

Cloud File Storage

Operation Guide

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

Operation Guide

- Access Management

- Managing File Systems

- Permission Management

- Using Tags

- Snapshot Management

 - What Is a File System Snapshot?

 - Use Limits

 - Creating a Snapshot

 - Scheduled Snapshots

 - Creating a File System Using a Snapshot

 - Deleting a Snapshot

- Guide for Cross-AZ and Cross-Network Access

- Automatically Mounting File Systems

- Data Migration Service

 - Features

 - Use Limits

 - Starting a Migration Task

 - Viewing Migration Task Result

- User Permission Management

- User Quotas

 - Operation Guide

 - Use Limits

 - Features

- Data Encryption

 - Operation Guide

 - Use Limits

 - Overview

- Data Lifecycle Management

 - Use Limits

 - Operation Guide

 - Features

- Upgrading Standard File Systems

Operation Guide

Access Management

Last updated : 2024-01-22 22:15:48

Overview

Cloud File Storage (CFS) supports access management at the resource level, i.e., allowing the root account to grant users and user groups permissions to manipulate specified resources. After authorization, specified users will be allowed or forbidden to perform operations using the CFS console or APIs based on the permissions granted. This document describes how to configure read-only, read/write, and custom policies for CFS users. For more information on how Cloud Access Management (CAM) works and can be used, see [CAM Overview](#).

Directions

Creating an access control policy

Log in to the [CAM console](#) and go to the policy management page.

To grant users permissions quickly, search for "CFS", select the preset read-only or read/write permissions, and associate them with the specified user group.

If you need to grant users permissions for specific operations, you can create a custom policy and associate it with the specified user group.

Full read/write permission policy

If you want to authorize users to perform all operations such as CRUD, associate them with the

`QcloudCFSFullAccess` policy. Below is the policy syntax for using the preset `QcloudCFSFullAccess` policy to grant collaborators or sub-users full read/write access to all CFS resources and VPC/subnet query permission:




```
{  
  "version": "2.0",  
  "statement": [  
    {  
      "action": [  
        "cfs:*"  
      ],  
      "resource": "*",  
      "effect": "allow"  
    },  
    {  

```

```
        "action": [
            "vpc:DescribeVpcEx",
            "vpc:DescribeSubnetEx"
        ],
        "resource": "*",
        "effect": "allow"
    }
]
```

Read-only permission policy

If you want to grant users permission to query but not create, modify, or delete resources, associate them with the `QcloudCFSReadOnlyAccess` policy. Below is the policy syntax for using the preset `QcloudCFSReadOnlyAccess` policy to grant collaborators or sub-users read-only access to all CFS resources and VPC/subnet query permission:



```
{
  "version": "2.0",
  "statement": [
    {
      "action": [
        "cfs:Describe*"
      ],
      "resource": "*",
      "effect": "allow"
    },
    {
```

```

        "action": [
            "vpc:DescribeVpcEx",
            "vpc:DescribeSubnetEx"
        ],
        "resource": "*",
        "effect": "allow"
    }
}

```

Custom policy

Custom policies allow more flexibility in permission management. The CAM console offers multiple methods for generating custom policies. This example shows you how to create a custom policy by using a **Policy Generator**. For other methods, please see [Creating Custom Policies](#).

The CAM policy generator is very user friendly. You simply need to select the desired parameters, and policy code will be generated automatically. This is especially suitable for first-time CAM users.

Log in to the [CAM Policies page](#), and select **Create Custom Policy > Create by policy generator**. Use the policy generator to create a custom policy to which you can add multiple statements. The configurations are described as below:

Parameter	Options and Effect
Effect	Allow or Reject
Service	Select CFS here
Action	All CFS-supported actions
Resource	<p>Specify the resources in six-segment format:</p> <p>For all resources in CFS, enter <code>*</code> .</p> <p>For all resources in a specified region, use the format <code>qcs::cfs:ap-guangzhou::*</code> .</p> <p>For all resources in all regions under a specified user account, use the format <code>qcs::cfs::uin/27700000:*</code> .</p> <p>For all file systems in a specified region under a specified user account, use the format <code>qcs::cfs:ap-guangzhou:uin/27700000:filesystem/*</code> .</p> <p>For file systems in a specified user group under a specified user account, use the format <code>qcs::cfs::uin/27700000:pgroup/pgroup-doxpcqh</code> .</p> <p>Note: The UIN in a policy must be a root account UIN. The file systems or permission group resources must belong to the root account.</p> <p>If a policy generator is used:</p> <p>Service: You can only enter "cfs".</p> <p>Region: Select a region from the drop-down list box as needed.</p> <p>Account: The current account information will be auto filled. If it is not filled, you can enter <code>uin/xxxxxxx</code> , where <code>xxxxxxx</code> is the UIN.</p>

	<p>Resource prefix: You can enter <code>filesystem</code> , <code>snap</code> , or <code>resource</code> , which represents file system instances, snapshot instances, or storage resource unit instances, respectively.</p> <p>Resource: Enter a resource ID, such as <code>cfs-xxxxxx</code> .</p>
Condition	Specifies under which condition this policy will take effect. For more information, see Conditions .

The APIs, API features, and notes for authorization are listed in the table below. You can set your resource permissions accordingly.

API Category	API Name	API Description	Permission Type	Note
Service APIs	SignUpCfsService	Activates the CFS service	Write	You do not need to specify resources when authorizing this API.
	DescribeCfsServiceStatus	Queries whether the CFS service is activated	Read	You do not need to specify resources when authorizing this API.
File system APIs	DescribeCfsFileSystems	Lists file systems	Read	You need to specify the resources as <code>*</code> when authorizing this API.
	CreateCfsFileSystem	Creates a file system	Write	You do not need to specify file system resources when authorizing this API.
	UpdateCfsFileSystemName	Updates the file system name	Write	You need to specify file system resources when authorizing this API.
	UpdateCfsFileSystemPGroup	Updates the permission group for a file system	Write	You need to specify file system resources when authorizing this API.

	UpdateCfsFileSystemSizeLimit	Updates the file system quota	Write	You need to specify file system resources when authorizing this API.
	DeleteCfsFileSystem	Deletes a file system	Write	You need to specify file system resources when authorizing this API.
	DescribeMountTargets	Queries mount targets	Read	You need to specify file system resources when authorizing this API.
	AddMountTarget	Creates a mount target	Write	You need to specify file system resources when authorizing this API.
	DeleteMountTarget	Deletes a mount target	Write	You need to specify file system resources when authorizing this API.
Permission group APIs	DescribeCfsPGroups	Lists permission groups	Read	You need to specify the resources as  when authorizing this API.
	CreateCfsPGroup	Creates a permission group	Write	You do not need to specify resources when authorizing this API.
	UpdateCfsPGroup	Updates the information of a permission group	Write	You need to specify permission group resources when authorizing this API.


	DeleteCfsPGroup	Deletes a permission group	Write	You need to specify permission group resources when authorizing this API.
	DescribeCfsRules	Lists permission group rules	Read	You need to specify permission group resources when authorizing this API.
	CreateCfsRule	Creates a permission group rule	Write	You need to specify permission group resources when authorizing this API.
	UpdateCfsRule	Updates the information of a permission group rule	Write	You need to specify permission group resources when authorizing this API.
	DeleteCfsRule	Deletes a permission group rule	Write	You need to specify permission group resources when authorizing this API.
Key APIs	DescribeKmsKeys	Queries KMS keys	Read	You need to specify the resources as <code>*</code> when authorizing this API.

Note:

As CFS file systems use VPC IPs, permissions for "vpc:DescribeVpcEx" and "vpc:DescribeSubnetEx" APIs are needed to create, list, and query file systems. We strongly recommend granting all VPC resources permissions for these two APIs in all your CFS authorization policies. See the `QcloudCFSReadOnlyAccess` policy statement to learn how to write the policy.

After configuring the above parameters, click **Add Statement** to add a statement to the custom policy. You can repeat this operation to add multiple statements. In case that a policy already exists or conflicts with other policies, see [Syntax Structure](#) for more information on whether and how they will take effect.


A policy should be written in the following format. There can be multiple statements in one policy.



```
{
  "version": "2.0",
  "statement": [{
    "effect": "Effect",
    "action": [
      "Action"
    ],
    "resource": "Resource"
  }]
```

```
}]  
{
```

For example, the policy syntax for prohibiting users from deleting certain file systems and updating quotas is as follows:



```
{  
  "version": "2.0",  
  "statement": [{  
    "effect": "deny",  
    "action": [
```

```
        "name/cfs:DeleteCfsFileSystem",
        "name/cfs:UpdateCfsFileSystemSizeLimit"
    ],
    "resource": [
        "qcs::cfs::uin/2779643970:filesystem/cfs-11111111",
        "qcs::cfs::uin/2779643970:filesystem/cfs-22222222",
        "qcs::cfs::uin/2779643970:filesystem/cfs-33333333"
    ]
  }
}
```

Authorizing a user or user group

If you want to grant an existing permission, you can search for `QcloudCFSFullAccess` , `QcloudCFSReadOnlyAccess` , or a custom policy and click **Bind User/Group** in the **Operation** column. Then, locate and select the user or user group that needs to be authorized and click **OK**.

Deauthorizing a user or user group

If you need to deauthorize a user or user group, click the policy name to go to the policy details page. Select the user or user group under the **User/User Group** tab and click **Remove User** or **Remove Group**. Click **OK** in the pop-up window. The CFS permissions of the user or user group will be revoked.

Managing File Systems

Last updated : 2024-01-22 22:15:47

Overview

You can view the created file systems and manipulate them in the CFS console, such as viewing the file system status and usage, file system details, and mount target information.

Note:

If a file system is **Creating**, you cannot view its details or delete it.

Prerequisites

Log in to the [CFS console](#) and click **File System** on the left sidebar to enter the file system list page.

Directions

Viewing file system status and usage

You can view the current file system usage and status on the file system list page. CFS also allows you to search by file system name, ID, VPCID, and IP.

Viewing a file system

Click a file system name on the file system list page to enter the file system details page, where you can view the basic information of the file system as well as information of its mount target and clients.

Basic information of a file system

The basic information of a file system includes its region, ID, name, file service protocol, file system status, and creation time. You can set the file system name on this page.

Mount target information

The mount target information of an NFS file system includes the network information, permission group, and recommended mount commands. You can modify the file system permission group on this page.

Mounted client information

You can select the **Mounted Clients** tab to view the information of the clients where the file system is mounted.

Note:

Client information display may have a delay of 1–3 minutes.

Snapshot chain

You can view and manage the snapshots of a file system on this page.

Renaming a file system

Click the file system to be renamed in the file system list to enter the file system details page, where you can click



to the right of the instance name to rename the file system.

Deleting a file system

If you no longer need a file system, you can find it in the file system list and select **Delete** in the **Operation** column to delete it.

Note:

To avoid system exceptions on clients, disconnect the file system from all clients before deleting it.

Permission Management

Last updated : 2024-01-22 22:15:48

Overview

A client must be in the same network as the file system, for which a permission group needs to be configured to manage the access and read/write permissions of the client. This document describes how to do so.

Directions

Step 1. Create a permission group

1. Log in to the [CFS console](#) and click **Permission Group** on the left sidebar.
2. On the permission group page, click **Create**. In the pop-up window, configure the permission group name and remarks.

Step 2. Add a permission group rule

Click the name of a permission group to enter the rule list page. You can add, edit, or delete rules in the rule list. If no rule is added to the permission group, all IPs will be allowed. The rules are described as below:

Field	Description
Access Address	You can enter a single IP or a CIDR block, such as 10.1.10.11 or 10.10.1.0/24. The default access address is <code>*</code> , indicating that all IPs are allowed. Please note that you need to enter the CVM instance's private IP here.
Read & Write Permissions	Read-only or read/write.
User Permissions	<p>The four options below are used for controlling the permissions of a user.</p> <p><code>all_squash</code>: Any user will be mapped to an anonymous user or user group.</p> <p><code>no_all_squash</code>: A user will be first matched with a local user, and if the match fails, it will be mapped to an anonymous user or user group.</p> <p><code>root_squash</code>: A root user will be mapped to an anonymous user or user group.</p> <p><code>no_root_squash</code>: A root user will be allowed to maintain root account permissions.</p> <p>Note:</p> <p>User permissions configuration is not supported for CIFS/SMB file systems and Turbo file systems and will not take effect.</p> <p>The default permission is 755 for each file system, and <code>nfsnobody</code> does not have write permission. Therefore, if there are no special needs, <code>no_root_squash</code> is recommended. If the root user creates a file directory and mounts the file system, when the access address is</p>

	set to <code>all_squash</code> or <code>root_squash</code> , an access IP can only read files. (This is because the mount path requires root permissions, but the access IP has been mapped to an anonymous user.)
Priority	<p>You can configure an integer between 1 and 100 as the priority level, where 1 indicates the highest priority. If the permission of a single IP conflicts with that of an IP within a CIDR block in the same permission group, the permission with a higher priority will apply. If their priority levels are the same, the permission of the single IP will apply. If two overlapping CIDR blocks are configured with different permissions and the same priority levels, the permissions of the overlapping CIDR blocks will take effect randomly. Please avoid configuring overlapping CIDR blocks.</p> <p>Note: Priority configuration is not supported for CIFS/SMB file systems and will not take effect.</p>

Step 3. Configure a permission group for a file system

The configuration of a permission group can be modified after the file system is created. You can choose to create a permission group first and select it when creating a file system. You can also select the default permission group when creating a file system and then go to the file system details page to change the permission group.

Note:

If the file system is mounted with the NFS v4 protocol, the modification to the permission group rules of the file system will take effect in 2 minutes.

Step 4. Modify the information and rules of a permission group

You can enter the permission group details page to modify the name, remarks, and rules of a permission group.

Note:

Permission group rules take effect asynchronously. Therefore, avoid adding individual IPs frequently.

We recommend you add a CIDR block or batch import IPs using a template.

Using Tags

Last updated : 2024-01-22 22:15:47

Overview

A [tag](#) is a key-value pair provided by Tencent Cloud to identify a resource in the cloud. You can use tags to categorize and manage cloud resources. Tencent Cloud will not use your tags. They are only used by you to manage your Tencent Cloud resources. This document describes how to edit CFS tags.

Use limits

Note the following limits when editing tags:

Quantity: Each file system allows up to 50 tags.

Tag key: A tag key can only contain digits, letters, and `+ = . @ -` . It cannot exceed 255 characters.

Tag value: A tag value can only contain empty strings or digits, letters, and `+ = . @ -` . It cannot exceed 127 characters.

Use cases and directions

Use case: An organization has purchased six file systems, of which the business group, scope, and owners are as follows:

File System ID	Business Group	Business Scope	Owner
cfs-abcdef1	E-commerce	Marketing campaigns	John Smith
cfs-abcdef2	E-commerce	Marketing campaigns	Chris
cfs-abcdef3	Gaming	Game A	Jane Smith
cfs-abcdef4	Gaming	Game B	Chris
cfs-abcdef5	Entertainment	Post-production	Chris
cfs-abcdef6	Entertainment	Post-production	John Smith

Taking cfs-abcdef1 as an example, we can add the following three sets of tags to the file system:

Tag Key	Tag Value
---------	-----------

dept	ecommerce
business	mkt
owner	John Smith

Similarly, you can add tag key-value pairs to other file systems based on the business group, scope, and owners.

Directions

Tagging a new file system

Prerequisites

Make sure you have a tag that you can use before creating a file system. If you do not have any tags, go to the [Tag console](#) and create one.

1. Log in to the [CFS console](#).
2. On the file system list page, click **Create**.
3. In the pop-up window, find **Tag** down below, and click **Add** to add a tag to the file system. Only existing tags can be added in this step.

✓ Select File System Type

>

2 Set Up Details

Storage Class

Standard

File System Name

Please enter no more than 64 Chinese characters, alphabets, numbers unde

Region

Guangzhou

Availability Zone

Guangzhou Zone 6

To decrease access latency, it's recommended that file system be in the same region with your CVM.

Protocol ⓘ

NFS

Data Source ⓘ

☐ Use a snapshot

Select Network

Select VPC

Please select a subnet

Permission Group

default (pgroubasic)

Permission group specifies a visiting allowlist with some permissions.[How to create?](#)

Scheduled Snapshot

☐ Set scheduled snapshots for the file system Recommended

Snapshot helps to recover data lost due to accidental deletion or virus.

Tag ⓘ

Tag Key

Tag Value

×

+ Add

Back

Create now


4. Click **Submit**, and the tag will be bound to the created file system.

Note:

For more information on creating file systems, see [Creating File Systems and Mount Targets](#).


Adding, modifying, or deleting a tag on an existing file system

1. On the file system list page, locate the file system and click **Edit Tag** in the **Operation** column.

Create		CFS client assistant							Separate multiple keywords by a v	
ID/Name	Monitor	Status ▾	Used/Total ↕	IA storage	Maximum Throughput ⓘ	Availabil ▾	VPC ID/CCN ID	IP		
cfs-4812ee604 frank		Available	0MiB/20TiB	0MiB	2GiB/s	Guangzhou...	vpc-8twhugxt	11.0.0.21		

2. Add, modify, or delete a tag as needed in the pop-up window.

Edit Tag ✕

Tags are used to manage resources by category in different dimensions. If the existing tags don't meet your requirements, you can [manage tags](#) .

1 resource(s) selected

Tag Key ▾

Tag Value ▾ ✕

[+ Add](#)

OK

Cancel

3. Click **OK**.

Filtering file systems by tag

To filter file systems by tag, follow the steps below:

1. Click the search box and select **Tag** from the drop-down menu.
2. Enter tag key and tag value, and click

 , or press Enter to search.

For example, if you want to find file systems tagged with a k1 tag key and v2 tag value, do a search for Tag:k1:v2.

File System

Guangzhou(18)

CFS's High-Performance Turbo is now available. [See Specifications](#)


CFS offers storage resource units available for all sub-products. [Package details](#)

Create

CFS client assistant

Tag: k1:v2

Separate multiple keywords by a vertical bar "|" and enter key.

ID/Name	Monitor	Status	Used/Total	IA storage	Maximum Throughput	Availability	VPC ID/CCN	Subnet	Storage Class
1 result found for "Tag:k1:v2" Back to list									
cfs-n8qmlol9 未命名		Available	0MiB/160TiB	-	100MiB/s	Guangzhou...	vpc-mkuolnm7	10.206.0.23	Standard

Total items: 1

20

Snapshot Management

What Is a File System Snapshot?

Last updated : 2024-01-22 22:15:48

Overview

Real-time replica of online data

Snapshots are fully usable copies of file systems. When a problem occurs in a file system for which a snapshot has been created, you can use the snapshot to quickly restore the file system to normal status. We recommend you create a snapshot of the file system before making any major changes to your business, 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 file systems for quick service deployment.

How snapshots work

A file system snapshot is a block-level clone or backup. In general, the snapshot size will be larger than the data size displayed in the file system because:

The underlying data block stores the metadata of the file system.

Deleting data modifies the blocks where data is written, which will be backed up to snapshots.

Use cases

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 critical 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 the operating system, 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 use an existing snapshot to create one or more file systems to quickly deploy the same business environment in batches, saving the time of repeated configuration.

Billing

See [Snapshot Billing Mode > Pay-As-You-Go \(Postpaid\)](#).

Quota limits

For details on snapshot quota limits, see [Use Limits](#).

Snapshot types

Manual snapshot

You can manually create a snapshot for a file system at a certain point in time. This snapshot can be used to create more file systems with identical data. For more information, see [Creating Snapshots](#).

Scheduled snapshot

If your business data changes frequently, you can use scheduled snapshots to provide continuous data backups. To achieve continuous backups of file system data over a certain time period, you only need to configure a backup policy and associate it with file systems, 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 all file system data. Refer to [Notes](#) to ensure snapshot data consistency.

Case review

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

Customer A has never created a snapshot for its file system. In May 2019, an operator performed a fio test on the file system, resulting in corruption of the file system. The data was permanently damaged.

Analysis: If customer A has created a snapshot for the file system before testing, the snapshot can be used to create a new file system and resume business immediately after the data corruption.

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

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

Analysis: If customer B has created scheduled snapshots for this file system, 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 file system 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 file systems. 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 file system exception occurs, resulting in major loss.

In situations where **a snapshot has been created**, data can be recovered when a server or file system 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.

Use Limits

Last updated : 2024-01-22 22:15:48

This document describes the use limits of Cloud File Storage (CFS) snapshots to help you use this feature more efficiently.

Item	Limit
Number of snapshots	100 per file system
Number of scheduled snapshot policies	30
Number of file systems that can be bound to a scheduled snapshot policy	200
Supported file system type	Standard/High-Performance

File system snapshots can be created by using either of the following methods: Copy-On-Write (COW) and Redirect-On-Write (ROW).

If the snapshot of a file system is **Migrating**, the snapshot has been taken (metadata captured and original file system data marked), and the migration is taking place.

During snapshot migration, the file system can be used properly. You can overwrite, modify, and delete files in the file system without affecting snapshot data migration or causing snapshot file loss.

Creating the first snapshot for a file system may take a while as all data in the file system needs to be migrated, and follow-up snapshots are faster to create, because only incremental data is migrated.

During snapshot migration, the I/O performance of the file system may decrease by about 15%. You are advised to configure proper scheduled snapshot policies, so that snapshots are taken during off-peak hours of your business.

Creating a Snapshot

Last updated : 2024-01-22 22:15:48

Overview

You can create snapshots for a file system to save its data at specific points of time. The file system snapshot feature adopts the incremental mode to create snapshots, and it creates a snapshot that records only the data changes compared with the last snapshot. This process is quick if the data changes a little. Although snapshots are created in incremental mode, if there are snapshots available, you can always use an undeleted snapshot to restore the data at the time when the snapshot is created.

You can create a snapshot of a file system in the normal state, but the snapshot can only capture the written data rather than data being written by an app or process. Therefore, according to your business needs, you can choose to temporarily stop all data writes and create a snapshot after data synchronization to obtain a complete snapshot.

Prerequisites

You have activated CFS.

Note:

Currently, only Standard and High-Performance file systems support the snapshot feature.

You have not reached the upper limits for the number and total size of snapshots in the current region. For more information, see [Use Limits](#).

Notes

A snapshot can only capture the written data but not cached data in the memory (such as files in the `/run` directory in a Linux system) of a file system. We strongly recommend that you ensure that all data in the memory is written to the file system and suspend read and write operations of the file system before creating a snapshot. The recommended method for writing data in the memory to the file system is as follows:

For better system performance, data is stored in the memory buffer before it is written to the file system at a proper moment. Therefore, the snapshot created for the file system does not contain data that is stored in the memory buffer and is not written to the file system. As a result, data inconsistency occurs.


To resolve this issue, run the `sync` command to forcibly write the data in the memory buffer immediately to the file system, and do not write new data to the file system before creating a snapshot. If no error message is returned after the command is executed, the data in the memory buffer has been successfully written to the file system.

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

Directions

1. Log in to the [CFS console](#).
2. Click **Create Snapshot** in the row of the target file system.
3. In the pop-up window, enter a snapshot name and click **OK**.

Create Snapshot



1. Snapshots only capture point-in-time data stored in the file system but not the memory. Therefore, to capture all data, you are advised to sync data cached in memory to disk before creating a snapshot.

File System ID

cfs-fkrk4rdn | bruce ▼

File System Name

bruce

Capacity

0MiB


Type

Standard

Snapshot Name

Enter a snapshot name

60 characters remaining. Supports Chinese characters, letters, and underscores ()

Tag 

[+ Add](#)

Confirm

Cancel

Scheduled Snapshots

Last updated : 2024-01-22 22:15:48

Features

CFS provides the **scheduled snapshot** feature that allows developers to flexibly configure backup task policies.

The following table describes the scheduled snapshot policies recommended for different business scenarios.

Scenario	Snapshot Policy	Snapshot Retention
Core product/service	Use scheduled snapshots, with the policy set to once per day.	7 to 30 days
Non-core and non-data product/service	Use scheduled snapshots, with the policy set to once per week.	7 days
Archive	Scheduled snapshot is not required. You can create snapshots manually whenever needed.	One month to several months
Test	Scheduled snapshot is not required. You can create snapshots manually whenever needed.	Deleted after being used

Policy description

The following table describes the content and features of scheduled snapshot policies, which help you better use snapshots in your businesses.

Item	Description
Object	All file systems
Policy	The point in time for scheduled snapshot creation can be accurate to every hour or every day. A scheduled snapshot policy is valid permanently after being set. If you modify a scheduled snapshot policy, it takes effect immediately.
Scheduled termination (Important)	Scheduled snapshots can be terminated periodically. After you set a snapshot lifecycle (1–30 days), scheduled snapshots are automatically deleted upon expiration, which reduces the backup costs. If you do not set a scheduled termination policy, scheduled snapshots will be stored permanently.
Batch	You can select multiple file systems and batch apply the same scheduled snapshot policy to them.

Naming rule	Scheduled snapshots are named in the format of "auto_policy_cfsidyyyyMMddHH", where "cfsid" is the file system ID, "yyyyMMdd" is the snapshot creation date, and "HH" is the hour. You can also modify snapshot names. For example, "auto_policy_cfs-2cj5yj0f2021090923" indicates a snapshot that was automatically created on 2021-09-09 23:00:00 for the file system cfs-2cj5yj0f.
Lifecycle (Important)	Snapshot lifecycles vary depending on the snapshot creation method: Manually created snapshots are permanently stored by default as long as the account balance is sufficient. Scheduled snapshots can be terminated periodically or permanently stored.
Snapshot conflict	Scheduled snapshots do not conflict with custom snapshots in use. However, they may conflict with each other on the creation time. When a snapshot is automatically created for a file system, you can create a custom snapshot only after automatic snapshot creation is complete, and vice versa. If there is a large amount of data in the file system and the time required for automatically creating a snapshot exceeds the interval between two time points for automatic snapshot creation, a snapshot is not automatically created at the next time point. For example, if you set 9:00, 10:00, and 11:00 as the time points for scheduled snapshot creation and it takes 70 minutes to create the snapshot at 9:00 (that is, the creation of this snapshot will be completed at 10:10), the next snapshot will not be created at 10:00, but at 11:00.
Snapshot quota	Each file system has a certain snapshot quota. If the number of snapshots on a file system has reached the quota limit, the scheduled snapshot task will be suspended or blocked. The snapshot quota is used to prevent continuous increase in the storage costs if developers forget a scheduled snapshot policy.
ASP	Indicates the scheduled snapshot policy, that is, Auto Snapshot Policy.
ASP quota	A single Tencent Cloud account can set a maximum of 30 ASPs in each region. An ASP can be associated with a maximum of 200 file systems.
Retention period	The console displays the repossession countdown for scheduled snapshots. You can manually change the retention period of scheduled snapshots to permanent. Manually created snapshots are permanently stored.
ASP pause	An ASP can be manually paused . After it is paused, snapshots are no longer automatically created. However, lifecycles of previously created snapshots are not affected. They will be periodically terminated or permanently stored based on configured rules.
Operations log	Displays the creation process of all scheduled snapshots, same as that of manually created snapshots.

Directions

Creating a scheduled snapshot policy

Note

A single Tencent Cloud account can set a maximum of 30 scheduled snapshot policies in each region.

1. Log in to the CFS console and go to the [Snapshot Policies](#) page.
2. Select a region.
3. Click **Create Policy**.
4. In the pop-up window, set parameters as needed and click **Confirm**.

Create Policy

i 1. The file system must be in Available status.

2. Snapshots only capture point-in-time data stored in the file system but not the memory. Therefore, to capture all data, you are advised to sync data cached in memory to disk before creating a snapshot.

3. If the data volume is too large to finish creating a snapshot before the next snapshot should start, the next one will be skipped. For example, you have scheduled snapshots to be created at 22:00, 23:00, and 01:00 (UTC+08:00), and the snapshot starting at 22:00 takes 85 minutes (finishes at 23:25). In this case, the snapshot scheduled at 23:00 will be skipped and the next snapshot time will be 01:00.

Name

Region

Shanghai

Backup frequency

Weekly

Monthly

Backup Day

☐ Every Mon

☐ Every Tues

☐ Every Wed

☐ Every Thur

☐ Every Fri

☐ Every Sat

☐ Every Sun

Backup Time

☐ 00:00

☐ 01:00

☐ 02:00

☐ 03:00

☐ 04:00

☐ 05:00

☐ 06:00

☐ 07:00

☐ 08:00

☐ 09:00

☐ 10:00

☐ 11:00

☐ 12:00

☐ 13:00

☐ 14:00

☐ 15:00

☐ 16:00

☐ 17:00

☐ 18:00

☐ 19:00

☐ 20:00

☐ 21:00

☐ 22:00

☐ 23:00

Snapshot Retention

☒ Delete after

30

day(s)

☐ Permanent

First Backup Time

Advanced settings

The actual snapshot time might be different from the scheduled one.

Confirm

Cancel

Parameter	Description
Name	Required. The name of a scheduled snapshot policy. It can contain up to 60 characters.
Region	Required. This parameter cannot be changed on the current page. For more information on how to set this parameter, see Step 2 .
Backup Day	Required. The days to automatically create snapshots. Value range: Every Sun to Every Sat
Backup Time	Required. The time to automatically create snapshots. Value range: 00:00 - 23:00. Due to server performance issues, the actual scheduled snapshot creation time may be slightly different from the time set here. In this case, data in a snapshot is based on the actual creation time.
Snapshot Retention	Required. Retained for a fixed number of days. The value range is 1 to 30 days. The default value is 30 days. Permanent

Assigning a file system

Note

A scheduled snapshot policy can be associated with a maximum of 200 file systems.

1. Log in to the CFS console and go to the [Snapshot Policies](#) page.
2. Select a region.
3. Click **Assign File System** in the row of the target policy.
4. In the **Assign File System** pop-up window, select the file systems to assign.

Assign File System

You have selected **1 snapshot policy**. [Collapse](#)

Snapshot Policy Name	Snapshot Policy ID	Assigned File Systems
aileen-test	asp-o8qe5vyv	0

Select file systems

Enter an ID/name Q

<input type="checkbox"/> ID/Name	Capacity	Scheduled Snapsho...
<input type="checkbox"/> cfs-4da614276 aileen-test	0MiB	-

Selected (0)

ID/Name	Capacity
---------	----------

↔

Confirm

Cancel

5. Click **Confirm**.

Enabling or disabling a scheduled snapshot policy

1. Log in to the CFS console and go to the [Snapshot Policies](#) page.
2. Select a region.
3. Find the target policy and click the button in the **Scheduled Snapshot** column to enable or disable the scheduled snapshot policy.

Modifying a scheduled snapshot policy

1. Log in to the CFS console and go to the [Snapshot Policies](#) page.
2. Select a region.
3. Find the target policy and click **Modify Policy** in the **Operation** column.
4. In the **Modify Policy** pop-up window, modify related parameters and click **Confirm**. For more information about parameter descriptions, see [Step 4](#).

Deleting a scheduled snapshot policy

1. Log in to the CFS console and go to the [Snapshot Policies](#) page.
2. Select a region.
3. Delete a scheduled snapshot policy as follows:
Click **More** > **Delete Policy** in the row of the target policy.

Setting permanent storage for scheduled snapshots

Note



If **Snapshot Retention** is set to **Permanent** for a scheduled snapshot policy, you do not need to perform the following operations for snapshots automatically created using this policy.

1. Log in to the CFS console and go to the [Snapshot Policies](#) page.
2. Select a region.
3. Click the ID of the target snapshot.
4. On the details page, click **Permanent** to permanently store the snapshot, and **Snapshot Retention** of the snapshot is changed to **Permanent**.

[←](#) **snapcfs-4057aa9fd** (aileen-test-cross)

Basic Info Operation Logs

Snapshot Information

Name	aileen-test-cross 
ID	snapcfs-4057aa9fd
Region	Shanghai
File System	cfs-45a6fb370 (brucetest)
Snapshot Size	Go to Snapshot Chain .
Retention	Permanent
Creation Time	2023-06-19 16:42:38
Tag	

Creating a File System Using a Snapshot

Last updated : 2024-01-22 22:15:48

Overview

Snapshots are an important way of data sharing and migration. File systems created using a snapshot own all data in the snapshot. You can use a snapshot to create a file system (to ensure data security, only new file systems can be created from snapshots).

This document describes how to use a snapshot to create a file system on the snapshot list page. When [creating a file system](#), you can configure the **Snapshots** parameter to specify a snapshot for creating the file system.

Directions

1. Log in to the CFS console and go to the [Snapshot List](#) page.
2. Click **Use** in the row of the target snapshot.
3. On the file system creation page, select the file system type, click **Next: Set Up Details**, and set the following fields.

Note:

The following describes the fields for creating a file system of Standard type as an example.

Field	Description
Billing Mode	Select a desired billing mode. Two billing modes are supported: pay-as-you-go and prepaid. Note: Only some products support the prepaid mode.
File System Name	Customize the name of the file system to create.
Region	Select the region where the file system is to be created.
Availability Zone	Select the availability zone where the file system is to be created.
Protocol	Select a protocol type, NFS or SMB, for the file system. NFS is more suitable for Linux and Unix clients and CIFS/SMB for Windows clients. The Turbo series can only be used by private clients and does not allow the selection of file system protocols.
Data Source	Optional. To create a file system using a snapshot, you need to select Use a snapshot and select the snapshot you want to use. If you create a file system using a snapshot, the initial data amount on the file system is the same as the snapshot size.

Permission Group	Each file system must be bound to a permission group. The permission group specifies an allowlist of the users that can access the file system and lists the read and write permissions.
Scheduled Snapshot	Optional. When creating a file system, you can select Scheduled Snapshot to create snapshots for the file system periodically based on the created scheduled snapshot policy. For more information, see Scheduled Snapshots .
Storage Capacity	Required only for the Turbo series. The Turbo series is an exclusive cluster and has restrictions on the minimum cluster scale and expansion step. For Standard Turbo, the minimum initial cluster scale is 40 TiB, and the expansion step is 20 TiB. For High-Performance Turbo, the minimum initial cluster scale is 20 TiB, and the expansion step is 10 TiB.
CCN	Required only for the Turbo series. Select an existing CCN instance or create a new one. For more information, see Cloud Connect Network .
IP Range	Required only for the Turbo series. This parameter allows you to reserve a CIDR block for Turbo related components. Ensure that the selected CIDR block does not conflict with those of other instances on the cloud that need to communicate with Turbo. To ensure the number of IP addresses in the CIDR block, the mask must have 16 to 24 bits, for example, 10.0.0.0/24.
Tag	If you already have a tag, you can add it to the new file system here. If you do not have a tag, log in to the Tags console to create a desired tag, and then bind it to the file system. You can also add a tag to the file system after it is created.

4. Click **Create Now**. You can see the newly created file system in the file system list.

Deleting a Snapshot

Last updated : 2024-01-22 22:15:48

Overview

When there is no need to use a snapshot again, you can delete it to release virtual resources.

Notes

When you delete a snapshot, only the data exclusive to the snapshot will be deleted, and the file system for which the snapshot is created will not be affected.

You can use a snapshot to restore a file system to the data status when the snapshot is created. Deleting a snapshot created earlier for a file system will not affect your use of snapshots created later.

Deleting a snapshot will also permanently delete all data in the snapshot. Deleted snapshots cannot be restored. Please proceed with caution.

Directions

1. Log in to the CFS console, and go to the [Snapshot Overview](#) page to view the snapshot quota usage in regions.

Snapshot Overview Guangzhou(22) ▼

- **Data protection:** keeps you safe from data loss due to accidental deletion or virus attacks. You can also restore your file system when problems occur during business upgrades or migration.
- **Easy management:** allows you to customize scheduled snapshot policies for different businesses. You can assign a Standard file system to a scheduled snapshot policy to create and delete snapshots periodically for the
- **Low cost:** except for the first snapshot, Tencent Cloud creates incremental snapshots that save only the data that has changed since the previous snapshot, to save storage space and reduce costs.
- This page displays all the statistics available for the root account.

Region	Total Snapshots	Total Snapshot Size
Guangzhou	22	13.41TiB

Snapshot Monitoring

2. On the [Snapshot List](#) page, click **Delete Snapshot** for the target snapshot.

Snapshot List

Shanghai(3)

Delete Snapshot

Create Snapshot

Separate multiple keywords by a vertical bar "|"

<input type="checkbox"/>	ID/Name	Status	Source File System ID/...	Snapshot Size	Snapshot type	Creation Time	Retention
<input type="checkbox"/>	snpcfs-4057aa9fd aileen-test-cross	Normal	cfs-45a6fb370 brucetest	0MiB	Turbo	2023-06-19 16:42:38	Permanent
<input type="checkbox"/>	snpcfs-465a1c964 aileen-test	Normal	cfs-4ece9d5d4 aileen-Test	0MiB	Turbo	2023-05-23 19:00:45	Permanent
<input type="checkbox"/>	snpcfs-m5so6k2f 宝物	Normal	cfs-lllt2hbx aileen-test	0MiB	Standard	2022-09-27 17:41:04	Permanent
<input type="checkbox"/>	snpcfs-iwl1x7h frank	Normal	cfs-ltyor6yd frnak	0MiB	Standard	2022-06-15 18:20:31	Permanent
<input type="checkbox"/>	snpcfs-m5lmh445 frank	Normal	cfs-ltyor6yd frnak	0MiB	Standard	2022-06-15 18:18:36	Permanent

Total items: 5

20

3. Confirm the information of the snapshot to be deleted:
- 3.1 Confirm the snapshot to be deleted, and click **Next**.

Delete Snapshot

1 List Details > 2 Confirmation > 3 Before You Start

You have selected 1 snapshot. [Collapse](#)

Snapshots	File System	Snapshot Time
snapcfs-4057aa9fd aileen-test-cross	cfs-45a6fb370 brucetest	2023-06-19 16:42:38

⚠ Deleted snapshots cannot be recovered.

☒ I understand that the deletion cannot be undone and the deleted data cannot be recovered.

Next

Cancel

3.2 Confirm the information, and click **Next**.

Delete Snapshot

✓ List Details > 2 Confirmation > 3 Before You Start

Resources to delete:

You have selected 1 snapshot. [Collapse](#)

Snapshot Name	Snapshot ID	File System	Snapshot T
aileen-test-cross	snapcfs-4057aa9fd	cfs-45a6fb370 brucetest	2023-06-19

Back

Next

Cancel

3.3 Read the precautions and click **Confirm** to execute the deletion.

Delete Snapshot

✓ List Details > ✓ Confirmation > 3 Before You Start



Note:

- Deleted data cannot be recovered. Ensure that you have backed it up.

Back

Confirm

Guide for Cross-AZ and Cross-Network Access

Last updated : 2024-01-22 22:15:47

Cross-AZ access in a VPC

If you need to have a Cloud File Storage (CFS) instance shared by multiple Cloud Virtual Machine (CVM) instances in different AZs in the same region, you can configure the CVM and CFS instances into the same VPC to achieve cross-AZ access to resources.

Assume that there is a CVM instance in Guangzhou Zone 1 that needs to use CFS, but file systems cannot be directly created in Guangzhou Zone 1 as the resources are sold out there.

If the CVM instance is in the "Guangzhou Zone 1" subnet of a VPC, you can log in to the [VPC console](#) to create a subnet whose AZ is "Guangzhou Zone 3" for the VPC.

Create a subnet

Network vpc-b9o4hppj(vpc-test-cfs-... | 192.168.0.0/16) ▼ 1 existing subnets

Subnet name	VPC IP range	CIDR ⓘ	Availability zone ⓘ
<input type="text" value="test"/> 4/60	<input type="text" value="192.168.0.0/16"/> ▼	192 . 168 . <input type="text" value="0"/> . 0 / <input type="text" value="24"/> ▼	<input type="text" value="Please select"/> ▼
+ New line			

[Advanced options ▶](#)

After the subnet is successfully created, go back to the CFS console and select this VPC and the subnet you just created to create resources in Guangzhou Zone 3. The CFS file system can be directly mounted to the CVM instance in the subnet of Guangzhou Zone 1 in this VPC. For more information, see [Using CFS File Systems on Linux Clients](#).

Cross-VPC and cross-region access

CFS supports the following scenarios for resource access.

A CFS instance needs to be shared by multiple CVM instances distributed in different VPCs.

Your CVM and CFS instances are in different VPCs.

Your CVM and CFS instances are in different regions (for better access performance, it is recommended that the CVM instances and CFS be in the same region).

You can use Cloud Connect Network (CCN) to achieve resource access across the CVM instances distributed in VPC-A/VPC-B and CFS instance distributed in VPC-C. For more information, see [Connecting Network Instances Under the Same Account](#).

Automatically Mounting File Systems

Last updated : 2024-01-22 22:15:48


Overview

You can configure a Linux or Windows client to which a Cloud File Storage (CFS) file system is mounted, so that the file system can be automatically mounted after the client is restarted.

Directions


Automatically mounting an NFS file system on Linux

1. Connect to the CVM instance that needs to automatically mount the file system by logging in to the CVM console or performing remote login. Then, open the "/etc/fstab" file (make sure that your login account has the root privileges).



```
// Run the following command to open the "fstab" file  
vi /etc/fstab
```


2. Enter "i" (insert) and add the following command to `/etc/fstab` . The mounting methods are as follows:



Mount the file system with NFS v4.0

<Mount target IP>:/ <target mount directory> nfs vers=4,minorversion=0,hard,timeo=6


Example: 10.10.19.12:/ /local/test nfs vers=4,minorversion=0,hard,timeo=600,retrans



Mount the file system with NFS v3.0

<Mount target IP>:/<fsid> <target mount directory> nfs vers=3,nolock,proto=tcp,hard

Example: 10.10.19.12:/djoajeo4 /local/test nfs vers=3,nolock,proto=tcp,hard,timeo=6



Mount the file system with Turbo

```
<Mount target IP>@tcp0:/<fsid>/cfs <target mount directory> lustre defaults,_netdev
```

```
Example: 172.16.0.7@tcp0:/01184207/cfs /root/turbo lustre defaults,_netdev 0 0
```

3. Press Esc and enter ":wq" to save the change. Restart the client, with the file system automatically mounted.

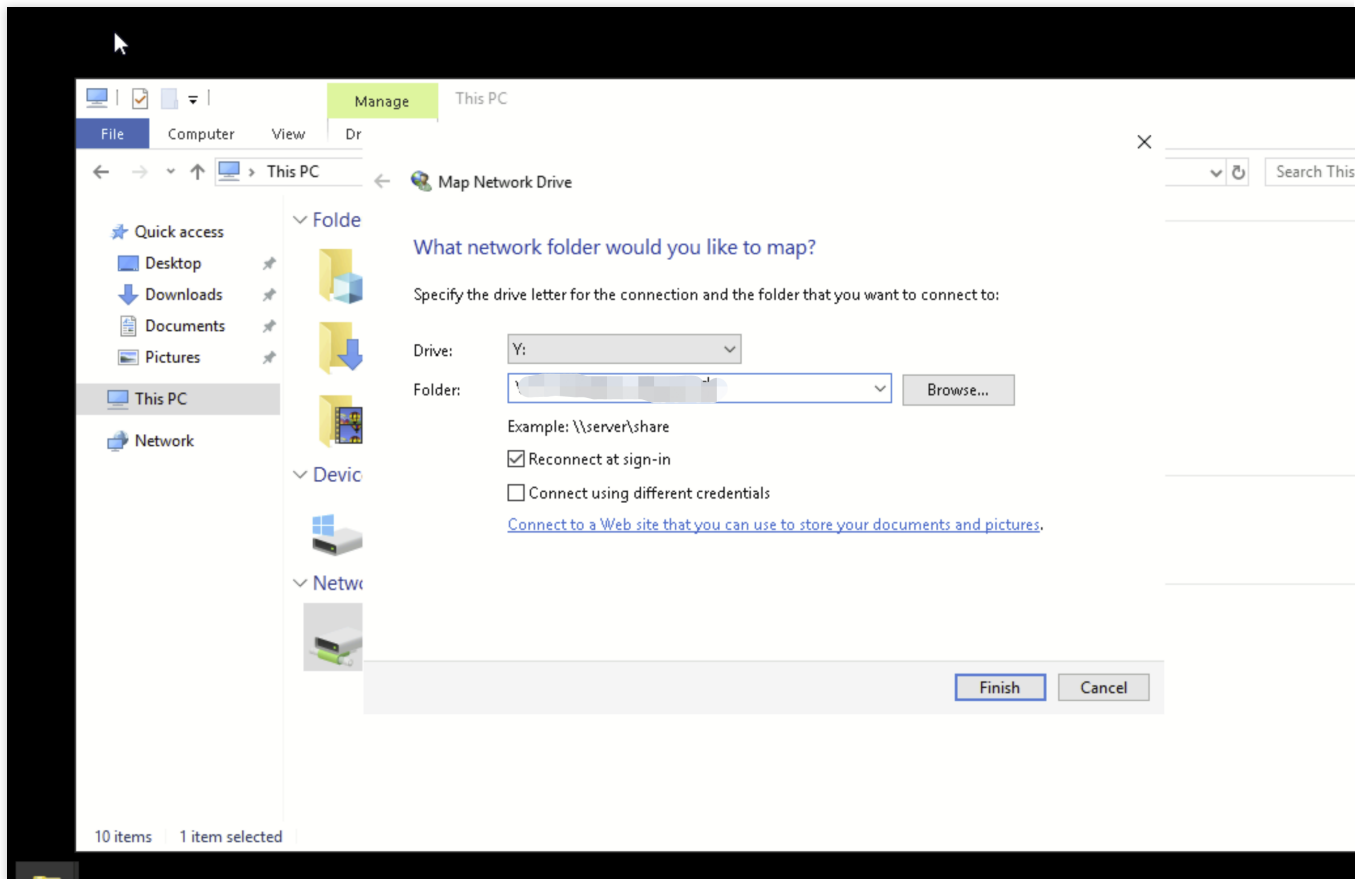
Note:

When the command for automatic mounting is added, but the shared file system is abnormal, Linux may not be started normally, because the automatic start command in `fstab` is not executed. To solve this problem, enter "Single

User Mode" upon startup, delete the automatic mounting command in `fstab` , and then restart the server.

Automatically mounting a file system on Windows

When mounting a file system, select "Reconnect at login". For more information, see [Using CFS File Systems on Windows Clients](#).



Data Migration Service

Features

Last updated : 2024-01-22 22:15:47

Overview

The data migration service is an online concurrent migration service provided by Cloud File Storage (CFS) for migration of massive amounts of data. It enables you to easily migrate data from sources such as object storage systems of Tencent Cloud and other mainstream providers to CFS, making your data transfer more efficient.

Use cases

Data loading

In machine learning and autonomous driving scenarios, certain datasets are originally stored in object storage. When frequent access is required for training and inference, data can be quickly loaded to High-Performance CFS through the migration service for more efficient data reads.

Cross-cloud data migration

Object storage can be directly accessed over the public network, which means file data from other clouds can be transferred to object storage and then to CFS through the data migration service.

Cross-account data migration

When CFS data is migrated from one account to another, the cross-account access capabilities of object storage can be used to implement cross-account data migration.

Advantages

Supports data migration from mainstream cloud services.

Supports migrating COS data to CFS, where all download traffic is over Tencent Cloud's private network, without incurring traffic fees.

Supports migrating data from mainstream object storage services to CFS.

Various data migration methods.

Bucket migration: Migrates all data according to the specified bucket path.

Inventory migration: Filters objects by time period and file prefix (specified directory) based on the object storage inventory list for finer-grained migration.

Supports overwriting by last modified time, full overwriting, and no overwriting.

Inventory migration: Filters objects by time period and file prefix (specified directory) based on the object storage inventory list for finer-grained migration.

More secure and efficient data transfer methods

Migration monitoring: Provides the real-time information of the running status, traffic, and progress of migration tasks.

Migration details: Provides detailed migration information such as the number of files, file size, file list, and status.

Automatic retry: Three automatic retries are performed in the case of various temporary errors to enhance the migration efficiency.

Use Limits

Last updated : 2024-01-22 22:15:48

The following table describes the use limits of the data migration service.

Item	Limit
Migration source	Data can be migrated only from object storage, including mainstream object storage services in the Chinese mainland.
Migration destination	Data can be migrated only to Cloud File Storage (CFS), which can be Standard (NFS), High-Performance (NFS), Standard Turbo, or High-Performance Turbo.
Progress display	Currently, the migration progress is updated once for every 1,000 files. If fewer than 1,000 files are migrated, the progress may deviate to some extent.

Note:

The data migration service is free of charge. If data migration crosses regions or clouds, the outbound traffic of the source object storage will be charged. No outbound traffic fees will be incurred only when data is migrated from COS to CFS in the same region.

If the source data is large in volume, you can use bucket migration or inventory migration. Specifically, use different bucket directories and inventories for different tasks to accelerate the migration. We recommend you run up to three tasks. Do not start multiple identical tasks for the same batch of data, which will lead to unnecessary metadata verification and prevent acceleration.

Starting a Migration Task

Last updated : 2024-01-22 22:15:47

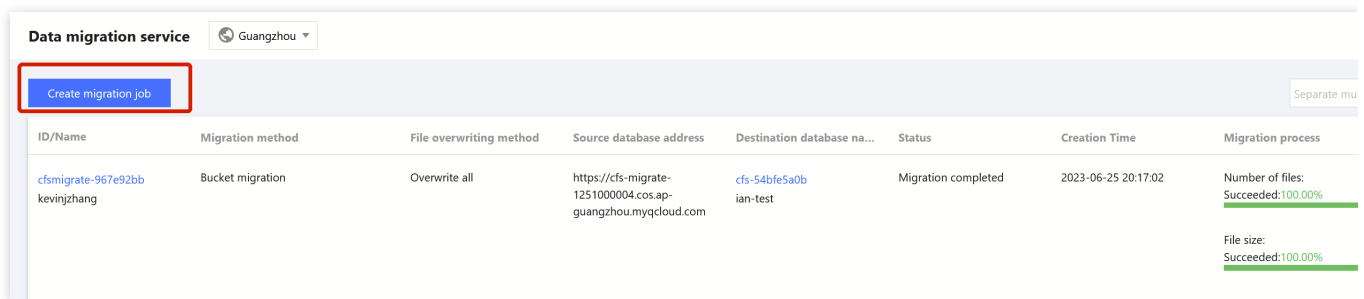
Overview

This document describes how to start a migration task.

Directions

Step 1. Create a migration task

1. Log in to the CFS console and go to the [Data Migration Service](#) page.
2. Click **Create migration job**.



Data migration service Guangzhou							
Create migration job Separate mu							
ID/Name	Migration method	File overwriting method	Source database address	Destination database na...	Status	Creation Time	Migration process
cfsmigrate-967e92bb kevinjzhang	Bucket migration	Overwrite all	https://cfs-migrate-1251000004.cos.ap-guangzhou.myqcloud.com	cfs-54bfe5a0b ian-test	Migration completed	2023-06-25 20:17:02	Number of files: Succeeded:100.00% File size: Succeeded:100.00%

Note:

Your account needs to be added to the allowlist before you can use this feature. If you need it, [submit a ticket](#) for application.

Step 2. Select the region for migration

Select the destination CFS region as the region for migration.

Step 3. Select the source service provider

Select **Tencent Cloud COS**.

Source
database

Tencent Cloud COS

Alibaba Cloud OSS

Hua

Step 4. Select the migration method

Currently, bucket migration and inventory migration are supported:

Migration
method

Bucket migration

Inventory migration

Bucket migration: It is suitable for migrating an object storage bucket or all objects in a certain directory.

Inventory migration: It is suitable for migrating objects within a specified time period. Currently, you can export the objects modified after a time point or within a time period through the object storage inventory for filtered data migration.

Step 5. Select the overwriting method

File
overwriting
method

LastModified

Overwrite all

Do not overwrite

For a source file has the same name as the destination file, the system compares the time indicated by "LastModified".

1. If the "LastModified" of the source file is earlier than that of the destination file, the source file will not be migrated and the destination file will be retained.
2. If the "LastModified" of the destination file is earlier than that of the source file, the source file will be overwritten.
3. If the two files have the same "LastModified", the system compares again:
 - If their sizes are different, the destination file will be overwritten.
 - If their sizes are the same, the source file will not be migrated.

Currently, overwriting by last modified time, full overwriting, and no overwriting are supported.

Overwriting by last modified time: When files with the same name exist, their `LastModified` values are checked, i.e., last modified time.

If the `LastModified` value of the file in the source address is earlier than that of the file in the destination address, no overwriting is performed.

If the `LastModified` value of the file in the source address is later than that of the file in the destination address, overwriting is performed.

If the two files have the same `LastModified` value, further check:

If their `Size` values are different, overwriting is performed.

If their `Size` values are the same, no overwriting is performed.

Full overwriting: Files with the same name are overwritten without any check.

No overwriting: In contrast to **full overwriting**, this method means no overwriting of files with the same name without any check.

Step 6. Enter the `SecretId` and `SecretKey`

Enter the `SecretId` and `SecretKey` required by access to object storage, which are encrypted on the platform. If you use a temporary key, make sure that the migration task can be completed during the validity period of the key; otherwise, part of the migration will fail.

SecretId	<input type="text"/>
SecretKey	<input type="text"/>

Step 7. Enter the source information (required for bucket migration only)

Currently, you can either select a bucket or enter the source bucket address (object storage endpoint) in the console. For data transfer from the root path, make sure that the source path is empty. For data transfer from a subpath, the source path should end with "/".

Source database address

☒ Select an available source bucket ☐ Enter a source bucket URL

Select a bucket

Source database path

Leave the source database path empty if data in a root path is migrated. Sub-path slash (/).

Step 8. Enter the inventory URL (required for inventory migration only)

Enter the URL of the `manifest.json` file of the object storage inventory. Specifically, after an inventory is generated, find the URL of the `manifest.json` file and enter it as the inventory URL.

11234-10000037 / httpd-2.4.38.tar.gz

Details Extraction

Basic Info

Object Name

Object Size

Modification Time

ETag

Specified Domain

Object Address

Default Endpoint

Want to preview files in the browser without download? You can do so by configuring the correct Content-Type for the files. For details, see [FAQ - Upload and Download](#).

After the object address is accessed, there will be request and traffic fees. For fee details, please see [Billing Description](#).

Temporary Link

Copy Temporary Link

Download Objects

Refresh

The temporary link carries the signature parameter, and can be used to access the object during the validity period of the signature, the signature is valid for 1 hour (2023-07-17 22:20:02).

Be sure to avoid leaking the temporary link, otherwise your objects may be accessed by other users.

Step 9. Select the destination file system

Select the destination file system and path. For data migration from the root path, leave the destination path empty. If the migration is from the `test` directory, enter either `/test` or `/test/`.

Viewing Migration Task Result

Last updated : 2024-01-22 22:15:48

Overview

This document describes how to view the results of migration tasks.

Prerequisites


You have logged in to the CFS console.

Directions

Viewing results in the data migration service panel

On the [Data Migration Service](#) page, select the target region and view the basic information of a migration task, such as migration settings and progress.

Data migration service

Guangzhou

Create migration job

ID/Name	Migration method	File overwriting method	Source database address	Destination database na...	Status	Creation Time	Migration
cfsmigrate-967e92bb kevinjzhang	Bucket migration	Overwrite all	https://cfs-migrate-1251000004.cos.ap-guangzhou.myqcloud.com	cfs-54bfe5a0b ian-test	Migration completed	2023-06-25 20:17:02	Number of Succeeded: File size: Succeeded:
cfsmigrate-ba8c181e ian4_copy_copy	Bucket migration	Overwrite all	https://cfs-migrate-1251000004.cos.ap-guangzhou.myqcloud.com	cfs-54bfe5a0b ian-test	Migration completed	2023-06-13 15:20:00	Number of Succeeded: File size: Succeeded:
cfsmigrate-8b676718 ian4_copy_copy	Bucket migration	Overwrite all	https://cfs-migrate-1251000004.cos.ap-guangzhou.myqcloud.com	cfs-3c24w3en brucetest	Migration completed	2022-11-17 12:36:06	Number of Succeeded: File size: Succeeded:
cfsmigrate-087eee04 ian4_copy	Bucket migration	Overwrite all	https://cfs-migrate-1251000004.cos.ap-guangzhou.myqcloud.com	cfs-3c24w3en brucetest	Migration completed	2022-10-10 11:59:27	Number of Succeeded: File size: Succeeded:

Viewing detailed information

On the [Data Migration Service](#) page, click the **ID/Name** of the target migration task to view its details, such as the total number/size of files, and the number/size of files to be migrated, migrated successfully, and failed to be migrated.

←

ian4_copy_copy

Basic Info

Basic Info

Region

Guangzhou

Job ID

cfsmigrate-ba8c181e

Job name

ian4_copy_copy

Migration mode

Full migration

Migration method

Bucket migration

File overwriting method

Overwrite all

Source database address

https://cfs-migrate-1251000004.cos.ap-guangzhou.myqcloud.com

Source database path

nodejs_source/

Destination database

ian-test(cfs-54bfe5a0b)

Destination database path

/ian-test

Status

Migration completed

Creation Time

2023-06-13 15:20:00

End time

2023-06-13 16:16:55

Details

Data	Progress	Total Quantity	To be migrated	
Number of files	Succeeded:100.00% <div></div>	2	0	2
File size	Succeeded:100.00% <div></div>	21.01MiB	0B	2
Data inventory	-	Export	-	E

User Permission Management

Last updated : 2024-01-22 22:15:48

This document describes how to set the access permissions of users and user groups based on the POSIX syntax in a file system, which can be Standard (NFS), High-Performance (NFS), Standard Turbo, or High-Performance Turbo.


Prerequisites

You have mounted the file system over the Turbo or NFSv3 protocol as instructed in [Using CFS Turbo on Linux Clients](#) or [Using CFS File Systems on Linux Clients](#).

Command description

Command	Description
<code>getfacl <filename></code>	View the current ACL of the file.
<code>setfacl -m g:cfsgroup:w <filename></code>	Set the write permission for the <code>cfsgroup</code> user group.
<code>setfacl -m u:cfsuser:w <filename></code>	Set the write permission for the <code>cfsuser</code> user.
<code>setfacl -x g:cfsgroup <filename></code>	Delete the permission of the <code>players</code> user group.
<code>getfacl file1 setfacl --set-file=- file2</code>	Copy the ACL of <code>file1</code> to <code>file2</code> .
<code>setfacl -b file1</code>	Delete all extended ACL rules and retain basic ACL rules (owner, group, and others).
<code>setfacl -k file1</code>	Delete all default rules from <code>file1</code> .
<code>setfacl -R -m g:cfsgroup:rw dir</code>	Grant the <code>cfsgroup</code> user group the permission to read/write files and directories in the <code>dir</code> directory tree.
<code>setfacl -d -m g:cfsgroup:rw dir</code>	Grant the <code>cfsgroup</code> user group the permission to read/write newly created files and directories in the <code>dir</code> directory tree.

Example



```
sudo useradd cfsuser # Create the `cfsuser` user
sudo useradd otheruser # Create the `otheruser` user
sudo groupadd cfsgroup # Create the `cfsgroup` user group
sudo usermod -g cfsgroup cfsuser # Allocate `cfsuser` to `cfsgroup`
sudo touch file1 # Create a file named `file1`
sudo setfacl -m g:cfsgroup:r-x file1 # Grant the `cfsgroup` user group the permission
sudo setfacl -m u:otheruser:rwx file1 # Grant the `otheruser` user the permission t
```

User Quotas

Operation Guide

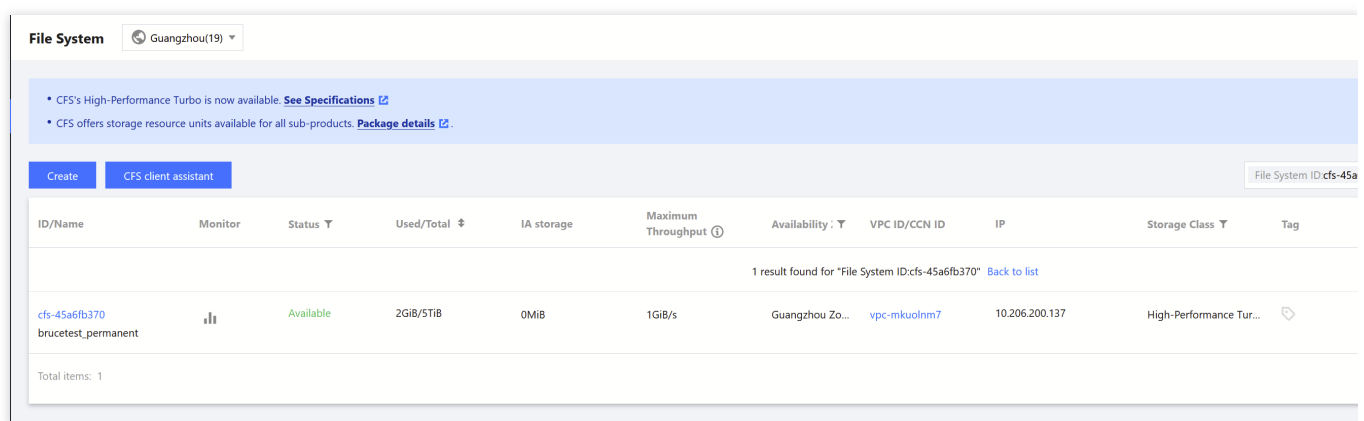
Last updated : 2024-01-22 22:15:47



Setting quotas

1. Go to the quota setting page: Log in to the [CFS console](#), and click **Create quota**.

Note:

Your account needs to be added to the allowlist before you can use this feature. If you need it, [submit a ticket](#) for application.



File System										
Guangzhou(19)										
<ul style="list-style-type: none">CFS's High-Performance Turbo is now available. See SpecificationsCFS offers storage resource units available for all sub-products. Package details										
Create CFS client assistant File System ID: cfs-45a										
ID/Name	Monitor	Status	Used/Total	IA storage	Maximum Throughput	Availability	VPC ID/CCN ID	IP	Storage Class	Tag
1 result found for "File System ID:cfs-45a6fb370" Back to list										
cfs-45a6fb370 brucestest_permanent		Available	2GiB/5TiB	0MiB	1GiB/s	Guangzhou Zo...	vpc-mkuolnm7	10.206.200.137	High-Performance Tur...	
Total items: 1										

2. Create a user quota: Set the quota type to Uid (user ID), enter an ID and the corresponding capacity quota and file count quota, and click **Confirm**.

Note:

Uid is a user ID in the Unix system. If no Uid is available, create one with the useradd command.

Add user/user group quota ×

File system instance

cfs-45a6fb370

Quota type *

Uid

ID *

Capacity quota *

—

10

+

GiB

Quota of files *

—

10000

+

Confirm

Cancel

3. Create a user group quota: Set the quota type to Gid (group ID), enter an ID and the corresponding capacity quota and file count quota, and click Confirm.

Note:

Gid is a group ID in the Unix system. If no Gid is available, create one with the groupadd command.

Add user/user group quota ×

File system instance

cfs-45a6fb370

Quota type *

Gid

ID *

Capacity quota *

—

10

+

GiB

Quota of files *

—

10000

+

Confirm

Cancel

Viewing quotas

Click the ID/name of a file system to go to its details page, and select **Quota information** to view quotas.

[←](#) **cfs-45a6fb370** (brucetest_permanent)

Basic InfoMount Target InfoSnapshot Chain**Quota information**Lifecycle Policies

Type	ID	Capacity quota	Quota of files	Usage
UID	100	10	10000	0

Total items: 1

Modifying quotas

Set quotas for the same Uid/Gid of the existing quotas to modify the quotas.

Add user/user group quota

File system instancecfs-45a6fb370

Quota type *

Gid

ID *

Capacity quota *

—

10

+

GiB

Quota of files *

—

10000

+

Confirm

Cancel

Deleting quotas

Click the ID/name of a file system to go to its details page, select **Quota information**, and click **Delete** to delete the quota.

cfs-45a6fb370

(brucetest_permanent)

Basic Info

Mount Target Info

Snapshot Chain

Quota information

Lifecycle Policies

Type	ID	Capacity quota	Quota of files	Usage
UID	100	10	10000	0
Total items: 1				

Use Limits

Last updated : 2024-01-22 22:15:48

Note:

This feature is available only to Turbo file systems. It is under beta testing. To try it, [submit a ticket](#).

The use limits of the user quota feature are as detailed below:

Item	Limit
Capacity quota	10 GiB to 1,000 TiB, with a minimum step of 1 GiB
File count quota	10,000 to 1 billion, with a minimum step of 10,000

Note

If quotas are set for both a user and their user group, both the quotas will take effect. If the limit of a quota is reached, it will trigger the system quota limit, and the "No space" error will occur for writes exceeding the limit.

If both capacity quota and file count quota are small, the capacity actually available for writes and number of files allowed will be slightly greater than the quota limits.

Features

Last updated : 2024-01-22 22:15:47

Overview

CFS user quota is a resource management feature for data management involving multiple users. It enables you to set capacity quotas and file count quotas for users and user groups as needed, addressing capacity allocation among multiple users and user groups and improving the overall resource utilization of the system.

Use cases

Allocating storage space among multiple departments/tenants

Different capacity usage limits need to be set for different end users to meet resource allocation or budget requirements. You can assign users and user groups to actual users, and allocate them resources using the user/user group quotas, enabling more efficient use of storage resources.

Strengths

User-specific capacity quota and file count quota

The user-specific capacity quota enables you to precisely allocate the storage capacity available to a user.

The user-specific file count quota enables you to restrict the usage of inode in the file system by a user, preventing them from saving too many small files that adversely affect the entire system.

User group-specific capacity quota and file count quota

The user group-specific capacity quota enables you to precisely allocate the storage capacity available to a user group. It is suitable for capacity and budget management in a multi-layer organization.

The user group-specific file count quota enables you to restrict the usage of inode in the file system by a user group, preventing them from saving too many small files that adversely affect the entire system.

Data Encryption Operation Guide

Last updated : 2024-01-22 22:15:48

Setting quotas

Step 1. Log in to the console

Log in to the [CFS console](#), and click **Create** to create a file system.

File System Guangzhou(19)											
<ul style="list-style-type: none">CFS's High-Performance Turbo is now available. See SpecificationsCFS offers storage resource units available for all sub-products. Package details											
<div>Create CFS client assistant</div>											
ID/Name	Monitor	Status ▾	Used/Total ⚡	IA storage	Maximum Throughput ①	Availability : ▾	VPC ID/CCN ID	IP	Storage Class ▾	Tag	Operation
cfs-ow7fp4ut frank	<div></div>	Available	64MiB/160TiB	-	100MiB/s	Guangzhou Zo...	vpc-74ggbalh	10.0.6.17	Standard	2	Edit Tag Create Snapshot Delet
cfs-4812ee604 frank	<div></div>	Available	0MiB/20TiB	0MiB	2GiB/s	Guangzhou Zo...	vpc-8twhugrt	11.0.0.21	Standard Turbo		Edit Tag Create Snapshot Expa
cfs-4904f2625 wyzx-thomasszha-test	<div></div>	Creating	0MiB/0GiB	0MiB	0MiB/s	Guangzhou Zo...	-	-	High-Performance Tur...		Edit Tag Create Snapshot Expa
cfs-r8qmkol9 未命名	<div></div>	Available	0MiB/160TiB	-	100MiB/s	Guangzhou Zo...	vpc-mkuolnm7	10.206.0.23	Standard	1	Edit Tag Create Snapshot Delet
cfs-mbxuoyj1	<div></div>	Available	0MiB/160TiB	-	100MiB/s	Guangzhou Zo...	vpc-mkuolnm7	10.206.0.118	Standard		Edit Tag Create Snapshot Delet

Step 2. Select the file system type (Turbo)

Select the Turbo file system, and click **Next: Set up Details**.

Create File System

1 Select File System Type

2 Set Up Details

Storage Class

Standard

Throughput: Max 300 MiB/s
IOPS: Max 15,000
Latency: Milliseconds
Capacity: 0-160TiB
✓ Cost efficiency
✓ Suitable for the small-scale general data storage such as log storage and data backup

High-Performance

Throughput: Max 1 GiB/s
IOPS: Max 30,000
Latency: Submillisecond
Capacity: 0-32TiB
✓ High performance and low latency
✓ Suitable for small-scale latency-sensitive core businesses such as DevOps, website application source code, and cloud desktop

Throughput

Throughput: Max 200 GiB/s
IOPS: Max 2,000,000
Latency: Milliseconds
Capacity: 0-1PiB
✓ High throughput and concurrency
✓ Suitable for read-intensive image rendering, game image distribution, etc.

Standard Turbo

Throughput: Max 100 GiB/s
IOPS: Max 2,000,000
Latency: Milliseconds
Capacity: 20TiB-100PiB
✓ High throughput and large capacity
✓ Suitable for businesses that require high throughput and mixed loads such as video rendering, AI rea

High-Performance Turbo

Throughput: Max 100 GiB/s
IOPS: Max 10,000,000
Latency: Submillisecond
Capacity: 10TiB-100PiB
✓ High throughput and IOPS
✓ Suitable for businesses that use large-scale small files such as AI training, auto driving, HPC computa

Next: Set Up Details

Step 3: Enable encryption

1. Enable encryption

Select File System Type

2 Set Up Details

Storage Class

Standard Turbo

File System Name

Please enter no more than 64 Chinese characters, alphabets, numbers under

Region

Guangzhou

Availability Zone

Guangzhou Zone 6

Data Source

☐ Use a snapshot

Storage Usage

20 TiB250 TiB500 TiB1000 TiB

Minimum capacity of Standard Turbo: 10 TiB

Network Type

☒ CCN☐ VPC

CCN Instance

cfs-turbo-test (ccn-22zyaepv)

VPCID	Subnet
vpc-pu8gqthp	172.16.0.0/16
vpc-grftzqyt	192.168.0.0/16
vpc-7eizj7p	10.0.0.0/16
vpc-acovo7jd	10.0.0.0/16
vpc-khmqcyp1	11.0.0.0/24

Total items: 155 / page1 / 3 pages

Turbo IPv4 CIDR

11.0.0.0/24

Permission Group

default (pgroupbasic)

Auto scaling policy

☐ Configure an auto scaling policy for the file system **Recommended**

Encryption

☐ Enable

Tags

+ Add

Fees

8.53CNY/hour

(For reference only. The actual cost depends on your usage)

Back

Create now

2. View encryption

You can view the encryption of a file system on its details page.

←

cfs-45a6fb370

(brucetest_permanent)

Basic Info

Mount Target Info

Snapshot Chain

Quota information

Lifecycle Policies

i

Currently, Turbo file systems can be mounted only to Linux clients.

Basic Info

Region

Guangzhou

Availability Zone

Guangzhou Zone 7

Instance ID

cfs-45a6fb370

Instance Name

brucetest_permanent

File System Protocol

TURBO

Storage Class

High-Performance Turbo

Instance Status

Available

Encryption status

Not encrypted

Creation Time

2023-05-17 19:57:44

Maximum Throughput*i*

1GiB/s

Used/Total

2GiB / 5TiB

IA storage

0MiB

Tag

Use Limits

Last updated : 2024-01-22 22:15:48

Cloud File Storage (CFS) encryption is subject to the following limitations:

Item	Limit
CFS	Available to the Turbo file systems only. Available for setting only during creation of a file system. Unable to change an encrypted Turbo file system to an unencrypted one.
Other	After the encryption option is selected, the system will use CFS encryption keys that cannot be specified, deleted, or changed.

Overview

Last updated : 2024-01-22 22:15:47

Overview

You can enable encryption for your Cloud File Storage (CFS) file systems if you need to encrypt the data stored in these systems to meet security or compliance requirements.

Tencent Cloud encrypts CFS data using keys based on the standard AES-256 algorithm.

CFS keys will be used for encryption by default, with no custom keys supported. Keys are stored in key management services that are well protected physically and logically, effectively preventing unauthorized access. Data keys in CFS are used in the memory of the host only, and will not be stored in any permanent storage medium in plaintext.

Notes

CFS encryption is free of charge.

CFS encryption is implemented on the server side, ensuring proper access by clients and no change to the original access method.

CFS encryption is available to Turbo file systems only.

The encryption feature involves data encryption and decryption and thus will compromise the file system performance by 10%-15%. We recommend you use it based on your actual needs.

Data Lifecycle Management

Use Limits

Last updated : 2024-01-22 22:15:48

Limits

The limits are as described below:

Item	Limit
Size of files subject to lifecycle management	Only files of 64 KB or larger can be transitioned, and the transition policies do not apply to files smaller than 64 KB.
File systems associated with a lifecycle management policy	A lifecycle management policy can be associated with up to 20 file systems.
Lifecycle management policies associated with a file system	A file system can be associated with up to 20 lifecycle management policies.
File system paths associated in a policy	Up to 50 paths in the same file system can be associated in a policy.
Policies associated with a directory	A directory and its parent directory and subdirectories can be associated with up to one lifecycle management policy.

Note:

This feature is available to Turbo file systems only. To try it, [submit a ticket](#).

Precautions

If a directory associated with a lifecycle management policy is renamed, the files in it will no longer be subject to the lifecycle management policy.

Files in the directory associated with a deleted lifecycle management policy will no longer be transitioned to IA storage, and those in the same directory that have been in IA storage will remain unchanged in their storage status. The lifecycle management feature is unavailable to clusters created before March 15, 2023. If you need to use this feature, [submit a ticket](#) to upgrade the clusters first.

Operation Guide

Last updated : 2024-01-22 22:15:47

Setting a lifecycle policy

Step 1. Create a lifecycle policy

1. Go to the [Lifecycle Policies](#) page and click **Create**.

Scheduled Snapshot Policies Guangzhou ▼

Note

- You can create a scheduled snapshot policy for a file system to create and delete snapshots periodically for your business. For details, see [Scheduled Snapshots](#).
- You are advised to create scheduled snapshots to protect your data for core businesses.
- Note: If your account has overdue payments, all snapshot operations will be forbidden. If the overdue payments last over 30 days, all snapshots in your account will be deleted.

Create Policy

Separate multiple key

ID/Name	Status ▼	Assigned File Syst ↕	Policy Details	Scheduled Snapshot	Creation Time ↕	Snapshot Ret
asp-o8qe5vyv aileen-test	Normal	0	Auto-create a snapshot at 02:00, 11:00, 12:00, 16:00, 17:00 every Tues, Wed and delete it after 30 days	<input type="checkbox"/>	2021-11-16 20:07:36	30 days
asp-6ud9sdej aileen-test-cross	Normal	0	Auto-create a snapshot at 01:00 every Mon and delete it after 0 days	<input checked="" type="checkbox"/>	2023-06-19 15:47:51	Permanen
asp-6ua2iscb ap-aileen-heihei	Normal	0	Auto-create a snapshot at 01:00 every Mon and delete it after 1 day	<input checked="" type="checkbox"/>	2023-06-19 16:29:58	1 day
asp-i39lyek3 brucetest	Normal	1	Auto-create a snapshot at 06:00 every Tues and delete it after 30 days	<input checked="" type="checkbox"/>	2023-06-20 00:52:25	30 days
asp-0cq3kc6r test	Normal	0	Auto-create a snapshot at 12:00, 21:00, 22:00 every day 1, 2, 4 of every month and delete it after 30 days	<input checked="" type="checkbox"/>	2023-03-22 15:52:57	30 days

Note:

This feature is available only to a Turbo file system. If you need it, [submit a ticket](#).

2. Specify the lifecycle policy by transition period or file size according to your business needs, and click **Next**.

Note:

You can set a file less than or equal to a value in the range of 1–64 MB as a small file.

Step 2. Apply the lifecycle policy

Apply the lifecycle policy to a path on the Turbo file system and click **Submit**.

Create lifecycle policy

1 Configure rule > 2 Apply

Up to 20 transition policies can be applied to a file system. Only one policy can be configured per directory. You can no longer apply a policy to a directory if a policy has already been applied to its parent directory or subdirectory. A policy can be applied to up to 20 file paths per file system).

[Add](#)[Remove](#)**File System****Directory path**

cfs-4812ee604(frank) ▼

Enter a directory path

Total items: 1

5 ▼ / page



1

/ 1 page

[Back](#)[Submit](#)

Note:

You can no longer apply a policy to a directory if a policy has already been applied to its parent directory or subdirectory.

If you want to apply the policy to files under /test/, you can set the target path to test, /test, or /test/. The system will automatically convert them to /test/.

Viewing lifecycle policies

Step 1. View a lifecycle policy

1. Go to the [Lifecycle Policies](#) page.

Lifecycle Policies Guangzhou

Create

ID/Name	Transition rule	Target directories	Storage Class	Creation Time	Operation
policy-qh4spptp frank	> 10 GiB: None ≤ 64 MiB: None Other file: 30 days	1	IA storage	2023-07-14 16:11:21	Configure Copy Delete

Total items: 1

2. Click a **policy ID** to view the policy.

← **policy-qh4spptp** (frank)

Transition rule

Policy Name

frank

Region

Transition rule

1. If the proportion of small files is not high, set "Transition period" to "None" for files ≤ 64 MiB. This occupies a small amount of storage space but greatly improves the IOPS performance of small files.

2. File extraction may take too long if large files are transitioned, so we recommend you not enable transitioning for large files over 10 GiB for read efficiency.

File category	File size	Transition period
Extra large file	> 10 GiB	None
Small file	≤ 64MiB	None
Other file	10 GiB ≥ Other file > 64 MiB	30 days

Transition log

File System/LifecycleLog

Target

Up to 20 transition policies can be applied to a file system. Only one policy can be configured per directory. You can no longer apply a policy to a directory if a policy has already been applied to its parent directory or subdirectory. A policy can be applied to up to 20 file systems (max 50 paths per file system).

File System

Directory path

Loading...

Total items: 1

Step 2. View the active lifecycle policies of a file system

1. Go to the [file system list page](#), and click the ID of the target file system.

File System Guangzhou(18)

• CFS's High-Performance Turbo is now available. [See Specifications](#)
 • CFS offers storage resource units available for all sub-products. [Package details](#)

[Create](#) [CFS client assistant](#) Separate multiple keys

ID/Name	Monitor	Status	Used/Total	IA storage	Maximum Throughput	Availability	VPC ID/CCN ID	IP	Storage Class	Tag
cfs-4812ee604 frank		Available	0MiB/20TiB	0MiB	2GiB/s	Guangzhou Zo...	vpc-8twhugt	11.0.0.21	Standard Turbo	
cfs-4904f2625 wyxz-thomasszha-zhao-test		Creating	0MiB/0GiB	0MiB	0MiB/s	Guangzhou Zo...	-	-	High-Performance Tur...	
cfs-n8qmlol9 未命名		Available	0MiB/160TiB	-	100MiB/s	Guangzhou Zo...	vpc-mkuolnm7	10.206.0.23	Standard	1
cfs-mtxuoqj1 cls-nlzhz4m_pvc-5efcd...		Available	0MiB/160TiB	-	100MiB/s	Guangzhou Zo...	vpc-mkuolnm7	10.206.0.118	Standard	
cfs-47cbb3591 kevinjzhang		Available	0MiB/20TiB	0MiB	2GiB/s	Guangzhou Zo...	vpc-khmqcyp1	11.0.0.2	Standard Turbo	

2. On the lifecycle policies tab page, view the active policies.

[←](#) **cfs-4812ee604** (frank)

[Basic Info](#)
[Mount Target Info](#)
[Snapshot Chain](#)
[Quota information](#)
[Lifecycle Policies](#)

Separate multiple keys

Directory path	Transition rule	Status	Op
/	> 10 GiB: None ≤ 64 MiB: None Other file: 30 days	Active	Co

Total items: 1

Modifying or deleting a lifecycle policy

1. On the details page of a lifecycle policy, click **Edit** in the upper-right corner.

←

policy-qh4spptp (frank)

Transition rule

Policy Namefrank

Region

Transition rule

1. If the proportion of small files is not high, set "Transition period" to "None" for files ≤ 64 MiB. This occupies a small amount of storage space but greatly improves the IOPS performance of small files.
2. File extraction may take too long if large files are transitioned, so we recommend you not enable transitioning for large files over 10 GiB for read efficiency.

File category	File size	Transition period
Extra large file	> 10 GiB	None
Small file	≤ 64MiB	None
Other file	10 GiB ≥ Other file > 64 MiB	30 days

Transition logFile System/LifecycleLog

Target

Up to 20 transition policies can be applied to a file system. Only one policy can be configured per directory. You can no longer apply a policy to a directory if a policy has already been applied to its parent directory or subdirectory. A policy can be applied to up to 20 file systems (max 50 paths per file system).

Search by

File System	Directory path
<div>cfs-4812ee604</div> <div>frank</div>	/

Total items: 1

2. Modify or delete the policy and click **Confirm**.

←

policy-qh4spptp (frank)

Transition rule

Policy Name

frank

Region

Guangzhou

Transition rule

1. If the proportion of small files is not high, set "Transition period" to "None" for files ≤ 64 MiB. This occupies a small amount of storage space but greatly improves the IOPS performance of small files.
2. File extraction may take too long if large files are transitioned, so we recommend you not enable transitioning for large files over 10 GiB for read efficiency.

File category	File size	Transition period
Extra large file	> 10 GiB	None ▾
Small file	≤ <div>64</div> MiB	None ▾
Other file	10 GiB ≥ Other file > 64 MiB	30 days ▾

Transition logFile System/LifecycleLog

Target

Up to 20 transition policies can be applied to a file system. Only one policy can be configured per directory. You can no longer apply a policy to a directory if a policy has already been applied to its parent directory or subdirectory. A policy can be applied to up to 20 file systems:

AddRemove

<input type="checkbox"/>	File System	Directory path
<input type="checkbox"/>	<div>cfs-4812ee604(frakn)</div>	/

Total items: 1

ConfirmCancel

Viewing the storage of transitioned data

Go to the [file system list page](#) to view the storage of transitioned data.

File System

Guangzhou(18)

- CFS's High-Performance Turbo is now available. [See Specifications](#)
- CFS offers storage resource units available for all sub-products. [Package details](#)


Create

CFS client assistant

ID/Name	Monitor	Status ▾	Used/Total ↕	IA storage	Maximum Throughput ⓘ	Availability : ▾	VPC ID/CCN ID	IP	Storage Class
cfs-4812ee604 frank		Available	0MiB/20TiB	0MiB	2GiB/s	Guangzhou Zo...	vpc-8twhugxt	11.0.0.21	Standard Turt
cfs-4904f2625 wyzx-thomaszhao-test		Creating	0MiB/0GiB	0MiB	0MiB/s	Guangzhou Zo...	-	-	High-Perform
cfs-n8qmlol9 未命名		Available	0MiB/160TiB	-	100MiB/s	Guangzhou Zo...	vpc-mkuolnm7	10.206.0.23	Standard
cfs-mtxuovj1 cls-nlz6hz4m_pvc-5efcd...		Available	0MiB/160TiB	-	100MiB/s	Guangzhou Zo...	vpc-mkuolnm7	10.206.0.118	Standard
cfs-47cbb3591 kevinjzhang		Available	0MiB/20TiB	0MiB	2GiB/s	Guangzhou Zo...	vpc-khmqcyp1	11.0.0.2	Standard Turt
cfs-n25gujvq kevinjzhang-test		Available	64MiB/160TiB	-	100MiB/s	Guangzhou Zo...	vpc-74ggbalh	10.0.6.7	Standard
cfs-jbya7s1r bruce15NFS		Available	1.25GiB/160TiB	-	100MiB/s	Guangzhou Zo...	vpc-mkuolnm7	10.206.0.197	Standard
cfs-57ba74639 bruce15NFS_31		Available	0MiB/160TiB	-	100MiB/s	Guangzhou Zo...	vpc-mkuolnm7	10.206.0.212	Standard
cfs-56a4ae6b5		Available	0MiB/160TiB	-	100MiB/s	Guangzhou Zo...	vpc-pu8gqxhp	172.16.32.13	Standard

Viewing the transitioning status of a file

To view the transitioning status of a file, use the following command:



```
lfs hsm_state /path/to/file
```

Note:

If "archived" is returned, the file has been transitioned to IA storage and has not been released from the file system.


If "released" is returned, the file has been transitioned to IA storage and has been released from the file system.

By default, a file is automatically released 1 hour after being archived. This frees up space on the file system.

Loading data


You can load files using the following commands:

For a single file:



```
sudo lfs hsm_restore /path/to/file
```

For all files in a directory:




```
nohup find /path/to/preload -type f -print0 | xargs -0 -n 1 sudo lfs hsm_restore &
```

Releasing data


After you have loaded data, you can release it using the following commands:

For a single file:



```
sudo lfs hsm_release /path/to/file
```

For all files in a directory:



```
nohup find /path/to/release -type f -print0 | xargs -0 -n 1 sudo lfs hsm_release &
```

Note:

The `hsm_release` operation cannot be performed on a file that has not been transitioned.

Features

Last updated : 2024-01-22 22:15:48

Overview

Data lifecycle management is an advanced data management feature provided by Tencent Cloud File System (CFS) to balance high performance and low cost in large-scale file storage. With this feature, you can set custom data lifecycle management policies, and the file system can automatically move cold data to lower-cost IA storage based on these policies. When the data is accessed, it is automatically restored to file storage, with the entire process being transparent to the business. This lowers storage costs.

Use cases

Support for storage of massive amounts of data

In scenarios involving massive data storage, traditional cloud file storage is often unable to meet users' cost control requirements due to its high unit price. In the past, the common solution was to manually store data on object storage using tools. However, with data lifecycle management, data transitioning can be completed transparently through simple configuration, greatly simplifying the operation and reducing the total cost of data storage.

Hot/Cold data

In scenarios such as autonomous driving, AI training, and offline analysis, the access frequency of data is different. For example, newly written data is often accessed very frequently, while the access frequency of older data gradually decreases over time. The data lifecycle management feature is well-suited for such access patterns. It can significantly reduce the cost of storing cold data, while meeting high-performance read and write requirements.

Strengths

Support for flexible lifecycle policies

Policies based on directories: Different businesses often have different demands for lifecycle policies. Setting policies based on directories can better meet the demands for diverse policies.

Policies based on file size: Large files often take a long time to restore after transitioning. If your business is sensitive to latency in reading large files, you can set lifecycle policies based on file size to meet your requirements for timeliness.

Policies based on access period: You can flexibly adjust the scope of data for transitioning according to the characteristics of your business, reducing the impact of frequent transitioning and restoration on your business and making better use of lifecycle policies.

Transparent to the business side

After a lifecycle policy is configured, the system automatically transitions and restores data, without any change required to the data access mode of the business side.

Reduced costs

The data lifecycle management feature achieves tiered storage of hot/cold data, which can reduce the unit cost by up to 70% in some scenarios.

Feature details

Lifecycle management policies

When creating a lifecycle management policy, you can configure rules to transition files that have not been accessed within 14, 30, 60, or 90 days to IA storage. Lifecycle management will determine whether to transition a file based on its access time (atime).

The following operations will update the access time of a file:

Reading a file

Writing a file

The following operations will not update the access time of a file:

Querying file metadata (such as by performing the ls or stat operation)

Renaming a file

Modifying file metadata such as user, group, or mode

Data transitioning/restoration

Data transitioning is a process in which data is moved from a Turbo file system to IA storage. When the trigger conditions of a policy are met, the system automatically adjusts the concurrency based on the current system load, and copies data to IA storage. Then, it releases the data in the Turbo file storage after 1 hour, with metadata information retained.

Data restoration is a process in which data is moved from IA storage to the Turbo file system. When data in IA storage is accessed for the first time, the system restores the data from IA storage to the Turbo file system. This process will take some time, depending on the file size and system load. When accessed later, this data will be obtained from the Turbo file system.

Upgrading Standard File Systems

Last updated : 2024-01-22 22:15:48

Overview

This document describes how to upgrade a file system.

Directions

1. Log in to the CFS console and go to the [File System](#) page.
2. In the pop-up window, click **Upgrade Now**.

Note:

If you select **Don't prompt again** and click **Not Now**, when you want to perform an upgrade later, you can view the upgradable file systems as instructed in [Related operations](#).

3. Select the file system to be upgraded and click **Upgrade File System**.
4. In the pop-up window, view the upgrade advantages and click **Next**.
5. Check the precautions on upgrade and click **Next**.
6. View the client connection status and click **Next**.
7. Select **I have confirmed that the file system has been unmounted from all clients** and click **Confirm Upgrade**.
8. In the pop-up window, click **Confirm Upgrade**.
9. Wait for the upgrade to complete, which will take 2–5 minutes.

Related operations

1. On the file system management page, click the **Filter** column.
2. Select **File System Upgrade** and **Upgradable**, and click **Confirm** to view the file systems upgradable in the region.

FAQs

1. How do I determine whether a CFS instance needs to be upgraded?

If an **Upgrade** button is displayed on the right of a file system, the instance needs to be upgraded.

2. What are the advantages of upgrade?

Higher performance

For legacy Standard instances, the maximum throughput of a file system is 100 MiB/s.

After they are upgraded to the new version, the minimum throughput of a file system will be increased to 100 MiB/s.

The bandwidth is calculated as follows: $\text{Min}[100 + \text{storage capacity (GiB)} * 0.1, 300]$.

Higher maximum storage capacity

The maximum storage capacity of file systems on the legacy version is 40 TiB, and that on the new version is increased to 160 TiB.

3. How long does the upgrade take?

The upgrade takes 2–5 minutes. Before the upgrade, you need to unmount the file system from all clients. After the upgrade, you need to remount the file system. Therefore, we recommend you spare at least half an hour for the upgrade.

4. Will the upgrade fail?

Generally, the upgrade will not fail. If the task fails on the backend under rare circumstances, an alarm will be triggered, and Tencent Cloud technical support will provide assistance promptly.

5. Can I purchase a new instance on the legacy version?

No. New Standard CFS instances are all on the latest version.

6. Will upgrading an instance affect the data and compatibility?

No.

7. When I access a file after the upgrade, I am prompted that "Unable to access xx. The handle is invalid.". What should I do?

This is because you did not perform unmounting from the client during the upgrade. Restart the client before mounting.