

Cloud File Storage

Product Introduction

Product Documentation



Copyright Notice

©2013-2024 Tencent Cloud. All rights reserved.

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

Trademark Notice



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

Service Statement

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

Contents

Product Introduction

Overview

Strengths

Storage Classes and Performance

Use Cases

Recommended Regions

Use Limits

Product Introduction

Overview

Last updated : 2024-01-22 21:58:08

Cloud File Storage (CFS) provides a scalable shared file storage service that can be used with Tencent Cloud services such as CVM, TKE, and BatchCompute. CFS offers standard NFS and CIFS/SMB file system access protocols to provide shared data sources for multiple CVM instances or other computing services. It supports elastic capacity expansion and performance scaling. CFS can be mounted on existing applications without modification. As a highly available and reliable distributed file system, CFS is suitable for various scenarios such as big data analysis, media processing, and content management.

CFS is easy to integrate, eliminating your need to adjust your business structure or make complex configurations. To integrate and use CFS, simply complete three steps: creating a file system, launching a file system client on a server, and mounting the created file system.

Product features

Integrated management

CFS supports NFS v3.0/v4.0 and CIFS/SMB2.0/SMB2.5/SMB3.0 protocols as well as POSIX access syntax (such as strong data consistency and file locking). You can mount a file system by running the standard mount command on the corresponding operating system.

Automatic expansion

CFS can automatically expand the storage capacity of a file system based on file size without interrupting requests and applications during the process, thereby ensuring exclusive use of storage resources while reducing management workload.

Security settings

CFS features extremely high availability and persistence. Each file stored in a CFS instance has 3 redundant copies. It supports access from VPC and classic network as well as access control.

Pay-as-you-go

CFS is billed by actual usage with no minimum fees or deployment or Ops fees. It allows multiple CVM instances to share the same storage capacity via the NFS and CIFS/SMB protocols, eliminating your need to purchase other storage services or care about cache.

Strengths

Last updated : 2024-01-22 21:58:08

Ease of use

Tencent Cloud compute resources can mount different Cloud File Storage (CFS) resources based on SMB/NFS/private protocols, and access them as if they were local disks using POSIX instructions.

CFS comes with a console where you can quickly create and configure a file system, reducing the time needed for deployment and the workload of file system maintenance.

Flexible scalability

CFS can be scaled out as needed, without interrupting requests and applications. This ensures continuous business operations and simplifies system management.

A single CFS namespace can provide petabytes of storage capacity, supporting data storage at a larger scale compared with self-built NAS, cloud hard disk, and other storage methods.

High security and reliability

CFS adopts a distributed three-copy storage mechanism that is highly reliable. The system guarantees that data is written in all three copies before returning a response of successful write. If any copy fails, a new copy can be created in the backend by using methods such as data migration, ensuring the availability of three data copies at all times. Meanwhile, the access layer guarantees the running of the HA mechanism to provide an overall reliability of up to 99.9999999% (9 9's), providing secure and reliable data storage services for you.

CFS can strictly control access to file systems, allowing you to control access with the help of security groups in classic networks or VPCs together with permission groups.

Best-in-class performance

CFS has the capabilities to provide 100-GB bandwidth, 1-million OPS, and 10-million IOPS, accommodating your data storage under all kinds of harsh performance conditions.

CFS supports concurrent access by thousands of clients. Based on intelligent load balancing, the cluster backend of CFS can ensure stable performance on the clients, fulfilling your storage needs involving high concurrency.

CFS can achieve sub-millisecond latency, meeting your business requirements in latency-sensitive scenarios.

Storage Classes and Performance

Last updated : 2024-01-22 21:58:08

Cloud File Storage (CFS) provides file systems with scalable storage, which can work with other Tencent Cloud services such as CVM, TKE, and BatchCompute. CFS offers the following storage classes, and you can select one based on your business needs.

Storage classes

Standard

Standard is a highly cost-effective file system that uses mixed media and accelerates data reads/writes through a data tiering mechanism. Three replicas on three independent physical servers on different racks are provided to guarantee strong consistency and successful storage of every data entry written to the file system. Its access server features hot data migration to ensure data reliability and high service availability, making it suitable for scenarios that require small-scale general data storage.

High-Performance

High-Performance is a low-latency file system that uses NVMe only and provides a high storage performance through a data tiering mechanism. Three replicas on three independent physical servers on different racks are provided to guarantee strong consistency and successful storage of every data entry written to the file system. Its access server features hot data migration to ensure data reliability and service high availability, making it suitable for small-scale core businesses that are latency-sensitive.

Standard Turbo

Standard Turbo is a parallel file system that uses mixed media and an asymmetric framework. Data nodes and metadata nodes are deployed independently. By allowing mounting with a private protocol, a single client can deliver performance like a storage cluster. In addition, underlying resources are isolated to ensure exclusive storage for the cluster. Three replicas on three independent physical servers on different racks are provided to guarantee strong consistency and successful storage of every data entry written to the file system. Its access server features hot data migration to ensure data reliability and service high availability, making it suitable for scenarios that require large-scale throughput and mixed loads.

High-Performance Turbo

High-Performance Turbo is a high-bandwidth, low-latency, parallel file system that uses NVMe only and an asymmetric framework. Data nodes and metadata nodes are deployed independently. By allowing mounting with a private protocol, a single client can deliver performance like a storage cluster. In addition, underlying resources are

isolated to ensure exclusive storage for the cluster. Three replicas on three independent physical servers on different racks are provided to guarantee strong consistency and successful storage of every data entry written to the file system. Its access server features hot data migration to ensure data reliability and service high availability, making it suitable for scenarios that use a large number of small files.

High-Throughput

High-Throughput is a parallel file system that uses a layered framework. It provides more flexible bandwidth scaling and access over the SMB protocol, meeting the storage requirements for a small capacity and high bandwidth. Three replicas on three independent physical servers on different racks are provided to guarantee strong consistency and successful storage of every data entry written to the file system. Its access server features hot data migration to ensure data reliability and high service availability, making it suitable for read-intensive scenarios such as rendering, game battle server, and non-linear editing.

Performance and specifications

General series

Product	Standard	High-Performance
Positioning	Cost-effective, suitable for small-scale general data storage	High performance and low latency, suitable for small-scale latency-sensitive core businesses
Scenario	Small-scale enterprise file sharing, data backup/archive, and log storage	Small-scale CI/CD development and testing environments, high-performance web services, OLTP databases, and high-performance file sharing
Storage capacity	0–160 TiB	0–32 TiB
Bandwidth (MiB/s)	$\text{Min}\{100 + 0.1 \times \text{capacity in GiB}, 300\}$	$\text{Min}\{200 + 0.2 \times \text{capacity in GiB}, 1024\}$
Read IOPS	$\text{Min}\{2,000 + 8 \times \text{size in GiB}, 15,000\}$	$\text{Min}\{2,500 + 30 \times \text{size in GiB}, 30,000\}$
Write IOPS	$\text{Min}\{2,000 + 8 \times \text{size in GiB}, 15,000\}$	$\text{Min}\{2,500 + 30 \times \text{size in GiB}, 30,000\}$
Maximum OPS	Read/Write: 10,000/1,000	Read/Write: 30,000/3,000
Latency	4K single-thread read: 3 ms 4K single-thread write: 7 ms	4K single-thread read: 1 ms 4K single-thread write: 1.5 ms

Cost	0.05 USD/GiB/month	0.2286 USD/GiB/month
Supported protocol	NFS/SMB	NFS
Scaling	Auto	Auto
Supported OS	Linux/Windows	Linux/Windows

Turbo series

Product	Standard Turbo	High-Performance Turbo
Positioning	High-throughput and large storage, suitable for businesses that require high throughput and mixed loads	High-throughput and high IOPS, suitable for businesses that use large-scale small files
Scenario	Non-linear media asset editing, image rendering, AI inferencing, OLAP business, and high-performance computing	High-performance and large-scale computation, AI training, OLTP databases, big data analysis, and OLAP services
Storage capacity	20 TiB to 100 PiB	10 TiB to 100 PiB
Bandwidth (MiB/s)	Min{0.1 x capacity in GiB, 100,000}	Min{0.2 x capacity in GiB, 100,000}
Read IOPS	Min{2 x capacity in GiB, 2 million}	Min{20 x capacity in GiB, 10 million}
Write IOPS	Min{1 x capacity in GiB, 1 million}	Min{5 x capacity in GiB, 3 million}
Maximum OPS	Read/Write: 300,000/20,000	Read/Write: 300,000/20,000
Latency	4K single-thread read: 0.2 ms 4K single-thread write: 3 ms	4K single-thread read: 0.2 ms 4K single-thread write: 1.5 ms
Cost	0.0857 USD/GiB/month	0.2 USD/GiB/month
Supported protocol	POSIX/MPI	POSIX/MPI
Scaling	Manual	Manual
Supported OS	Linux	Linux
Consistency	Strong consistency	Strong consistency

Infrequent Access (IA)

Product	IA
Positioning	Storage of infrequently accessed warm and cold data
Scenario	Used together with Standard Turbo or High-Performance Turbo to achieve automatic hot-cold data tiering, reducing storage costs
Storage capacity	0-1 EiB
Bandwidth	600 MiB/s
Cost	Storage usage: 0.0171 USD/GiB/month Data transfer: 0.0085 USD/GiB
Scaling	Auto

Note:

An IA file system cannot be mounted directly for access. It must be used together with Standard Turbo or High-Performance Turbo to achieve automatic hot-cold data tiering.

High-Throughput

Product	High-Throughput
Positioning	High throughput and large capacity, suitable for large-scale read-intensive businesses
Scenario	Read-intensive scenarios such as video rendering, game battle server, and non-linear editing
Storage capacity	0-1 PiB
Bandwidth (MiB/s)	0-200 GiB/s (dependent on the deployment workload)
Read IOPS	Min{2 x capacity in GiB, 2 million}
Write IOPS	Min{1 x capacity in GiB, 1 million}
Maximum OPS	Read/Write: 300,000/20,000
Latency	4K single-thread read: 5 ms 4K single-thread write: 10 ms
Cost	Capacity: 0.1428 USD/GiB/month Bandwidth: 428.571 USD/GiB/s/month

Supported protocol	SMB
Scaling	Auto
Supported OS	Windows

Notes

In the performance-related formulas, the capacities of Standard Turbo and High-Performance Turbo refer to the capacities purchased for the cluster. For Standard and High-Performance, the capacities refer to the storage that is actually used by the instances.

The table above shows the capabilities of the file system. To reach the performance upper threshold, you usually need to perform multi-threaded reads/writes using multiple compute nodes.

The performance benchmark is tested in interruption-free conditions. The results of mixed tests or other loads may vary.

OPS indicates the file system's ability to process metadata per second, which is not the same as IOPS.

Currently, High-Throughput CFS is not available for purchase in the console. If you need to purchase it, please [submit a ticket](#).

Use Cases

Last updated : 2024-01-22 21:58:08

Enterprise file sharing

Cloud File Storage (CFS) provides storage services suitable for organizations with a large number of employees who need to access and share the same datasets. It allows the administrators of an organization to create file systems and set read/write permissions for internal clients.

High-performance computing and big data analysis

CFS provides the scale, performance, high throughput of computing nodes, read-after-write consistency, and low-latency file operations required by high-performance computing and big data applications, making it particularly suitable for such scenarios as machine learning, AI training, and centralized server log processing and analysis.

Streaming media processing

For media workflows such as video editing, video production, broadcast processing, sound design, and rendering, shared storage is generally used to work with large files. The strong data consistency model of CFS, coupled with high throughput and shared file access, helps reduce the time it takes to complete such tasks.

Content management and web services

As a highly persistent, high-throughput file system, CFS can be used for various content management systems. It stores and provides data for various applications such as websites, online distribution, and archiving.

Dedicated software environment

CFS provides the foundation for the migration of traditional service architectures to the cloud for government, education, and healthcare sectors. Generally, a dedicated software program needs to share the same file storage system and only supports POSIX standard protocol operations.

Recommended Regions

Last updated : 2024-01-22 21:58:08

The following table describes regions where Cloud File Storage (CFS) is available.

Note:

If you need to use CFS in a region where resources are sold out for your CVM instances or services, you can choose a VPC in that region, create a subnet in an availability zone (AZ) where resources can be purchased, and then create a CFS file system in that subnet. For more information, see [FAQs](#).

This table recommends AZs for resource deployment based on the resource-adding capacity and resource adequacy of CFS. This doesn't mean CFS resources are always available for purchase in these AZs. If you purchase resources in recommended AZs, subsequent addition of resources can be guaranteed.

Note:

Refer to the information displayed in the console to see whether you can purchase. This table **does not represent that all types of CFS services** have resources in the listed regions.

Region	AZ
Beijing	Beijing Zone 6 Beijing Zone 7
Shanghai	Shanghai Zone 4 Shanghai Zone 5 Shanghai Zone 8
Guangzhou	Guangzhou Zone 6 Guangzhou Zone 7
Nanjing	Nanjing Zone 1 Nanjing Zone 2 Nanjing Zone 3
Chengdu	Chengdu Zone 1 Chengdu Zone 2
Chongqing	Chongqing Zone 1
Hong Kong (China)	Hong Kong (China) Zone 1 Hong Kong (China) Zone 2
Shanghai Finance	Shanghai Finance Zone 1 Shanghai Finance Zone 2
Beijing Finance	Beijing Finance Zone 1

Shenzhen Finance	Shenzhen Finance Zone 1 Shenzhen Finance Zone 2
Singapore	Singapore Zone 1 Singapore Zone 2 Singapore Zone 3 Singapore Zone 4
Tokyo	Tokyo Zone 1
West US	Silicon Valley Zone 1
East US	Virginia Zone 2
Mumbai	Mumbai Zone 1
Thailand	Bangkok Zone 1 Bangkok Zone 2
Seoul	Seoul Zone 1
Toronto	Toronto Zone 1
Frankfurt	Frankfurt Zone 1 Frankfurt Zone 2
Indonesia	Jakarta Zone 1

Use Limits

Last updated : 2024-01-22 21:58:08

Limits and Notes

Limits

Product type	Standard		High-Performance		Standard Turbo	
	Specification	Recommended	Specification	Recommended	Specification	Recommended
Maximum system capacity	160 TiB	140 TiB	32 TiB	24 TiB	100 PiB	40 PiB
Minimum system capacity	None		None		40 TiB	
Maximum system bandwidth	300 MiB/s	240 MiB/s	1 GiB/s	800 MiB/s	100 GiB/s	100 GiB/s
Maximum number of system files	Min[15,000 x used capacity (GiB), 1 billion]	Min[10,000 x used capacity (GiB), 0.8 billion]	Min[20,000 x used capacity (GiB), 15 billion]	Min[15,000 x used capacity (GiB), 1 billion]	Min[15,000 x deployed capacity (GiB), 1 billion]	Min[15,000 x deployed capacity (GiB), 1 billion]
Maximum number of system directories	10 million	8 million	15 million	10 million	10 million	10 million
Maximum length of filename	255 bytes	255 bytes	255 bytes	255 bytes	255 bytes	255 bytes
Maximum length of absolute path	4096 bytes	4096 bytes	4096 bytes	4096 bytes	4096 bytes	4096 bytes
Maximum number of directory levels	1000	16	1000	16	1000	1000
Maximum number of files/subdirectories per directory	1 million	0.8 million	1 million	0.8 million	1 million	1 million

Maximum number of concurrently opened files	65536	1000	65536	1000	65536	1
Maximum number of locks per file	512	512	512	512	512	5
Maximum number of clients	1000	100	1000	100	2000	1
Maximum bandwidth per client	300 MiB/s	300 MiB/s	500 MiB/s	500 MiB/s	10 GiB/s	1
Maximum number of mounted file systems per client	1000	16	1000	16	16	8
Billing	Billed by the actual usage (excluding prepaid)		Billed by the actual usage (excluding prepaid)		Billed by the purch	
Supported protocol	NFS/SMB		NFS		POSIX/MPI	
Supported OS	Linux/Windows				Linux	

Notes

Turbo series

The Turbo series is mounted using a client. After you run the `mount` command on the client installed, you can use the file system the same way as a local file system.

The Turbo series is billed according to the capacity purchased. For example, if you purchased a 40 TiB file system of the Turbo storage class, you will be billed at the 40 TiB rates by hour, regardless of how much you actually use. For example, if you use the file system for 1 hour, the fee will be calculated as follows: $40 \times 1024 \times 0.0857/24/30 = 4.876$ USD. The file system can be terminated anytime.

To ensure the cloud load balance of the file system after scaling up, we recommend that you scale up when around 80% of the capacity has been used. Online scale-up is supported and will be imperceptible during the whole process. The Turbo series cannot be scaled down. You can create a Turbo instance, migrate your data, and then delete the old instance.

Because the self-deployed cluster needs to be set up again, the initial creation of the Turbo series will take about 20 minutes.

It is recommended to mount and use a Turbo file system only on a client in the same availability zone (AZ). Cross-AZ delays may cause issues such as client mount timeouts and disconnections.

If you need the Turbo series with higher specifications (supporting more files, directories, etc.), please [submit a ticket](#).

UID and GID

When the NFS v3.0 protocol is used, if the UID or GID of the file does not exist in the local account, then the UID and GID will be displayed directly. Otherwise, the relevant username and group name will be displayed based on the mapping relationship of the local UID and GID.

When the NFS v4.0 protocol is used, if the Linux version is above 3.0, the UID rules and the GID rules will be the same as those in the NFS v3.0 protocol. Otherwise, the UID and GID of all files will be displayed as `nobody` .

Note:

When you mount a file system to a client whose Linux version is below 3.0 by using the NFS v4.0 protocol, we recommend that you refrain from performing "change owner" or "change group" on the file or directory. Otherwise, its UID and GID will become `nobody` .

Supported CIFS/SMB protocols

Supported protocol versions: CIFS/SMB 1.0 and later are supported. However, SMB 1.0 is not recommended for mounting, because it is inferior in terms of performance and features to SMB 2.0 and later and because Windows has stopped its technical support service for Windows versions supporting SMB 1.0 or earlier.

You cannot use NFS and SMB to access the same file system at the same time or directly access an SMB file system via WAN.

Read/write ACL is provided only at the file system level. No ACL is provided at the file/directory level.

IOCTL/FSCCTL operations such as sparse files setting, file compression, ENI status query, and reparse point setting are not supported.

Alternate Data Streams are not supported.

Some protocol features in SMB 3.0 or later such as SMB Direct, SMB Multichannel, SMB Directory Leasing, and Persistent File Handle are not supported.