encrypted snapshot and images created by it cannot be replicated across regions.
Other limitations	<ul style="list-style-type: none">• The cloud disk encryption feature relies on KMS in the same region. If you have no other operation requests, there is no need to perform additional actions in the KMS console.• 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.• 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.

Billing Description

The cloud disk encryption feature and the CMK used do not incur additional costs, nor do read and write operations on the data in the cloud disk. However, when managing operations

involving encrypted cloud disks, whether through the console or using the API for encrypted cloud disk management operations, KMS will be used in the form of an API, and will be counted towards your KMS call count in that region. KMS itself charges based on the number of calls. For billing details, please refer to [Key Management Service Billing Overview](#).

Management operations on an encrypted cloud disk include:

- Creating an encrypted cloud disk
- Attaches cloud disks
- Detaching cloud disks
- Creating a Snapshot
- Rolling 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 an Encrypted Cloud Disk Using the Console

1. Log in to the [Cloud Block Storage Console](#), select the appropriate region, and then click **Create**.
2. In the **Purchase Data Disk** dialog box, select the cloud disk encryption option.

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 **Confirm**.
4. Upon completion of your purchase, you can view the created encrypted cloud disk on the [Cloud Disk List](#) page.
The newly created encrypted cloud disk is in the **Pending Mount** state. You can refer to [Attaching Cloud Disks](#) to mount the cloud disk to a CVM in the same availability zone.

Creating Encrypted Cloud Disk Using Snapshot

Please refer to [Creating Cloud Disk from Snapshot](#) . By selecting an encrypted snapshot to create a cloud disk, you can create an encrypted cloud disk that already contains relevant data.

Creating an Encrypted Cloud Disk using API

You can create an encrypted cloud disk using the following two methods via the [CreateDisks Interface](#) :

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

Changing data encryption status

- If you need to transition the existing data on your cloud disk from an unencrypted state to an encrypted state, we recommend using the `rsync` command in Linux or the `robocopy` command in Windows to copy the data from the unencrypted disk to the newly created encrypted disk.
- If you need to convert the existing data on your cloud disk from an encrypted state to a non-encrypted state, we recommend using the same command to copy the data from the encrypted disk to a newly created non-encrypted disk.

Cloud Disk Data Backup Point

Last updated: 2023-09-19 17:44:29

The 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 data versions.

Unlike the [Cloud Disk Snapshot](#), the data backup point entirely follows the lifecycle of the cloud disk and cannot exist independently. When the cloud disk is deleted, the data backup point will also be automatically deleted and cannot be used to restore data from that backup point.

Features of the Data Backup Point

- **Historical Data Backup**

Once the data backup point is configured, Tencent Cloud will **automatically** retain a historical data backup for the cloud disk at any given moment from the previous calendar day. In the event of data anomalies caused by viruses, intrusions, or misoperations, the automatically created historical data backup point can be used to restore the system to its pre-issue state.

- **Conversion to Long-Term Persistent Backup**

The data backup point at any time point can be converted into a snapshot to retain the data backup for a longer time independent of the cloud disk's lifecycle.

Differences between Data Backup Points and Snapshots

What You Need to Know	Note
Cloud Disk Data Backup Point	<p>The Data Backup Point is a historical data backup of a past moment for the cloud disk, automatically retained by Tencent Cloud. The Tencent Cloud platform is responsible for the automatic rotation and updating of backups.</p> <ul style="list-style-type: none"> • Strengths: <ul style="list-style-type: none"> ○ Once configured, no manual intervention is required as the cloud platform automatically manages data backup and rotation. ○ The cost is only related to the size of the cloud disk and the maximum number of backup points retained. It supports prepaid payment, and the cost can be estimated. • Notes:

	<ul style="list-style-type: none">○ The lifecycle of the data backup point is entirely tied to the cloud disk. Once the cloud disk is deleted, the backup data cannot be retained.
Cloud Disk Snapshots	<p>A Cloud Disk Snapshot is a complete data backup of the cloud disk, created either manually or through an automated policy. Users need to configure the backup generation and retention policies themselves. For more information, please refer to Cloud Disk Snapshot. The lifecycle of a snapshot is independent of the cloud disk, meaning that the backup data can be retained even if the cloud disk is deleted.</p> <ul style="list-style-type: none">● Notes:<ul style="list-style-type: none">○ There is a certain operational threshold as it requires manual creation of snapshots or manual configuration of automatic policies.○ The billing method for Cloud Disk Snapshots is pay-as-you-go, billed hourly, which may have a certain degree of unpredictability.

Utilizing Cloud Disk Data Backup Points

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

- The number of data backup points indicates the maximum number of data backups retained. Tencent Cloud will **automatically** retain a historical data backup for the cloud disk configured with backup points at any given moment from the **previous calendar day**. For example:

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

- On May 3rd, the cloud disk will retain historical data from any given moment on May 2nd.
- On May 4th, the cloud disk will retain historical data from any given moment on May 3rd, and the historical data backup from May 2nd will be automatically deleted.
- The process goes on until the cloud disk reaches the end of its lifecycle.

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

- On May 3rd, the cloud disk will retain historical data from any given moment on May 2nd.
- On May 4th, this cloud disk will retain historical data from any given moment on May 3rd and May 2nd.

- On May 5th, this cloud disk will retain historical data from any given moment on May 4th, May 3rd, and May 2nd.
- On May 6th, this cloud disk will retain historical data from any given moment on May 5th, May 4th, and May 3rd.
- The process goes on until the cloud disk reaches the end of its lifecycle.

Billing description

Pricing information

The billing method for data backup points is solely related to the size of the cloud disk and the maximum number of backup points retained, and is unrelated to other factors. The billing mode for data backup points follows the cloud disk billing mode. For specific prices, please refer to the [Cloud Disk Data Backup Point Price Overview](#).

Note

- For individual users, the **first** backup point for each prepaid cloud disk is available at a limited-time offer of 90% off. The promotional period runs from June 30, 2023, to December 31, 2023.
- All other users receive a 20% discount for 3–6 cloud disk backup points and a 50% discount for 7 or more backup points. The promotion period is from June 30, 2023, to June 30, 2024.

The aforementioned benefits cannot be combined with other promotions (such as business discounts or other activities) for cumulative effect.

Fees after the adjustment

When purchasing data backup point quotas, or increasing or decreasing the maximum number of data backup points retained, please refer to [Fee for adjusting time-specific data backup points](#) for specific cost rules and billing examples.

Supports and Limits

- The data backup point is not billed like a snapshot.
- The data backup point will cease to exist with the end of the cloud disk's lifecycle, and it will no longer be possible to retain the data backup. If long-term retention is required, please convert the data backup point into a [Snapshot](#).

Directions

Notes	Reference

Enabling Data Backup Point When Purchasing a New Cloud Disk	Creating Cloud Disk
Increase, decrease/return existing data backup points of the cloud disk	Adjusting Data Backup Points
Viewing Cloud Disk Data Backup Point Information	Viewing Cloud Disk Information
Restoring Cloud Disk Data Using an Existing Data Backup Point	Using Data Backup Point to Restore Cloud Disk Data
Convert the data backup point into a snapshot to retain important data backups for a long term.	Converting Data Backup Points into Snapshots
Deleting Data Backup Point	Deleting Data Backup Point

Enhanced SSD Performance

Last updated: 2023-09-22 11:38:23

Commencing from January 2021, the individual performance of the Enhanced SSD Cloud Disk is composed of both **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

The performance of each Enhanced SSD Cloud Disk is composed of both **Basic Performance** and **Extra Performance**. Regardless of the allocation, the current performance of the Enhanced SSD Cloud Disk follows these rules:

Performance of Enhanced SSD Cloud Disk	Maximum value
Random IOPS	100000
Maximum throughput (MB/s)	1000

Basic Performance

For the Enhanced SSD Cloud Disk, the basic performance varies linearly with capacity and cannot be adjusted independently. The calculation formula is as follows:

Basic Performance	Basic Performance Calculation Formula	Maximum Basic Performance
Basic Performance – Random IOPS	$\text{Min}\{1800 + \text{Capacity in GB} \times 50, 50,000\}$	50000
Basic Performance – Maximum Throughput (MB/s)	$\text{Min}\{120 + \text{Capacity in GB} \times 0.5, 350\}$	350

According to the formula,

- Upon reaching a capacity of 460GB, the throughput limit is first achieved. At this point, the basic performance values are: 24,800 Random IOPS, with a maximum throughput of 350MB/s.

- Upon reaching a capacity of 964GB, the random IOPS limit is achieved. At this point, the basic performance values are: 50,000 random IOPS, with a maximum throughput of 350MB/s.

Extra Performance

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

Preparations

The following requirements must be met to enable extra performance:

- Currently, only **Enhanced SSD** and **ulTra SSD** support extra performance. Other types of cloud disks do not support this feature.
- Extra performance can only be configured once any aspect of the basic performance reaches its maximum value. According to the basic performance calculation formula, it is only supported when the capacity is > 460GB.

Extra performance formula

Refer to the following formula to configure the extra performance.

Extra Performance	Extra performance formula	Maximum Extra Performance
Extra Performance – Random IOPS	$\text{Min}\{\text{Configured Value} \times 128, 50,000\}$	50000
Extra Performance – Maximum throughput (MB/s)	$\text{Min}\{\text{Configured Value} \times 1, 650\}$	650

Extra performance price

For the price of extra performance, please refer to [Pricing Overview](#).

Sample

Example 1: Enhanced SSD Cloud Disk, with a capacity requirement of 2000GB and a throughput requirement of 500MB/s.

From the basic performance description, it is clear that the configured capacity has reached the upper limit of basic performance (throughput). To meet the demand, an additional performance of $(500-350)/1 = 150$ needs to be purchased. Therefore, the overall configuration to be purchased is: 2000GB Enhanced SSD Cloud Disk, with an additional performance configuration of 150.

Example Two: Enhanced SSD Cloud Disk, with a capacity requirement of 1000GB and an IOPS requirement of 50000.

From the basic performance description, it is known that the configuration capacity has reached the upper limit of basic performance (random IOPS). However, if the user does not require additional IOPS performance, there is no need to configure extra performance. Therefore, the overall configuration to be purchased is: 1000GB Enhanced SSD Cloud Disk, with no extra performance.

Cloud Disk Snapshots

Snapshot Overview

Last updated: 2023-09-19 20:46:35

Overview

- **Real-time replica of online data**

A snapshot is a fully functional copy of a cloud disk. In the event of an issue with a cloud disk for which a snapshot has been created, the snapshot can be used to swiftly restore the disk to its pre-issue state. It is recommended to create a snapshot of the relevant cloud disk prior to significant business changes, allowing for rapid data recovery in case of a failed business change.

- **Persistent backup for key milestones**

Snapshots can serve as a long-lasting backup for business data, preserving milestone states of the data.

- **Rapid business deployment**

You can use snapshot files of your business to quickly clone multiple cloud disks, achieving the goal of rapid server deployment.

Scenarios

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**

Before performing significant operations such as changing the operating system, upgrading application software, or migrating business data, you can create one or more snapshots. If any issues arise during the change process, you can promptly restore business data using the created snapshots.

- **Application of multiple replicas of production data**

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

- **Rapid environment deployment**

You can create a snapshot for a cloud server and use this system snapshot to create a custom image. You can create one or multiple instances from the created image, allowing

for the rapid and bulk deployment of cloud servers with the same environment, saving time on repeated configurations.

Billing

For detailed information about snapshot billing, see [Snapshot Billing Overview](#) and [Snapshot Pricing Summary](#).

Quota Limits

For detailed information on snapshot quota limits, please refer to [Use Limits on Snapshot](#).

Snapshot Types

- **Manual Snapshot**

Manually create a snapshot of the cloud disk data at a specific point in time. This snapshot can be used to quickly create more cloud disks with the same data, or to restore the cloud disk to the state at that point in time in the future. For specific operations, please refer to [Creating Snapshots](#).

- **Scheduled snapshots**

When your business is continuously updated, you can use scheduled snapshots to provide ongoing backup functionality. Simply create a backup policy and associate it with a cloud disk to achieve continuous data backup within a certain period, significantly improving data security. For detailed operations, please refer to [Scheduled Snapshots](#).

Note

During the snapshot creation process, there may be a situation where some application data is stored in memory and not yet persisted. This phenomenon can result in the snapshot not capturing the most recent and complete cloud disk data. Please refer to [Precautions](#) to ensure snapshot data consistency.

Case Review

Case One: Data loss due to failure to manually create a snapshot before executing high-risk operations.

For instance, Customer A has never created a snapshot for their cloud disk. One day in May 2019, due to an fio test conducted by an operator, the file system was damaged and the data could not be retrieved.

Analysis: If Customer A had created a snapshot for the cloud disk before conducting the test, they could have quickly enabled snapshot rollback after the data was damaged, promptly restoring the business.

Case 2: Failing to create scheduled snapshots for important data disk, causing data loss

For instance, Customer B had created snapshots for multiple cloud disks, but after January 2019, they did not create snapshots for newly purchased cloud disks due to cost considerations. One day in June 2019, data on a cloud disk that was not protected by a snapshot was accidentally deleted at the file system level, resulting in irretrievable data loss. **Analysis:** If Customer B had configured regular snapshot protection for this cloud disk, they could have restored the data to the state at the time of the last snapshot after the data was accidentally deleted, thus minimizing the loss. After the incident, Customer B proactively created a snapshot for this cloud disk, significantly enhancing data protection.

Case 3: Rolling back with scheduled snapshot to restore business after a misoperation

For instance, Customer C has configured snapshot protection for all cloud disks. On a certain day in May 2019, an abnormal startup occurred due to an operational error.

Analysis: Customer C promptly used a scheduled snapshot from two days prior to restore data, preventing any damage to the business.

The above are all examples of data loss due to operational errors. By comparison, it can be observed that:

- In the absence of **created snapshots**, data recovery can be extremely difficult when issues arise with servers or cloud disks, potentially leading to significant losses.
- In the case of **having created snapshots**, when issues arise with the server or cloud disk, data can be mostly recovered, resulting in minimal loss.

It is recommended to regularly create snapshots for different types of businesses to enhance data security and achieve low-cost, high-disaster tolerance for your operations.

Others

If you still encounter other [snapshot-related issues](#), we provide a collection of related questions for your reference, allowing you to quickly locate and resolve your concerns.

Cloud Access Management

CAM Overview

Last updated: 2023-09-22 11:39:59

If you utilize services such as Cloud Virtual Machine (CVM), Cloud Block Storage (CBS), Virtual Private Cloud (VPC), and databases within Tencent Cloud, and these services are managed by different individuals yet all share your cloud account key, the following issues may arise:

- Your key will be easily compromised because it is shared by several users.
- You cannot restrict the access of others, which may lead to inadvertent actions and pose significant security risks.

You can circumvent these issues by using sub-accounts to allow different individuals to manage different services. By default, sub-accounts do not have the right to use CVM or operate CVM-related resources. However, you can create policies to grant sub-accounts the necessary resources or permissions.

[Cloud Access Management \(CAM\)](#) is a web-based service provided by Tencent Cloud, primarily designed to assist customers in securely managing access permissions to resources under their Tencent Cloud accounts. With CAM, you can create, manage, and terminate users (or user groups), and control who can use which Tencent Cloud resources through identity management and policy administration.

When utilizing CAM, you can associate a policy with a user or a group of users, enabling or denying them the use of specified resources to accomplish certain tasks. For more fundamental information about CAM policies, please refer to [Policy Syntax](#). For more detailed usage information about CAM policies, please refer to [Policy](#).

If you do not require to manage the access permissions to CBS-related resources for sub-accounts, you may bypass this section. Skipping these parts will not impact your comprehension and utilization of the remaining portions of the documentation.

Quick Start

A CAM policy must either authorize the use of one or more CBS-related operations or deny the use of one or more CBS-related operations. It must also specify the resources that can be operated on (which could be all resources, or for some operations, only a subset of resources). The policy can also include conditions set for operating resources.

Task scheduling	Document
Basic policy structure	Policy Syntax
Operation definition in a policy	CBS Related Operations
Resource definition in a policy	Relevant CBS Resource Paths
Apply conditions to restrict policies.	CBS Condition Keys
Resource-level permissions supported by CBS	Resource-level permissions supported by CBS
Console Examples	Console Samples

Authorizable Resource Type

Last updated: 2023-09-19 17:50:38

Resource-level permission refers to the ability to designate which resources a user can manipulate. Cloud Block Storage (CBS) supports resource-level permissions, meaning that for CBS operations that support resource-level permissions, you can dictate when a user is permitted to perform operations or use specific resources. The types of resources that can be authorized in Cloud Access Management (CAM) are as follows:

ResourceType	Resource Description Method in Authorization Policy
Related to CBS Instances	<code>qcs::cvm:\$region::volume/*</code>

[Related to CBS Instances](#) outlines the current CBS API operations that support resource-level permissions, as well as the resources and condition keys supported by each operation.

When setting the resource path, you need to replace variables such as `$region` and `$account` with your actual parameter information. You can also use the `*` wildcard in the path. For operation examples, please refer to [Access Management Examples](#).

Note

CBS API operations not listed in the table do not support resource-level permissions. For these API operations, you can still grant users the permission to use them, but the resource element in the policy statement must be specified as `*`.

CBS APIs

Configuring using API	Resource Path	Condition Key
Attach Cloud Disks AttachDisks	<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 Cloud Disks CreateDisks	<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>
Describe Disk Operation Logs	<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>

DescribeDisks – Retrieve Cloud Disk List	qcs::cvm:\$region:\$account:volume/* qcs::cvm:\$region:\$account:volume/\$diskId	cvm:region cvm:zone cvm:disk_type
Detach Cloud Disks – DetachDisks	qcs::cvm:\$region:\$account:volume/* qcs::cvm:\$region:\$account:volume/\$diskId	cvm:region cvm:zone cvm:disk_type
Modify Disk Attributes ModifyDiskAttributes	qcs::cvm:\$region:\$account:volume/* qcs::cvm:\$region:\$account:volume/\$diskId	cvm:region cvm:zone cvm:disk_type
Modify the billing mode of a cloud disk ModifyDisksChargeType	qcs::cvm:\$region:\$account:volume/* qcs::cvm:\$region:\$account:volume/\$diskId	cvm:region cvm:zone cvm:disk_type
Modify the renewal flag of a cloud disk ModifyDisksRenewFlag	qcs::cvm:\$region:\$account:volume/* qcs::cvm:\$region:\$account:volume/\$diskId	cvm:region cvm:zone cvm:disk_type
Renew Cloud Disk RenewDisk	qcs::cvm:\$region:\$account:volume/* qcs::cvm:\$region:\$account:volume/\$diskId	cvm:region cvm:zone cvm:disk_type
Expand Cloud Disk ResizeDisk	qcs::cvm:\$region:\$account:volume/* qcs::cvm:\$region:\$account:volume/\$diskId	cvm:region cvm:zone cvm:disk_type
Return Cloud Disks TerminateDisks	qcs::cvm:\$region:\$account:volume/* qcs::cvm:\$region:\$account:volume/\$diskId	cvm:region cvm:zone cvm:disk_type

Authorization Policy Syntax

Last updated: 2023-09-22 14:06:43

Policy Syntax

CAM policy:

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

- **Version:** This is a mandatory field. Currently, only the value `"2.0"` is permitted.
- **statement** describes the details of one or more permissions. It contains a permission or permission set of multiple other elements such as `effect`, `action`, `resource`, and `condition`. One policy has only one `statement`.
 - **Action:** Describes the permitted or denied operations. An operation can be an API (described with a 'name' prefix) or a feature set (a specific set of APIs described with a 'permid' prefix). **This element is mandatory.**
 - **Resource:** Describes the specific data authorized. Resources are described in a six-segment format. The detailed resource definitions vary by product. **This element is mandatory.**
 - **Condition:** Describes the constraints under which the policy takes effect. A condition consists of an operator, an operation key, and an operation value. The condition value can include information such as time and IP address. Some services allow you to specify other values in the condition. This element is optional.
 - **Effect:** Describes whether the result of the statement is "allow" or "explicitly deny". It includes two scenarios: allow (permission) and deny (explicit rejection). **This element is mandatory.**

CBS Operations

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

If you need 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 wildcards to specify multiple actions. For example, you can specify all actions whose names begin with the word "Describe", as follows:

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

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

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

CBS Resource Path

Each CAM policy statement has its own applicable resources. Resource paths are generally in the following format:

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

- **project_id:** Describes project information, included only for compatibility with early CAM logic, and does not need to be filled in.
- **service_type:** Refers to the product abbreviation, such as `CVM`.
- **region:** Region information, such as "bj".
- **account:** Describes the root account of the resource owner, such as `uin/164256472`.
- **Resource:** Detailed resource information of each product, such as "volume/diskid1" or "volume/*".

For instance, you can specify a specific CBS resource (disk-abcdefg) in the statement as illustrated below:

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

You can also use the wildcard `"*"` to designate all CBS resources belonging to a specific account, as illustrated below:

```
"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 permission control, you can use the wildcard "*" in the Resource element as shown below:

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

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

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

CBS Condition Keys

In policy statements, you can optionally specify conditions to control the effective time of the policy. Each condition contains one or more key-value pairs. Condition keys are case-insensitive.

- If you specify multiple conditions or specify multiple keys in one condition, they will be evaluated with the logical operation "AND".
- If you specify a key with multiple values in a single condition, it will be evaluated using the logical OR operation. All conditions must be matched to grant permissions.

The table below describes the CBS condition keys used for specific services:

Condition keys	Reference types	Key value pairs
cvm:region	String	cvm:region= region <ul style="list-style-type: none"> • Here, region refers to the region (for example, "ap-guangzhou").
cvm:disk_type	String	cvm_disk_type= disk_type <ul style="list-style-type: none"> • Wherein, disk_type refers to the type of disk (for example, CLOUD_PREMIUM).

Relevant Products

Last updated: 2023-09-19 17:52:38

Product name	Relationship with Cloud Block Storage
Cloud Virtual Machine	Cloud Block Storage can serve as both the system disk and data disk for a Cloud Virtual Machine, providing an elastically scalable block storage device.
GPU Cloud Computing	SSD Cloud Disk and Premium Cloud Disk can serve as both the system disk and data disk for GPU Cloud Computing, providing an elastically scalable block storage device.
Tencent Cloud Observability Platform	Performance metrics of Cloud Block Storage can be viewed through Tencent Cloud's Observability Platform, including Disk Read IOPS, Disk Read Throughput, Disk Write IOPS, Disk Write Throughput, Disk I/O Wait Time, Disk I/O Service Time, and Disk I/O Busy Ratio.
Tags	Tags are utilized to identify cloud resources, enabling you to categorize and search for Cloud Block Storage resources.

Predecessor

Last updated: 2023-09-19 17:54:02

HDD cloud disk

Note

Currently, the sale of HDD cloud disks has been discontinued. The following content is for reference only for users who are still using HDD cloud disks.

HDD cloud disk is the previous generation of cloud disk types provided by Tencent Cloud, suitable for low I/O load business scenarios where data is infrequently accessed. It uses magnetic media as the storage medium and employs a three-copy distributed mechanism to achieve highly reliable data storage.

Performance Metrics

Performance Metrics	HDD cloud disk
Random IOPS	Offers a range of specifications from 10 GB to 16,000 GB, along with hundreds of random IOPS performance.
Throughput (MB/s)	Supports an I/O throughput performance of 50 MB/s.
Latency	-

Scenarios

For scenarios where requests are relatively infrequent, such as system logs, enterprise work files, data warehouses, small-sized blogs, and BBS, we recommend that you select HDD cloud disks to reduce costs.