Cloud File Storage
Operation Guide

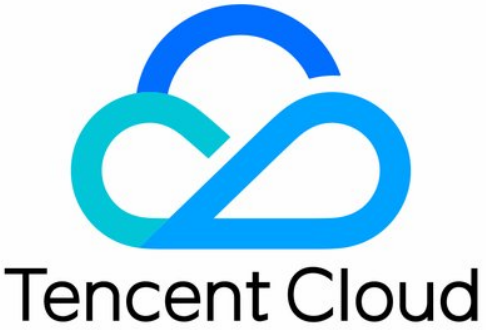

Copyright Notice

Trademark Notice

Service Notice

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

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

Contact Us

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

Contents

Operation Guide
Cloud Access Management

Last updated: 2023-08-29 10:12:28

Scenario

Cloud File Storage (CFS) now supports resource-level access management, allowing the primary account to grant specified operational permissions to designated users and user groups for specific resources. Once authorization is completed, both the CFS console and API will permit or prohibit user operations based on the granted permissions.
This guide will explain how to authorize read-only, read/write, and custom policies for CFS users. For more information on the principles and guidelines of Tencent Cloud Access Management, please refer to Access Management.

Instructions

Creating an Access Control Policy

Log in to the Access Management Console policy management page.

- To grant users permissions quickly, do a 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 provide a more flexible way to authorize users, and the Access Management Console offers various methods to generate these policies. This example will demonstrate

how to create a new custom policy using the "Create by Policy Generator" method (for other methods, please refer to the Policy documentation).
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.
On the Policy Management Console policy page, select **Create Custom Policy > Create by Policy Generator**. On the new policy page, use the policy generator to add multiple statements to a custom policy. The configuration is as follows:

Category	Parameter	Options and Effect
Effect	Effect	Allow or Reject
Sevice	Service	Select CFS here
Action	Action	All CFS-supported actions
Resources	Resource	Specify the resources in six-segment format: • The notation for all resources in Cloud File Storage is `*`. • The notation for all resources in a specified region is `qcs::cfs:ap-guangzhou::*`. • The syntax for all resources in all regions under a specified user is `qcs::cfs::uin/27700000:*`. • The syntax for specifying all file systems in a specific region under a specific user is `qcs::cfs:ap-guangzhou:uin/27700000:filesystem/*`. • The system syntax for a specific user group under a specified user is `qcs::cfs::uin/27700000:pgroup/pgroup-doxpcqh`. • 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. If a policy generator is used: • Service: You can only enter "cfs". • Region: Select a region from the drop-down list box as needed. • Account: The current account information will be auto filled. If it is not filled, you can enter `uin/xxxxxxx`, where `xxxxxxx` is the UIN. • Resource prefix: You can enter `filesystem`, `snap`, or `resource`, which represents file system instances, snapshot

		instances, or resource unit package instances, respectively. • Resource: Enter a resource ID, such as `cfs-xxxxx`.
Conditions	Condition	To determine under what conditions this policy will take effect, please refer to Effective Conditions for the setup method.

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	Supports and Limits
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 `*` when authorizing this API.
	CreateCfsFileSystem	Create 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	Write	You need to specify file system resources

		quota		when authorizing this API.
	DeleteCfsFileSystem	Deleting 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 point	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 permission group	Write	You do not need to specify resources when authorizing this API.
	UpdateCfsPGroup	Updating the information of a permission group	Write	You need to specify permission group resources when authorizing this API.
	DeleteCfsPGroup	Deletes 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 `*` 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 polices. See the `QcloudCFSReadOnlyAccess` policy statement to learn how to write the policy.

After completing the above parameter settings, click **Add Statement** to add a statement to this custom policy. You can repeat the above steps to add multiple statements. If there are duplicate or conflicting policies, please refer to Syntax Structure for their relationship and effective results.

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/User Group

If you choose the permissions provided by the system, you can directly search for QcloudCFSFullAccess, QcloudCFSReadOnlyAccess, or other custom policies in the policy list. Then, click **Associate Users/Groups** in the operation column on the right. In the pop-up window, find and select the users or user groups that need to be authorized. Finally, click **Confirm** to complete the authorization.

Deauthorizing a User/User Group

To revoke the permissions of an authorized user, go to the **Associated Users/Groups** list on the corresponding policy detail page, select the user/user group for which you want to cancel the authorization, and then click **Disassociate User/User Group**. After confirming the revocation, the user/user group will lose the permission to operate CFS resources.

Managing File Systems

Last updated: 2023-08-29 10:13:55

Scenario

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 point information.

> **Note**
> If a file system is in the **Creating** status, you cannot view its details or delete it.

Preparations

Log in to the Cloud File Storage console and navigate to the file system list page.

Instructions

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 for items 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 point and mounted clients.

- **Basic Information of the File System**

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

  The mount point information of an NFS file system includes the network information, permission group, and recommended mount command. 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 manage and view snapshots under a file system on the file system list page.

Renaming a File System

On the file system list page, click the file system whose name you want to change, enter the file system details page, and click on the right side of the instance name to make modifications.

Deleting a File System

When you no longer need a file system, you can locate it in the file system list and click **Delete** on the right to remove it.

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

Permission Management

Last updated: 2024-12-12 15:31:40

Scenario

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.

Instructions

Step 1. Create a permission group

1. Log in to the File Storage Console and click on **Permission Group** in the left sidebar.
2. On the Permission Group page, click **Create** to create a new permission group. In the pop-up window, configure the permission group name and remarks.

Step 2. Add a permission group rule

Click on **Permission Group Name** to access the rule list. Here, you can add, edit, or delete rules. If no rules are added to the permission group, all will be allowed. The rules are explained as follows.

Parameter	Description
Access Address	You can enter a single IP or a single network segment, such as 10.1.10.11 or 10.10.1.0/24. The default visitor address is *, which means all are allowed. Please note that you need to enter the internal IP of the CVM here.
Read & Write Permissions	Read-only or read/write.
User Permissions	The four options below are used for controlling the permissions of a user. • all_squash: Any user will be mapped to an anonymous user or user group. • no_all_squash: 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. • root_squash: A root user will be mapped to an anonymous user or user group.

	• no_root_squash: A root user will be allowed to maintain root account permissions. **⚠ Note:** • User permissions configuration is not supported for CIFS/SMB file systems and Turbo file system will not take effect. • The default permission is 755 for each file system, and nfsnobody does not have write permission. Therefore, if there are no special needs, `no_root_squash` is recommended. If the root user creates a file directory and mounts the file system, when the access address is set to `all_squash` or `root_squash`, the 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	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. **⚠ Note:** Priority configuration is not supported for CIFS/SMB file systems and will not take effect.

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 using a template.

Using Tags

Last updated: 2023-08-29 10:18:18

Scenario

Tags are key-value pairs provided by Tencent Cloud for identifying cloud resources. Tags can assist you in conveniently categorizing cloud resources from various perspectives, such as business, usage, owner, etc. It's important to note that Tencent Cloud does not utilize the tags you set; they are solely for your management of Tencent Cloud resources. This document guides you through the process of editing tags for file storage resources.

Usage Limits

There are several limits on tags:

- **Quantity Limit**: Each file system can have up to 50 tags.
- **Tag Key Limit**: Tag keys can only consist of `numbers`, `letters`, and `+=.@-`, with a maximum length of 255 characters.
- **Tag Value Limit**: A tag value can only be a `null string or a number`, `a letter`, or `+=.@-`, and the maximum length of a tag value is 127 characters.

Operation Examples

Case Description: A company has purchased six file systems. The information about the departments using these systems, their business scope, and the person in charge is 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 3 sets of tags to the file system:

Tag Key	Tag Value
dept	ecommerce
business	mkt
owner	zhangsan

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

Instructions

Tagging a new file system

Preparations

Ensure that you have the necessary tags before creating a file system; if you do not have any tags, please first create them in the Tag Console.

1. Log in to the File Storage console.
2. On the file system management page, click **Create**.
3. In the "**Create New File System**" configuration window that appears, locate the **Tags** configuration item below. Click **Add** to append tag information to this file system. (Only existing tags can be added here.)

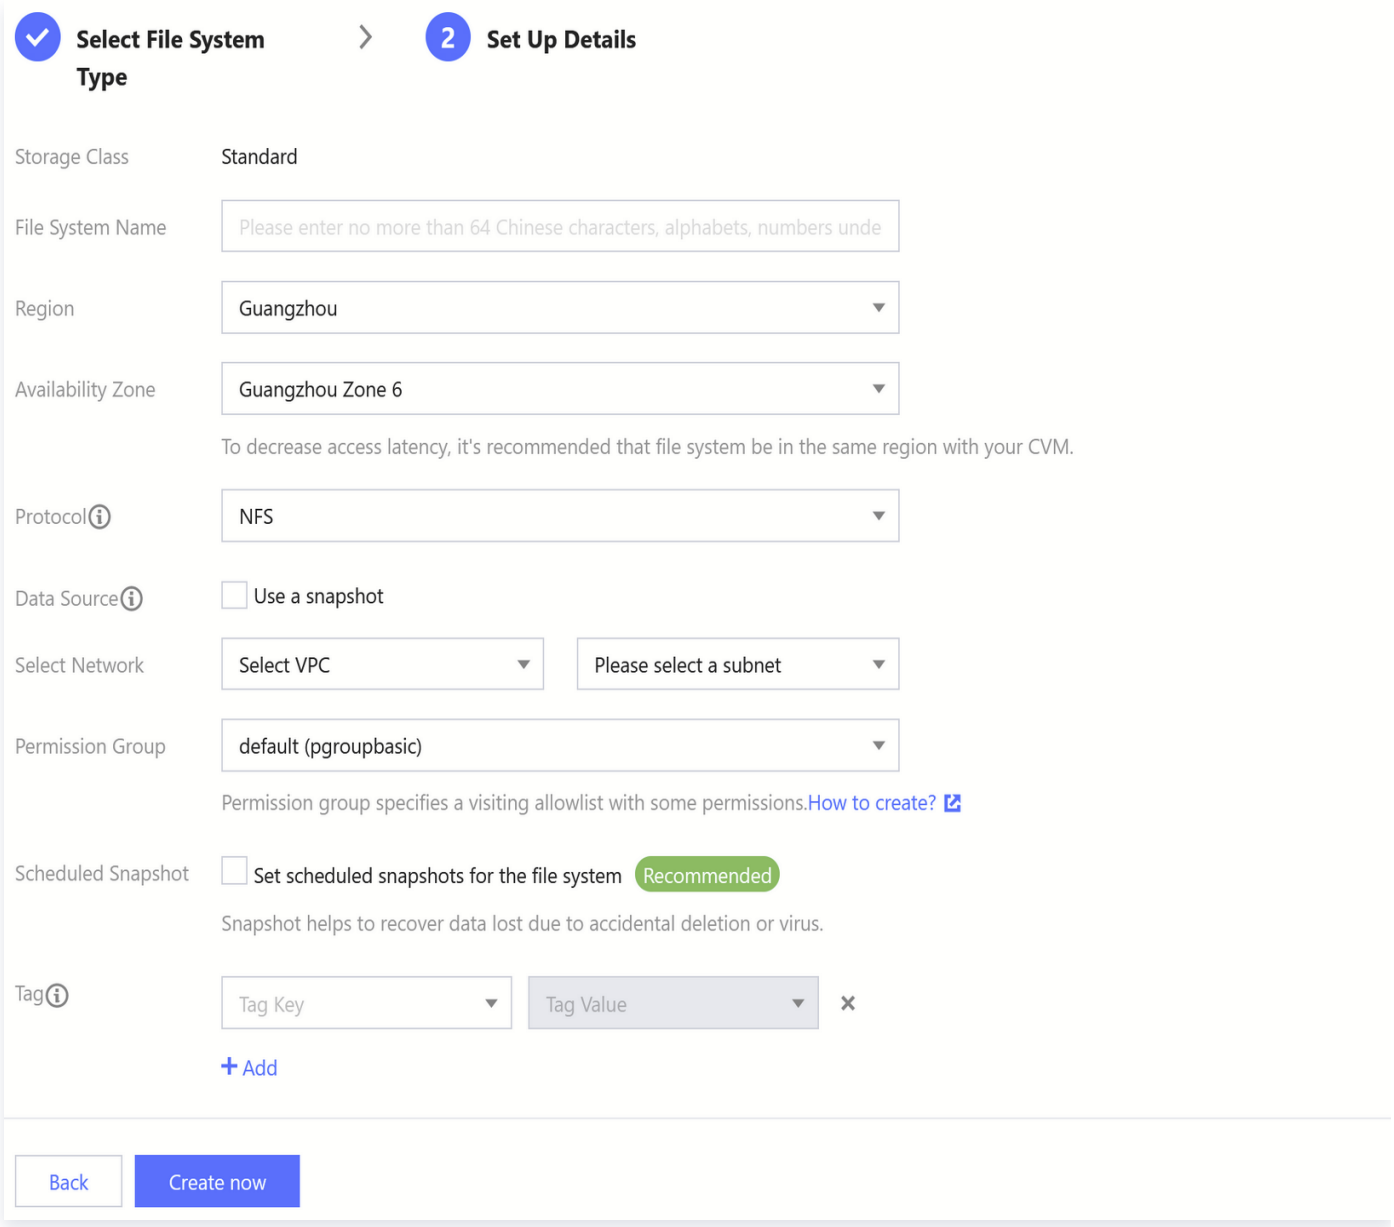

4. Click **Create Now**, the corresponding tags will be bound once the file system is successfully created.

> **Note**
> For a detailed explanation of the configuration items related to creating a file system, see the Creating a File System and Mount Target document.

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

1. On the file system management page, select the file system for which you want to edit tags, then click **Edit Tag** in the operation column.

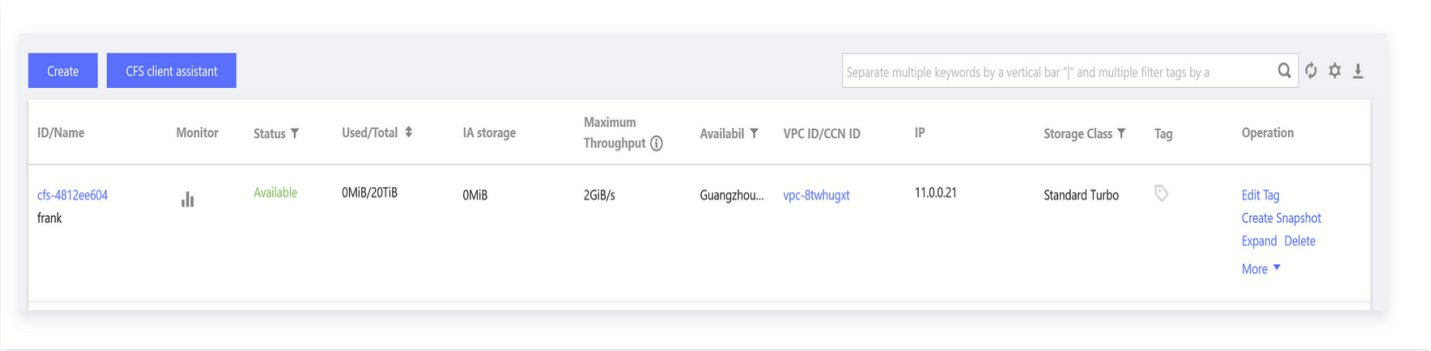

2. Add, modify, or delete a tag as needed in the pop-up window.

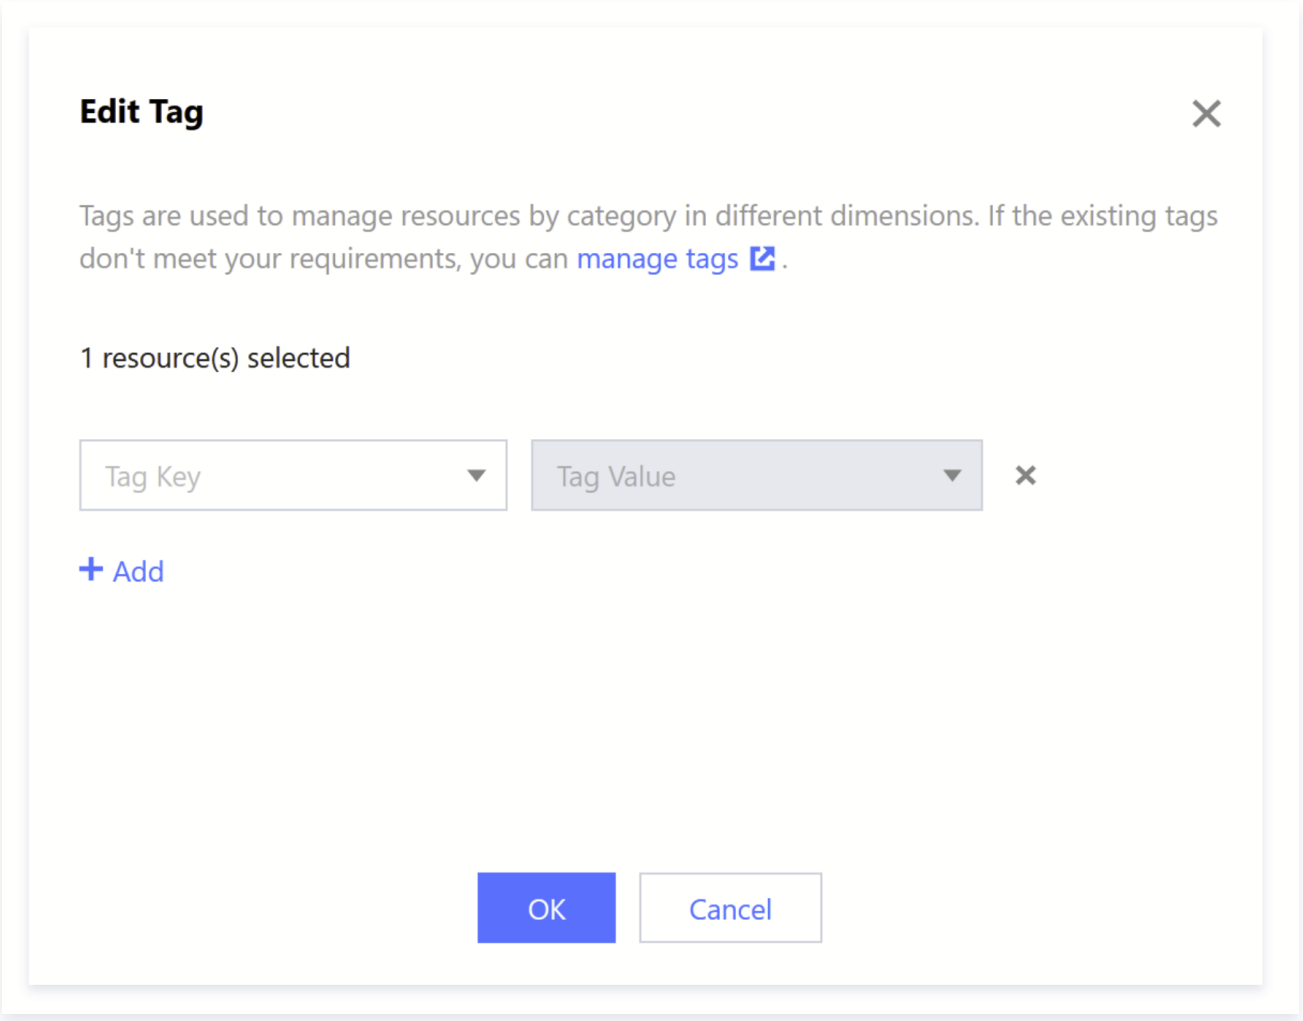

3. Click **OK** to finalize the addition, modification, or deletion of tags.

Filtering file systems by tag

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

1. In the search box, select **Tags**.

2. Enter the tag key and value after **Tag:** then click 🔍 or press ENTER to search, as shown below:
   For instance, if you wish to filter file systems with the tag key as 'tke-created' and the tag value as 'eee', you can enter `Tag:tke-created:yes`.

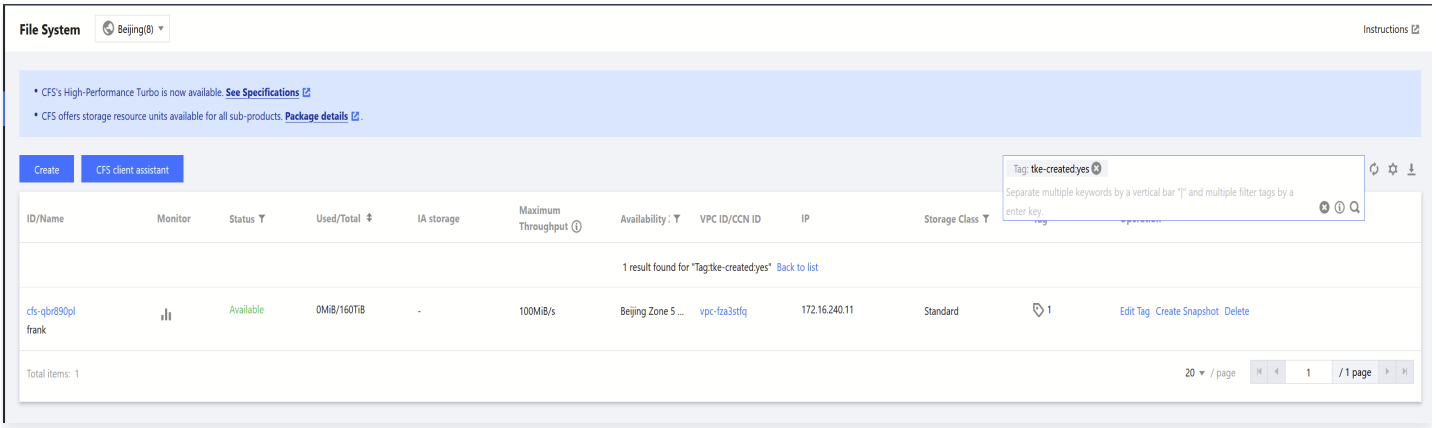

File System Expansion

Last updated: 2023-08-29 10:18:45

This document will guide you on how to perform relevant expansion operations on Turbo file systems via the console. Currently, it supports two methods: **manual expansion** and **automatic expansion policy**.

Note:

1. Standard Turbo/High-Performance Turbo are billed based on the purchased capacity. To prevent business disruptions due to insufficient capacity, it is recommended to carry out relevant expansion operations promptly when the capacity usage rate reaches 85%.
2. Standard Turbo/High-Performance Turbo support online expansion, and business access is not affected during the expansion process. However, the subsequent automatic rebalance service by the system may have a minor impact on performance (most businesses will not notice this).
3. Standard Turbo/High-Performance Turbo only support expansion and do not support reduction. If a smaller capacity is needed, it is recommended to create a new cluster and then perform data migration. For data migration operations, please refer to File System Data Copy Scheme.

Configure Automatic Expansion Policy (Recommended)

1. When creating a Turbo file system, select the option to configure an automatic expansion policy.

Note:
The default auto-expansion policy at creation is to start at an 85% threshold and expand to a 70% threshold. That is, when the capacity utilization rate reaches 85%, auto-expansion is initiated, reducing the capacity utilization rate to below 70%.

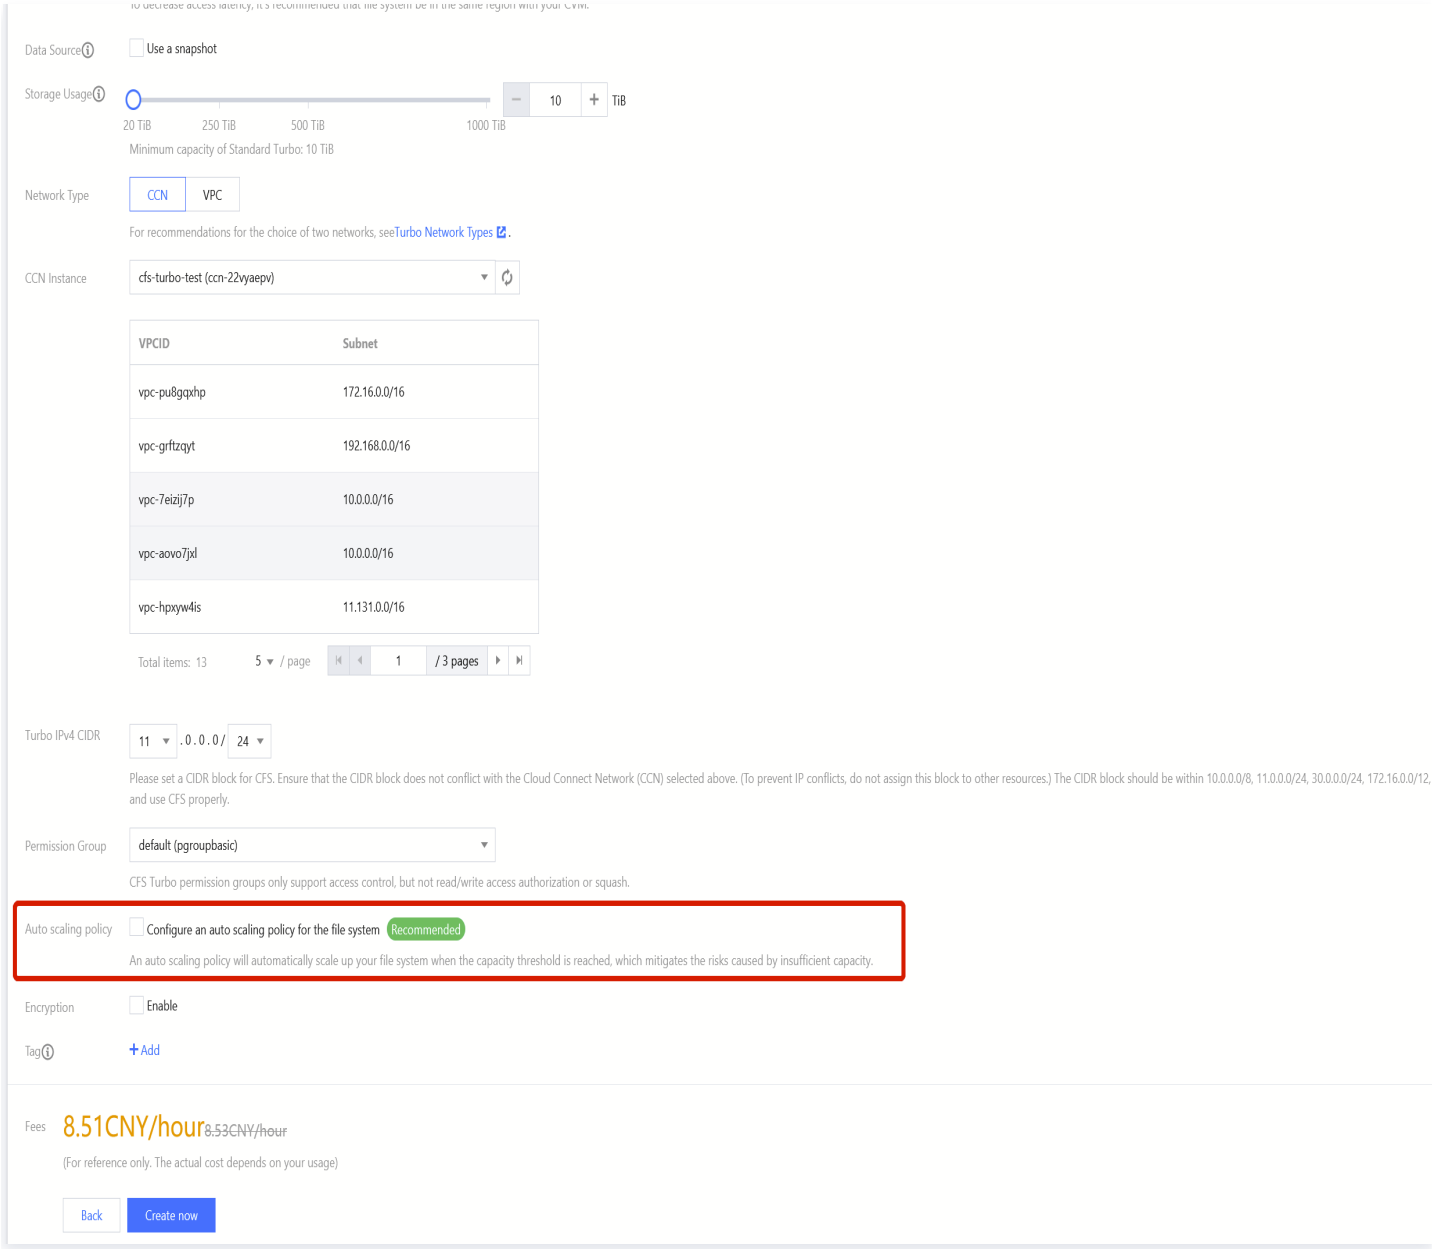

2. If you need to adjust or add an expansion policy later, you can click on "More" in the console and select "Configure auto scaling policy".

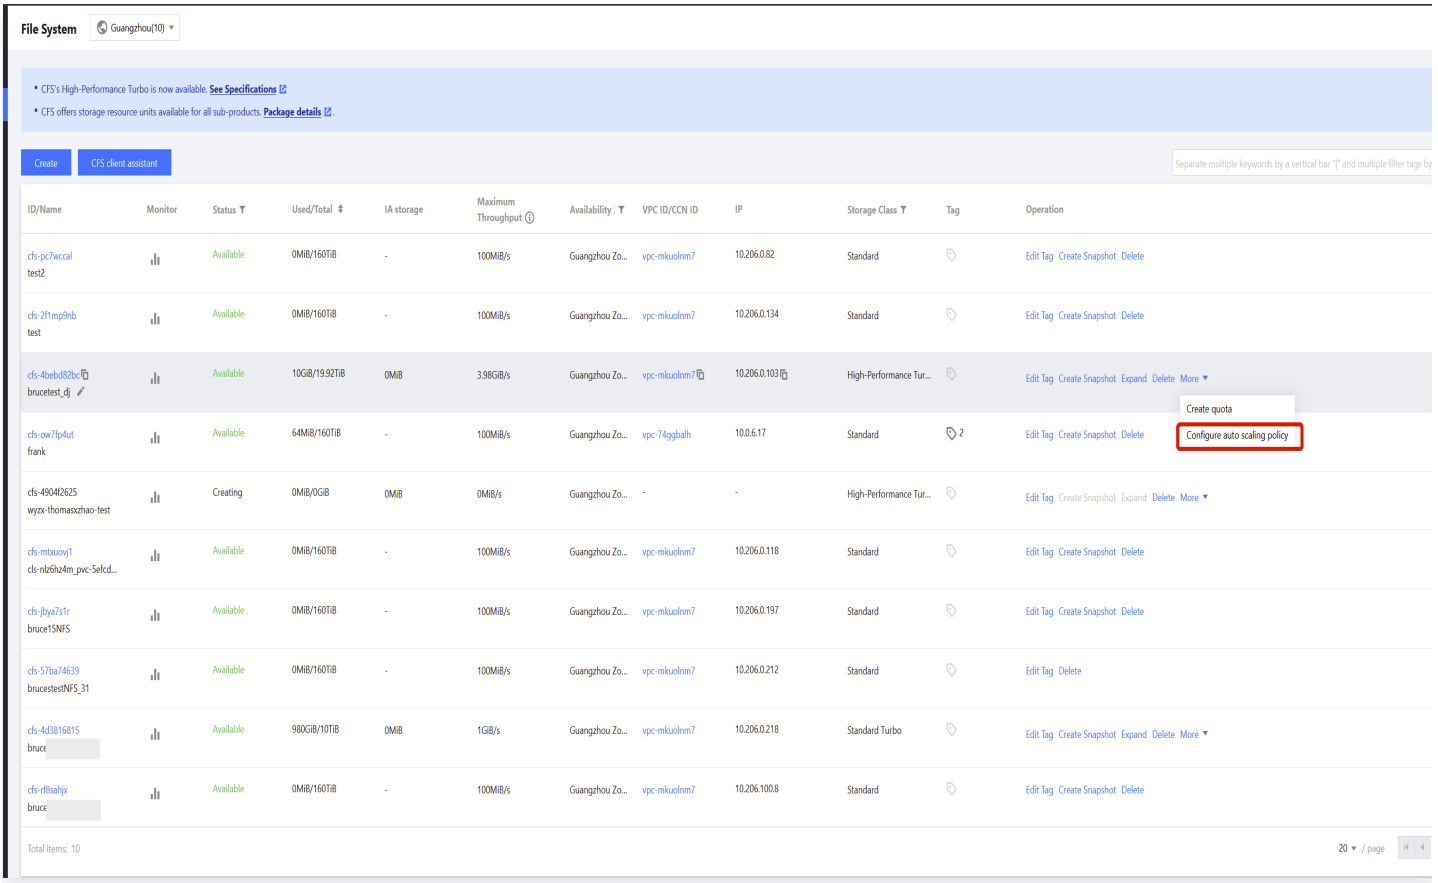

3. Depending on the requirements, you can enable/disable the expansion policy or adjust the expansion threshold.

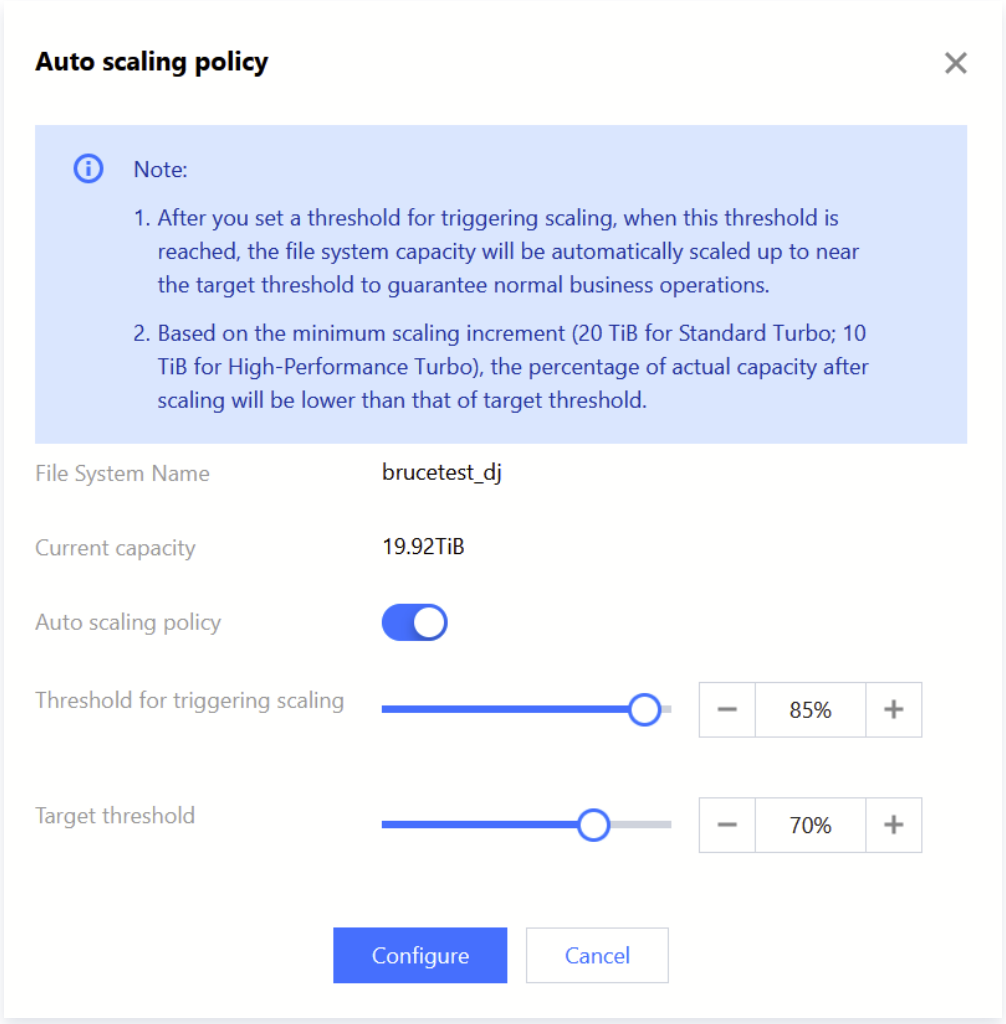

Manual

1. Log in to the CFS Cloud File Storage Console.
2. Perform expansion operations on the Turbo file system that requires expansion.

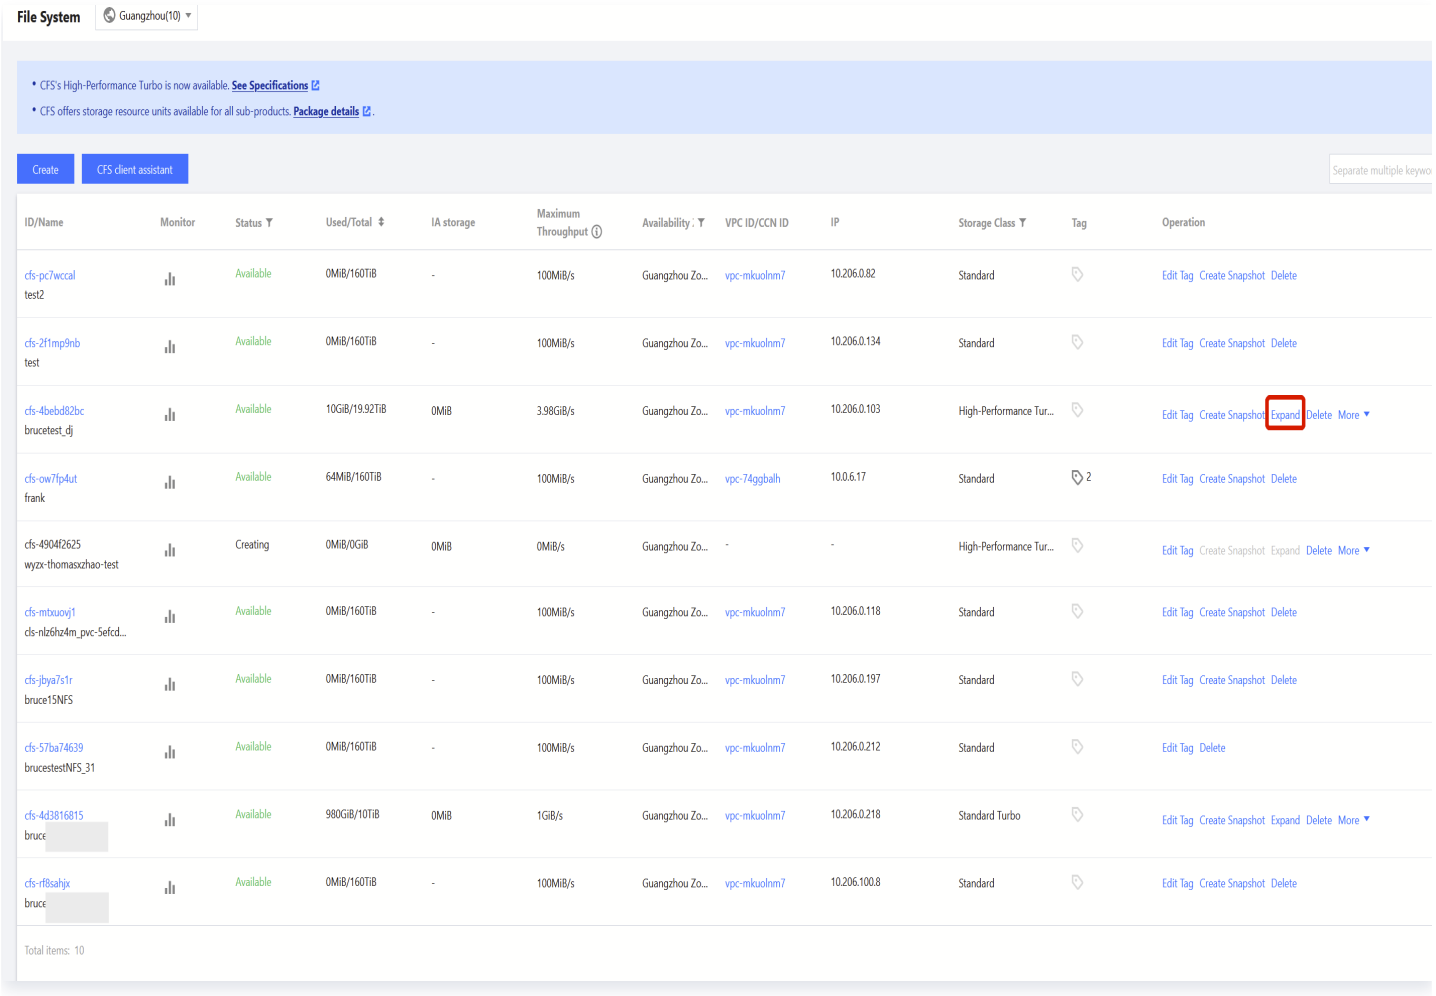

3. Select the target capacity and click **Expand**.

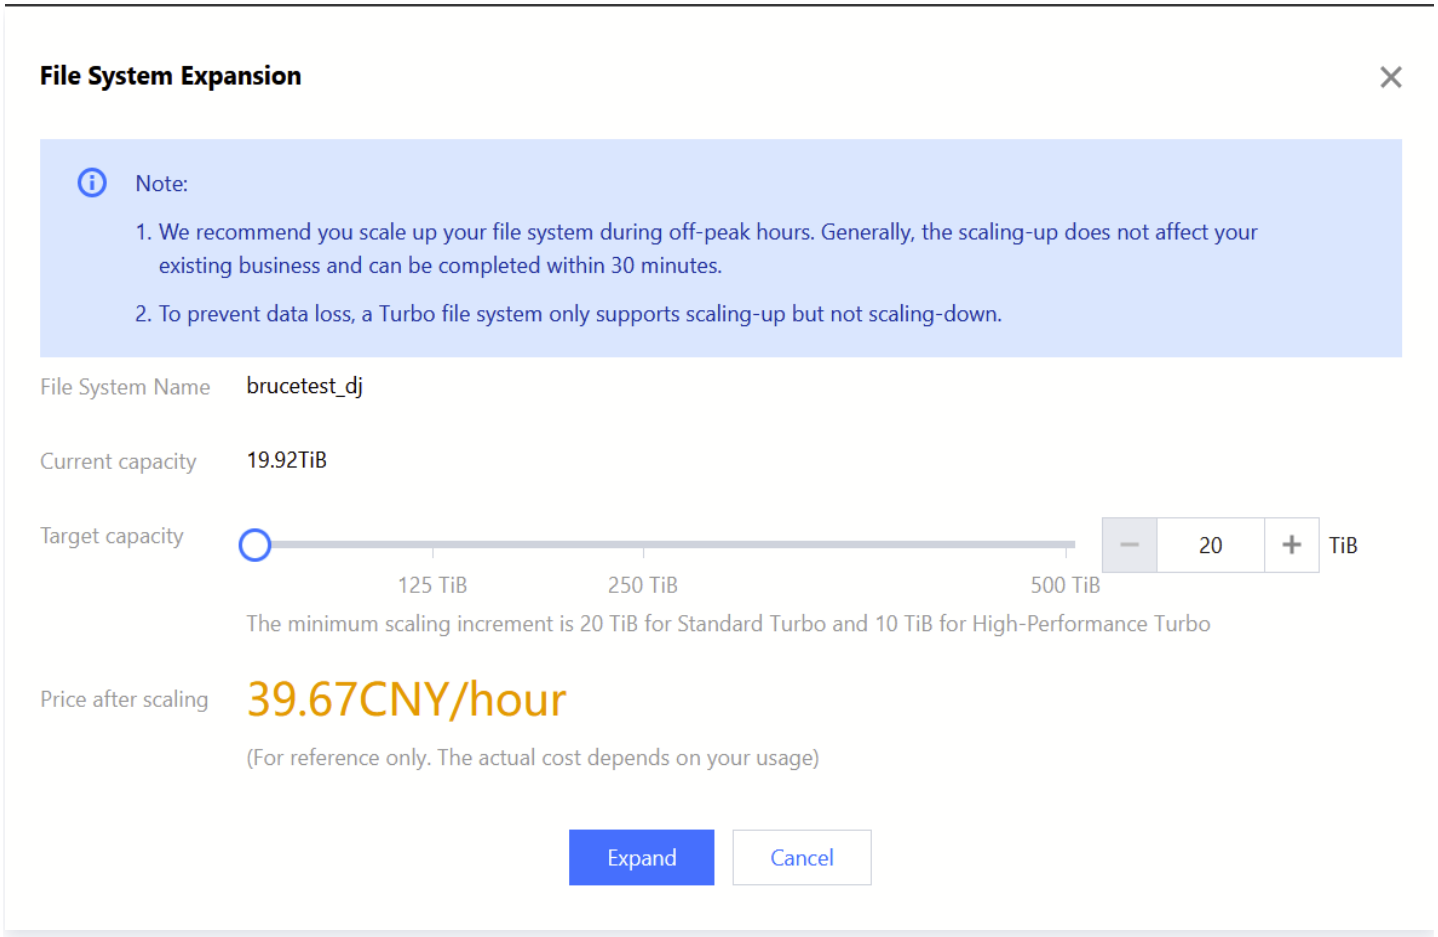

Snapshot Management
What Is a File System Snapshot?

Last updated：2023-08-29 10:21:30

Overview

- **Real-time replica of online data**
  A snapshot is a fully functional copy of a file system. In the event of an issue with the file system for which a snapshot has been created, the system can be swiftly restored to its pre-issue state using the snapshot. It is recommended to create a snapshot of the relevant file system prior to significant business changes, allowing for rapid data recovery in case of a failed business transformation.
- **Persistent backup for key milestones**
  Snapshots can serve as a long-lasting backup for business data, preserving milestone states of the data.
- **Quick business deployment**
  You can use snapshot files of your business to rapidly clone multiple file systems, thereby achieving the goal of swift 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 in, which will be backed up to snapshots.

Scenarios

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

- **Daily data backup**
  You can use snapshots to regularly back up important business data to avoid data loss caused by incorrect operations, attacks, and viruses.
- **Quick data recovery**
  You can create one or more snapshots before performing significant operations such as changing the operating system, upgrading application software, or migrating business data. If any issues arise during the change process, you can promptly restore business data using the created snapshots.

- **Application of multiple replicas of production data**
  You can create snapshots of production data to provide near-real-time, authentic production data for applications such as data mining, report querying, and development testing.
- **Quick environment deployment**
  You can create one or more file systems from an existing snapshot, allowing for the rapid, bulk deployment of identical business environments, thus saving time on repeated configurations.

Billing

Please see Snapshot Billing Method > Pay-as-You-Go (Postpaid)

Quota Limits

For detailed information on snapshot quota limits, please refer to Usage Limits.

Snapshot Types

- **Manual Snapshots**
  Manually create a snapshot of the file system data at a specific point in time. This snapshot can be used to rapidly create more file systems with the same data. For detailed operations, please refer to Creating Snapshots.
- **Scheduled snapshots**
  When your business is continuously updated, you can use scheduled snapshots to provide ongoing backup functionality. By simply formulating a backup strategy and associating it with the file system, you can achieve continuous backup of file system data within a certain cycle, significantly enhancing data security. For specific operations, please refer to Scheduled Snapshots.

> **Note**
> During the snapshot creation process, there may be instances where some application data is stored in memory and not persistently stored. This situation can result in the snapshot not capturing the most recent and complete file system data. Please refer to Considerations to ensure the consistency of snapshot data.

Case Review

Case 1: Failing to manually create snapshots before a high-risk operation, causing data loss

For instance, customer A has never created a snapshot of the file system. One day in May 2019, due to an fio test conducted by an operator, the file system was damaged and the data could not be retrieved.
Analysis: If customer A had created a snapshot of the file system before conducting the test, they could have swiftly activated the snapshot to create a new file system and promptly restore the business after the data was damaged.

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

For instance, Customer B had created snapshots for multiple file systems, but after January 2019, they did not create snapshots for newly purchased file systems due to cost considerations. One day in June 2019, data was lost and could not be recovered from a file system that was not protected by a snapshot due to accidental deletion of data at the file system level.
Analysis: If Customer B had implemented regular snapshot protection for this file system, the data could have been restored to the state at the time of the last snapshot after the accidental deletion, thus minimizing the loss. After the incident, Customer B proactively created a snapshot for this file system, significantly enhancing data protection.

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

For instance, customer C has snapshot protection for all file systems. One day in May 2019, an error caused a boot anomaly.
Analysis: Customer C promptly used a scheduled snapshot from two days prior to restore the data, preventing any business damage.
These cases all involve data loss due to incorrect operations, but the results are different. By comparison, we can find that:

- In the absence of a **created snapshot**, data retrieval can be extremely challenging when issues arise with the server or file system, potentially leading to significant losses.
- In the event of a server or file system issue, data can be largely recovered with minimal loss, provided that a **snapshot has been created**.

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

Limits

Last updated: 2023-08-29 10:21:46

This document describes limitations on CFS snapshots to help you use snapshotting more efficiently.

Item	Usage Limits
The number of snapshots	Each file system can have up to 100 snapshots
Number of scheduled snapshot policies	30
Number of file systems that can be bound to a scheduled snapshot policy	200
Supported file system type	StandardHigh-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 status of a file system is **Migrating**, the snapshot of the file system 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 nor 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 they are incremental or differential backups.
- During snapshot migration, the I/O performance of the file system may decrease by about 15%. You are advised to configure scheduled snapshot policies and take snapshots during off-peak hours of businesses.

Creating Snapshot

Last updated: 2023-08-29 10:22:16

Scenario

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 for 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.

Preparations

- You have activated CFS.

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

Supports and Limits

A snapshot only retains the data that has been written to the file system at that moment and does not include data that is in memory but has not been written to the file system (for example, files under the `/run` directory in the Linux system). It is strongly recommended that you ensure that the data in memory has been written to the file system and that read and write operations on the file system are suspended before creating a snapshot. The recommended operations for flushing data from memory are as follows:

- For better system performance, data is stored in the memory buffer before it is written to the file system at the 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.
- Execute the `sync` command to force immediate write of data in the memory buffer to the file system and avoid writing new data before creating a snapshot. If no error message is returned after the command is executed, it means that the data in the cache has been

written to the file system. As shown in the figure below:

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

Instructions

1. Log in to the CFS console.
2. Click **Create Snapshot** on the right side of the target file system row.
3. In the pop-up dialog box, enter the snapshot name and click **Confirm**, as shown in the figure below:

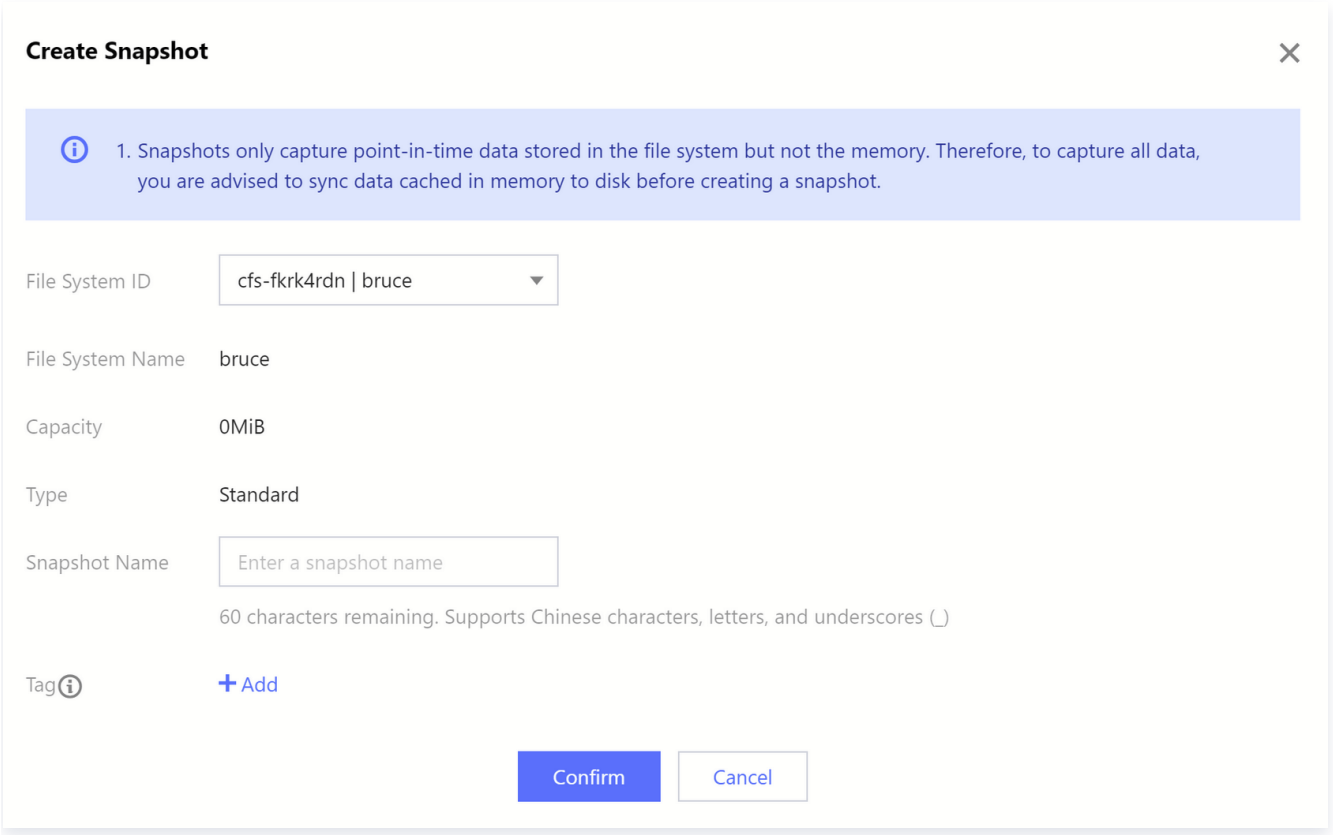

Scheduled Snapshots

Last updated: 2023-08-29 11:45:57

Overview

Tencent Cloud File Storage offers a **Scheduled Snapshot** feature, which allows developers to flexibly set backup task policies. It is recommended to adopt different scheduled snapshot policies for different businesses, as suggested in the following table:

Scenario	Snapshot Policy	Recommended 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	Note
Objects	All file systems
Execution 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 Deletion (**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 for them.
Naming rule	The naming convention for automatic snapshots is auto_policy_cfsidyyyyMMddHH. Here, cfsid is the file system ID, yyyyMMdd is the date, and HH is the hour. You can also manually modify the snapshot name. For instance, auto_policy_cfs-2cj5yj0f2021090923 represents an automatic snapshot created for the file system cfs-2cj5yj0f at 23:00 on September 9, 2021.
Lifecycle (**Important**)	Snapshot lifecycles vary depending on the snapshot creation method: • Manually created snapshots are set to **long-term storage** by default and can be retained indefinitely, provided the account balance is sufficient. • Scheduled snapshots, based on the creation rules, can be set to a **regularly scheduled deletion** time, or they can be configured for long-term retention.
Snapshot conflict	Scheduled snapshots do not conflict with custom snapshots in use. However, they may conflict with each other on the creation time. • While an automatic snapshot is being created for a file system, you must wait for the automatic snapshot to complete before you can create a custom snapshot, and vice versa. • If the data volume of the file system is large and the duration of a snapshot exceeds the interval between two automatic snapshot points, the next automatic snapshot will be skipped. For instance, if a user sets 9:00, 10:00, and 11:00 as automatic snapshot points, and the snapshot at 9:00 takes 70 minutes (i.e., it is completed at 10:10), then the automatic snapshot at 10:00 will not be executed, and the next snapshot point will be at 11:00.
Snapshot quota	Each file system has a certain snapshot quota. If the number of snapshots for a file system reaches the quota limit, the automatic snapshot task will be suspended and blocked. The snapshot quota is primarily designed to prevent developers from forgetting about an automatic snapshot policy, which could lead to an endless increase in storage costs.
ASP	Indicates the scheduled snapshot policy, that is, Auto Snapshot

	Policy.
ASP quota	Under a single Tencent Cloud account, a maximum of 30 ASP policies can be set for each region. A single ASP can be associated with up to 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	The ASP Scheduled Snapshot Policy provides a manual **Pause** feature. Once paused, no new snapshots will be automatically created. However, the lifecycle of already generated automatic snapshots is not affected by the pause feature and will still be periodically destroyed or preserved long-term according to the set rules.
Operations log	Show the creation process of all scheduled snapshots, same as that of manually created snapshots.

Instructions

Creating a scheduled snapshot policy

> **Note**
> A single Tencent Cloud account can establish up to 30 scheduled snapshot policies within the same region.

1. Log in to the File System console and navigate to the Snapshot Policy page.
2. Select a region.
3. Click **Create Policy**.
4. In the "Create Snapshot Policy" page, set the following parameters and click **Confirm**. As illustrated below:

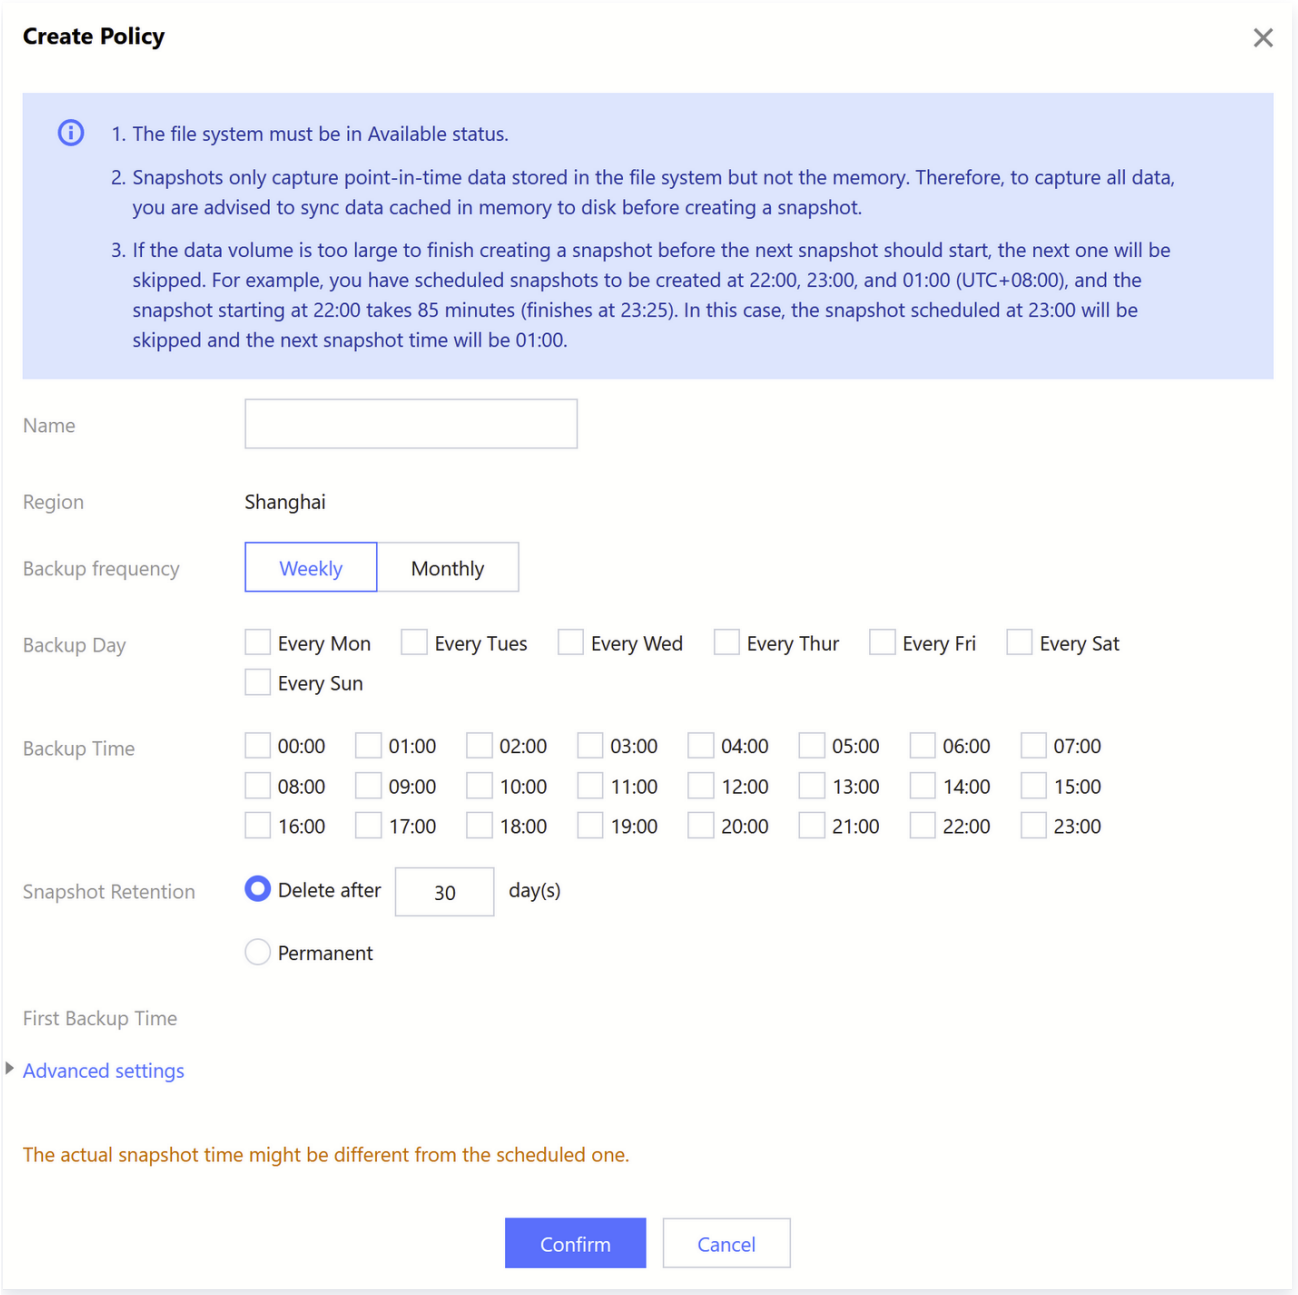

Configuration Item	Description
Name	Required. The name of a scheduled snapshot policy, supporting up to 60 characters.
Region	Required. The parameter on this page cannot be modified. For specific setting methods, please refer to Step 2.
Backup Day	Required. The date to execute the scheduled snapshot can be selected on a weekly/monthly basis.

Backup Time	Required. The time to execute scheduled snapshots can be selected within the range of 00:00 – 23:00 at each whole hour. Please note that due to backend operational conditions, there may be a discrepancy between the set snapshot time and the actual snapshot creation time in the console. The data within the snapshot is based on the actual creation time in the console.
Recommended Snapshot Retention	Required. • Automatically deleted after a fixed number of days, ranging from 1 to 30. The default retention period is 30 days. • Permanently stored.

Assigning a file system

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

1. Log in to the File System console and navigate to the Snapshot Policy page.
2. Select a region.
3. Click on the **Target Policy** name to enter the detailed interface, and then click on **Associate File System** below.
4. On the "Assign File System" page, select the file systems you wish to associate, as shown below:

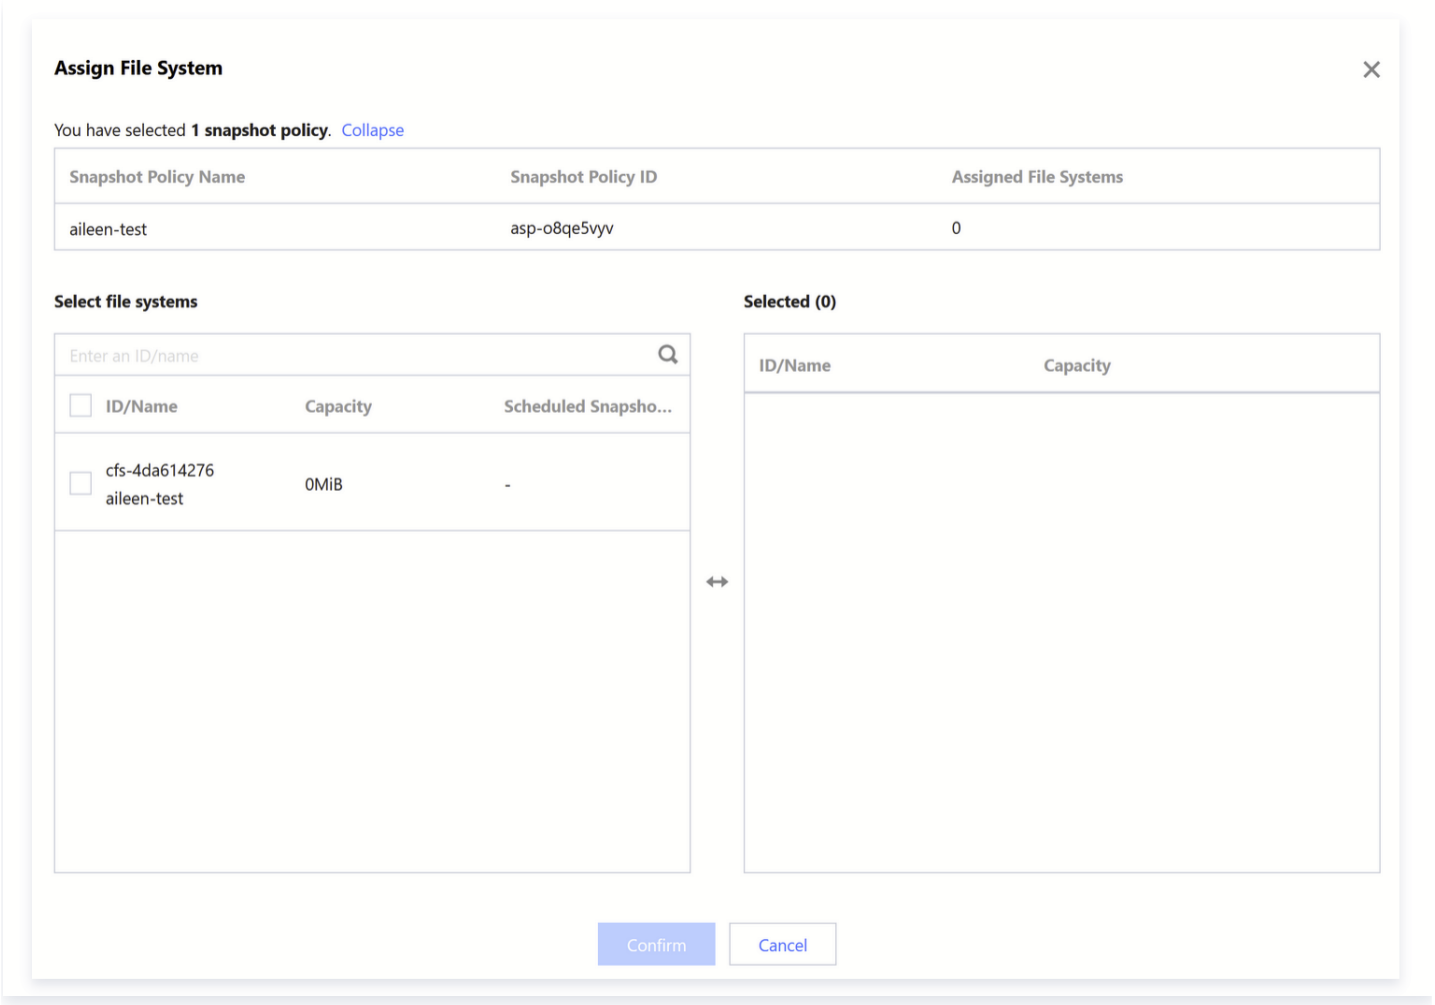

5. Click **Confirm**.

Enable/Disable scheduled snapshot policies

1. Log in to the File System console and navigate to the Snapshot Policy 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 scheduled snapshot policies

1. Log in to the File System console and navigate to the Snapshot Policy page.
2. Select a region.
3. Locate the row of the target policy and click **Edit Policy**.
4. In the "Modify Snapshot Policy" page, adjust the relevant parameters (for parameter descriptions, please refer to Step 4) and click **OK**.

Deleting scheduled snapshot policies

1. Log in to the File System console and navigate to the Snapshot Policy page.
2. Select a region.
3. Delete a scheduled snapshot policy as follows:
   - Single Deletion: Select **More > Delete** on the right side of the target policy row.
   - Batch Deletion: Select the scheduled snapshot policies you wish to delete and click **Delete** at the top of the list.

Turning automatic snapshots into long-term stored snapshots

> **Note**
> If the **Snapshot Retention Period** in the automatic snapshot policy is set to **Permanent**, there is no need to perform the following operations on the snapshots generated by this policy.

1. Log in to the File System console and navigate to the Snapshot Policy page.
2. Select a region.
3. Click the ID of the target automatic snapshot.
4. In the details page, click on **Retain Permanently** to set the automatic snapshot to be permanently retained. The snapshot retention period will then be set to permanent, as shown in the figure below:

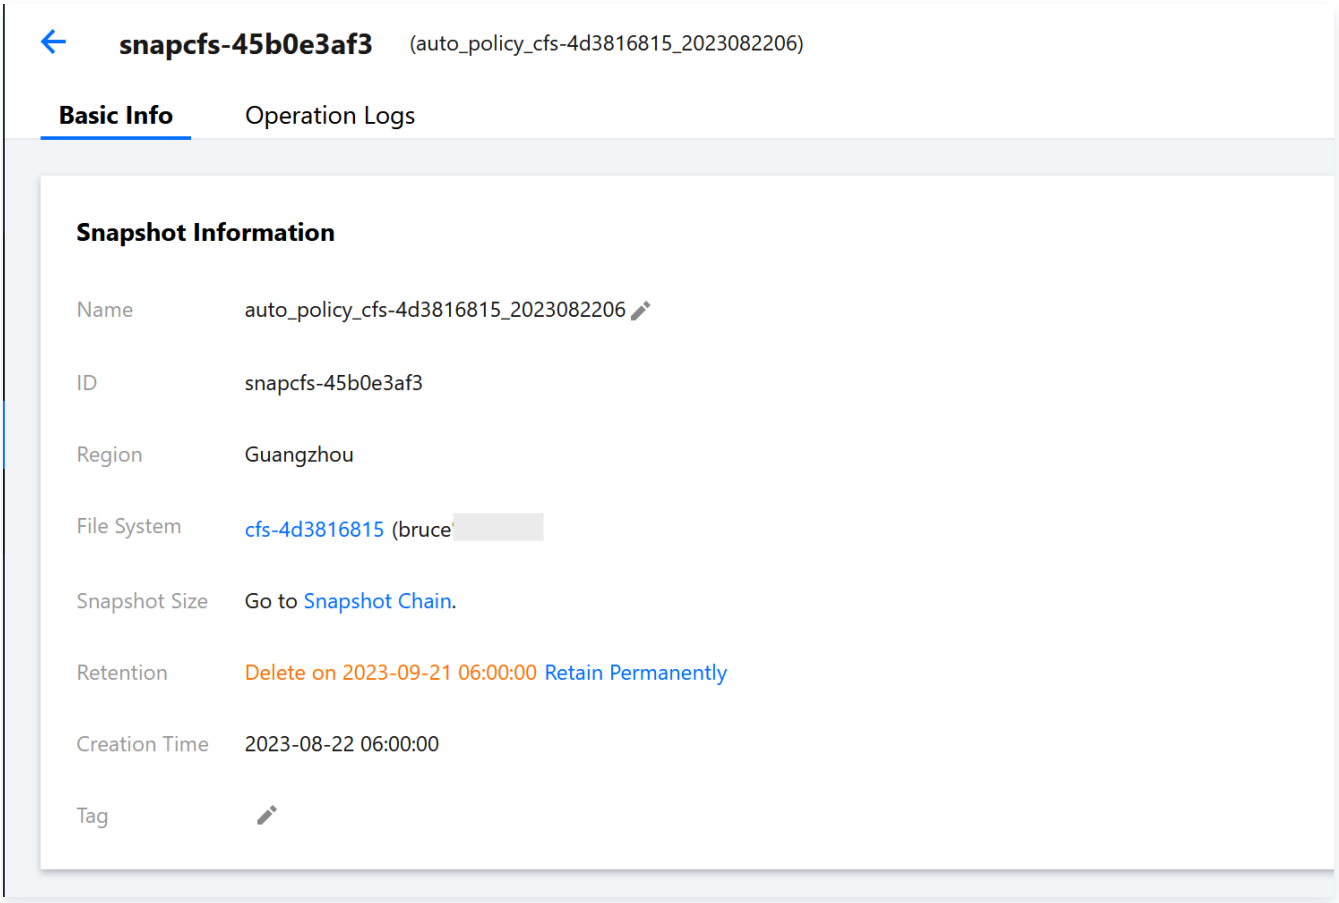

Creating a File System Using a Snapshot

Last updated: 2023-08-29 10:23:11

Scenario

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 guide will walk you through the process of creating a file system from a snapshot on the snapshot list page. Additionally, you can specify a snapshot to create a file system by configuring the **snapshot** parameter during the file system creation process.

Instructions

1. Log in to the file system console and navigate to the Snapshot List page.
2. In the row of the target snapshot, click **Use**.
3. On the **Create File System** page, select the file system type, click **Next: Detailed Settings**, and configure the following parameters.

> **Note**
> The following describes the process of creating a Standard type file system using a snapshot as an example.

Parameter	Description
Billing Mode	Select the 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.
Regions	Select the region where the file system is to be created.
Availability Zones	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 parameters. When creating a file system using a snapshot, you need to check **Create a file system using a snapshot** and select the snapshot you want to use. If you choose to create a file system from a snapshot, the initial data volume in the file system will be consistent with the size of the snapshot.
Permission Group	Each file system must be bound to a permission group. The permission group specifies an allowlist that can access the file system and lists the read and write permissions.
Scheduled Snapshot	Optional. During the file system creation, you can choose to schedule snapshots based on the already created snapshot policies. This will periodically create snapshots for the file system. For more detailed information about scheduled backups, please refer to 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	For Turbo series only, a Cloud Connect Network (CCN) must be specified. You can either select an existing CCN or create a new one. For more information, refer to Introduction to CCN.
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 yet have a tag, please first create the necessary tags in the Tag Console, and then bind the tags to the file system. Alternatively, you can add tags to the file system after it has been created.

4. Click **Create Now** to view the newly created file system in the file system list.

Deleting Snapshots

Last updated: 2023-08-29 10:23:24

Scenario

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

Supports and Limits

- 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 the continued use of snapshots created later.
- **When a snapshot is deleted, all data within the snapshot will be simultaneously erased and cannot be retrieved. Once deleted, snapshots cannot be restored. Please proceed with caution.**

Instructions

1. Log in to the file system console, go to the Snapshot Overview page, and check the usage of snapshots in various regions.

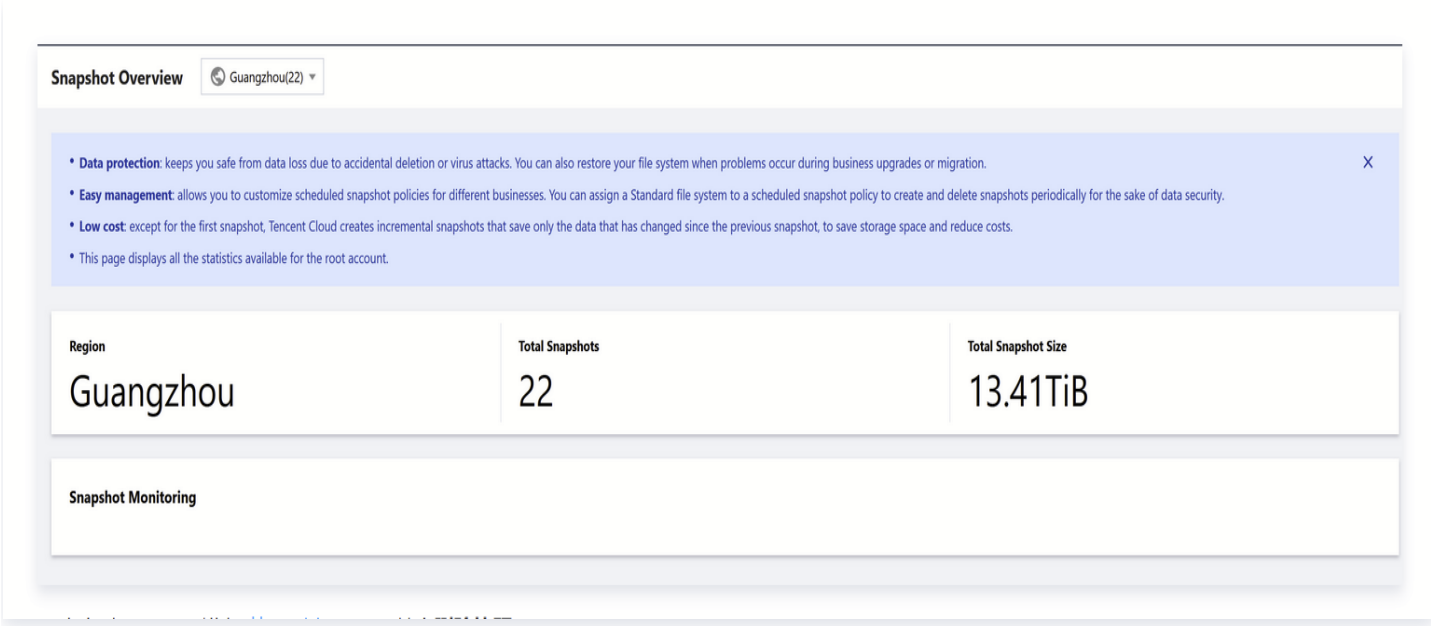

2. Based on your actual needs, navigate to the Snapshot List page and click **Delete Snapshot**.

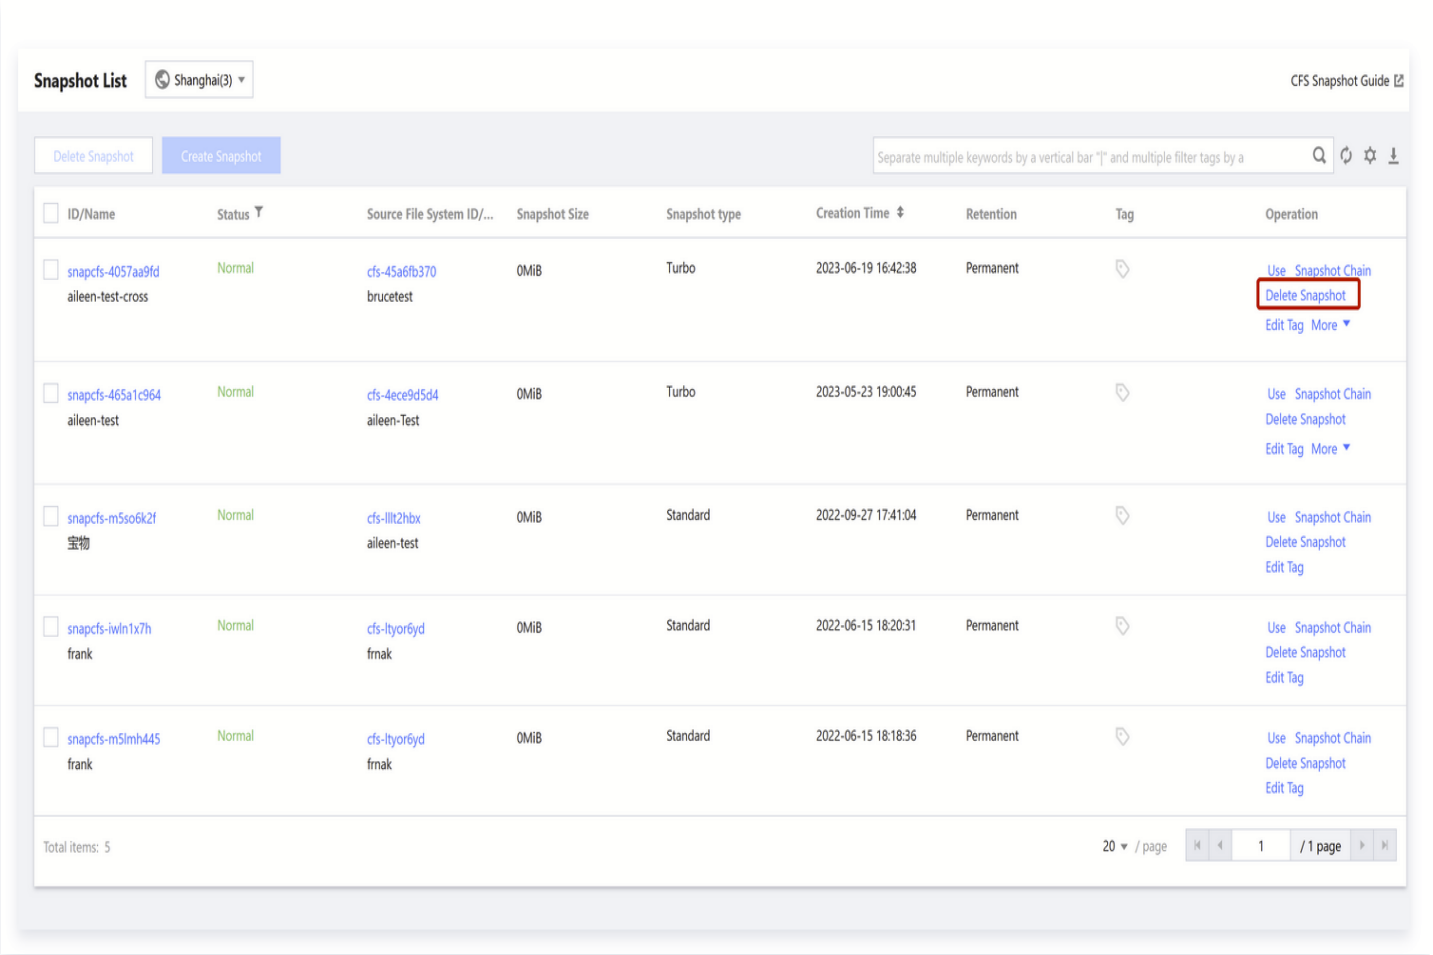

3. Confirm the information of the snapshot to be deleted:

   3.1 After confirming the details of the list to be deleted, click **Next**.

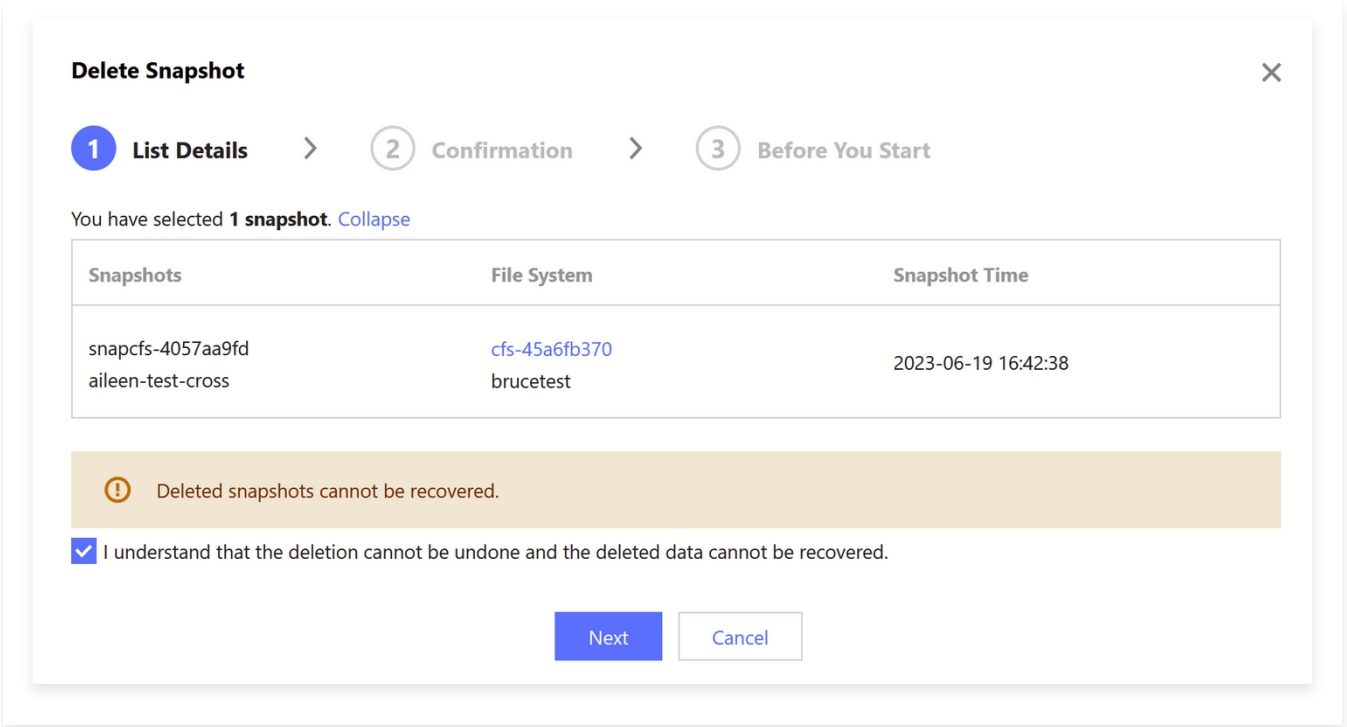

3.2 Confirm the information and click **Next.**

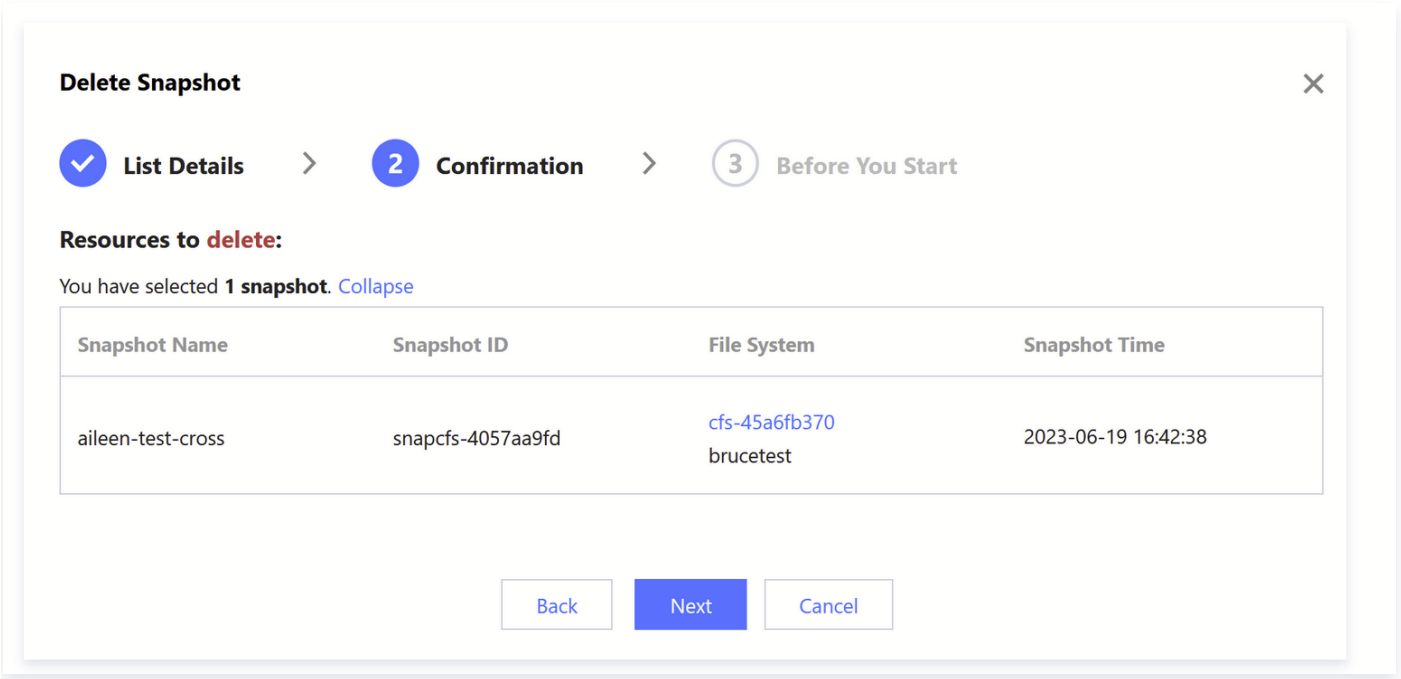

3.3 After reading the operation instructions, click **Confirm** to delete.

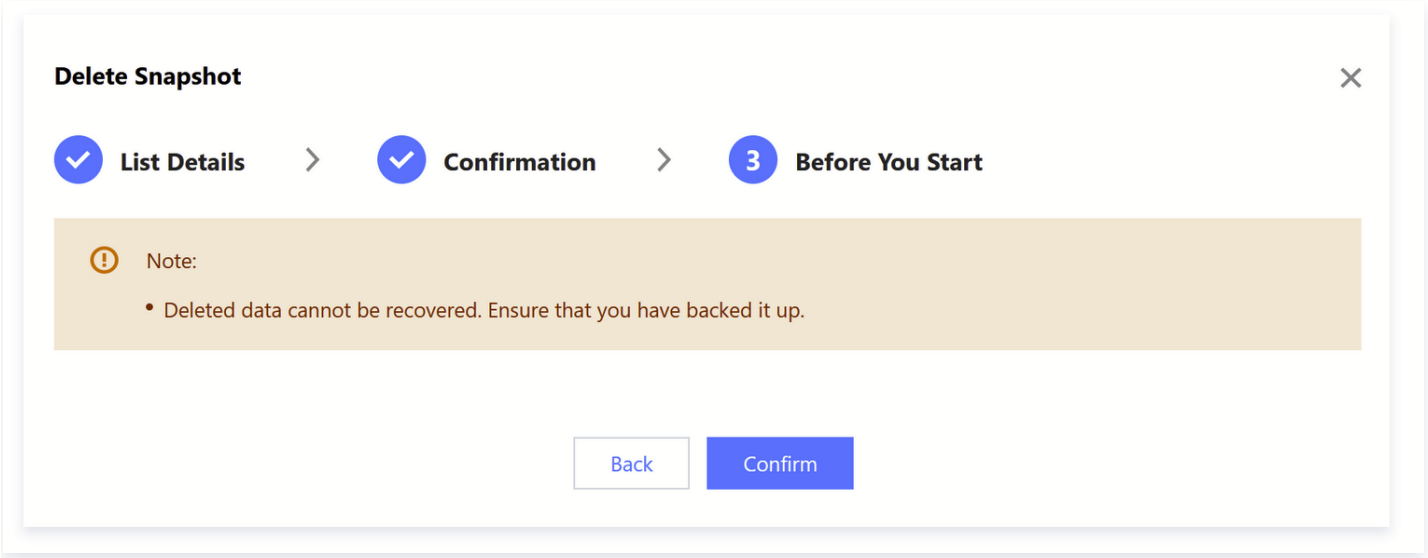

Guide for Cross-AZ and Cross-Network Access

Last updated: 2023-08-29 10:27:17

Cross-AZ access in a VPC

If you need to have a CFS instance shared by multiple 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.
Taking Guangzhou as an example, if you already have a CVM instance in Guangzhou Zone 1 and need to use CFS, but the resources in Guangzhou Zone 1 are sold out and a file system cannot be created directly. You can log in to VPC subnet and create a subnet for this VPC in "Guangzhou Zone 3".

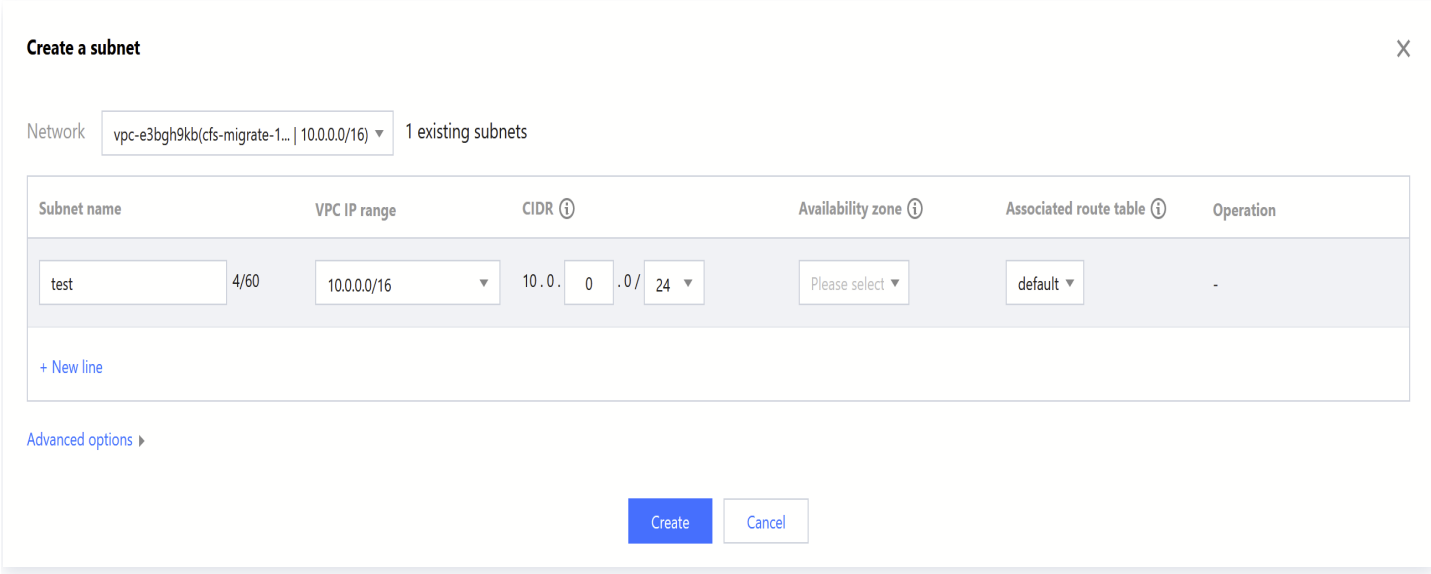

After the subnet is successfully created, return to the CFS console and select this VPC and the newly created subnet when creating resources in Guangzhou Zone 3. At this point, the CVM instance in the Guangzhou Zone 1 subnet of this VPC can directly mount the CFS. Refer to File System Mounting Help for more information.

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 interconnect CVM instances in VPC-A and VPC-B with a CFS instance in VPC-C by establishing a Cloud Connect Network (CCN), enabling cross-access among VPC-A, VPC-B, and VPC-C. For more information, refer to CCN Quick Start Guide.

Automatically Mounting File Systems

Last updated: 2023-08-29 10:27:32

Scenario

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

Instructions

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. Then, 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 Point IP>:/ <Target Directory to be Mounted> nfs
vers=4,minorversion=0,hard,timeo=600,retrans=2,_netdev,noresvport 0 0
Example: 10.10.19.12:/ /local/test nfs
vers=4,minorversion=0,hard,timeo=600,retrans=2,_netdev,noresvport 0 0
```

```
Mount the file system with NFS v3.0
<Mount Point IP>:/<fsid> <Target Directory to be Mounted> nfs
vers=3,nolock,proto=tcp,hard,timeo=600,retrans=2,_netdev,noresvport 0
0
Example: 10.10.19.12:/djoajeo4 /local/test nfs
vers=3,nolock,proto=tcp,hard,timeo=600,retrans=2,_netdev,noresvport 0
0
```

```
Mount the file system with Turbo
<Mount Point IP>@tcp0:/<fsid>/cfs <Target Directory to be Mounted>
lustre defaults,_netdev 0 0
```

```
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**
> If an automatic mount command has been added, but the shared file system is in an abnormal state, the Linux system may not start normally, as it needs to successfully execute the automatic start commands in fstab to start properly. In this case, you can enter "single-user mode" during system startup, delete the automatic mount command in fstab, and then restart the host.

Automatically mounting a file system on Windows

During the mounting process, select the "Reconnect at sign-in" option, as shown below. For more mounting assistance, please refer to Using CFS File System on Windows Client.

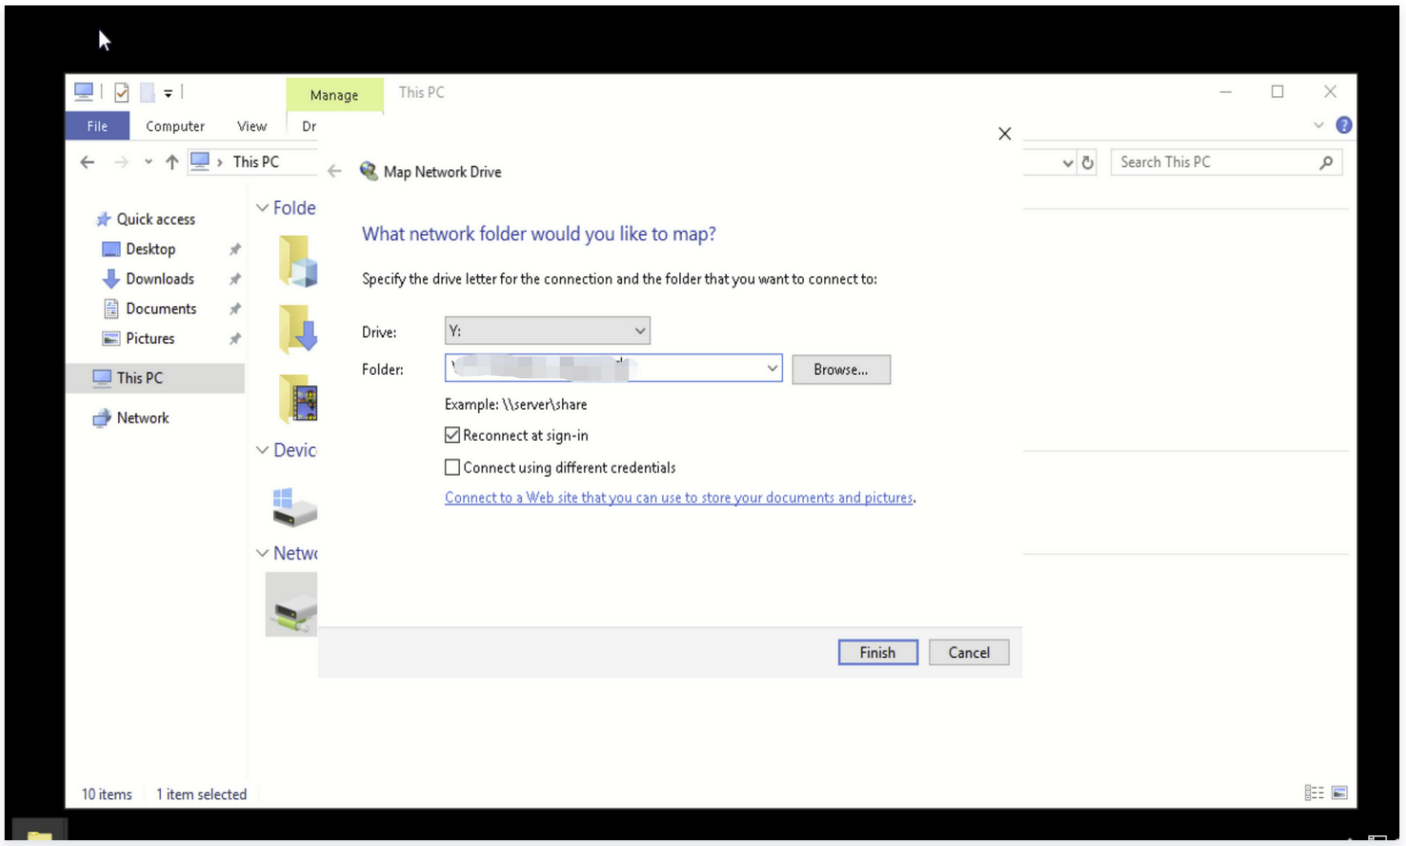

Data Migration Service Features

Last updated: 2023-08-29 10:27:48

Overview

The data migration service is an online concurrent migration service provided by 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.

Scenarios

Loading data

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.

Strengths

- 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.
- Flexible overwriting methods.
  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.

Limits

Last updated: 2023-08-29 10:28:00

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

Item	Description
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 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: 2023-08-29 10:28:19

Scenario

This document describes how to start a migration task.

Instructions

Step 1: Establish a Migration Task

1. Log in to the CFS console and navigate to the Data Migration Service page.
2. Click on **Create Migration job**.

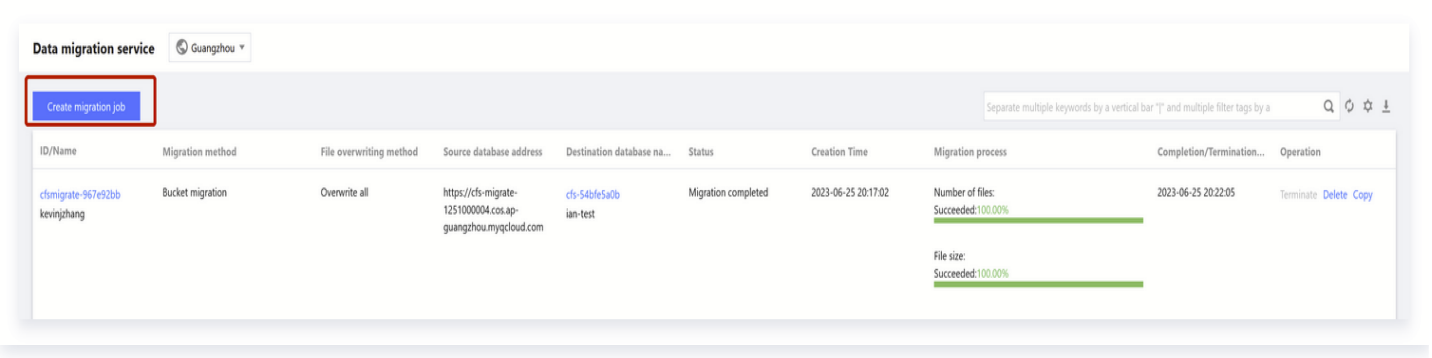

> **Note**
> This feature is currently available on a whitelist basis. If you wish to use it, please submit a ticket to contact us.

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**.

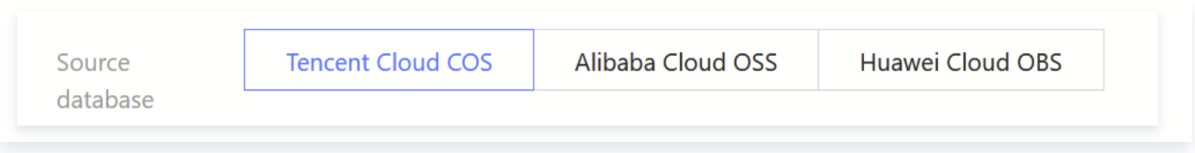

Step 4. Select the migration method

Currently, bucket migration and inventory migration are supported:

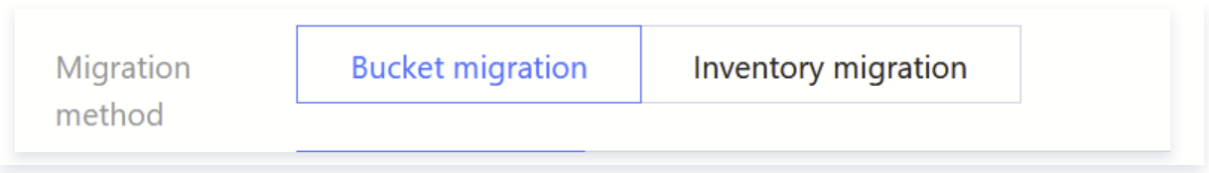

- Bucket migration: It is suitable for migrating an object storage bucket or all objects in a certain directory.
- Inventory migration: It is suitable for migrating objects within a specified time period. Currently, you can export the objects modified after a time point or within a time period through the object storage inventory for filtered data migration.

Step 5. Select the overwriting method

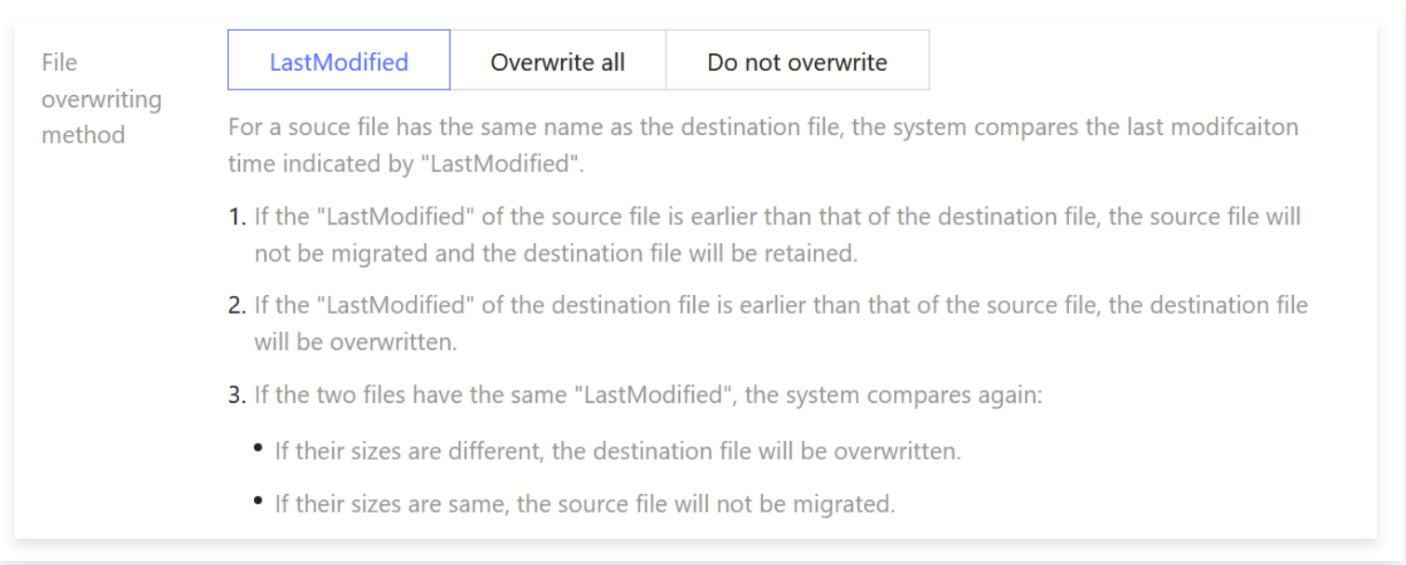

Currently, overwriting by last modified time, full overwriting, and no overwriting are supported.

- 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` 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` 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 Overwrite: Contrary to the **Full Overwrite** policy, files with the same name are skipped 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.

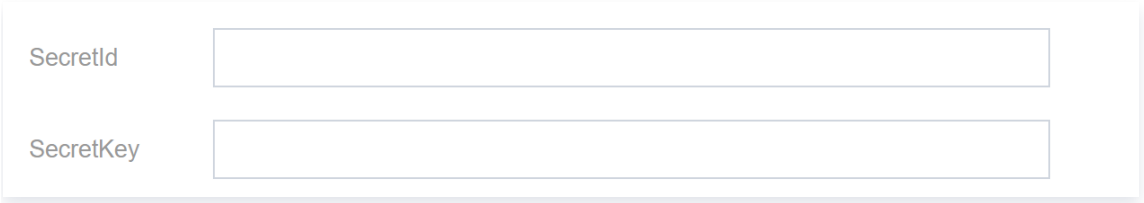

Step 7. Enter the source information (required for bucket migration only)

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 and other subpaths end with "/".

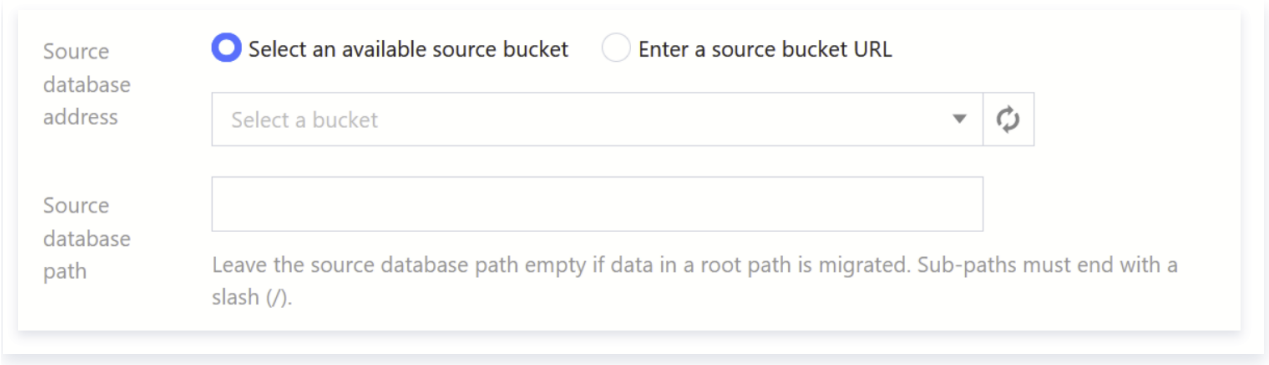

Step 8. Enter the inventory URL (required for inventory migration only)

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 the inventory URL.

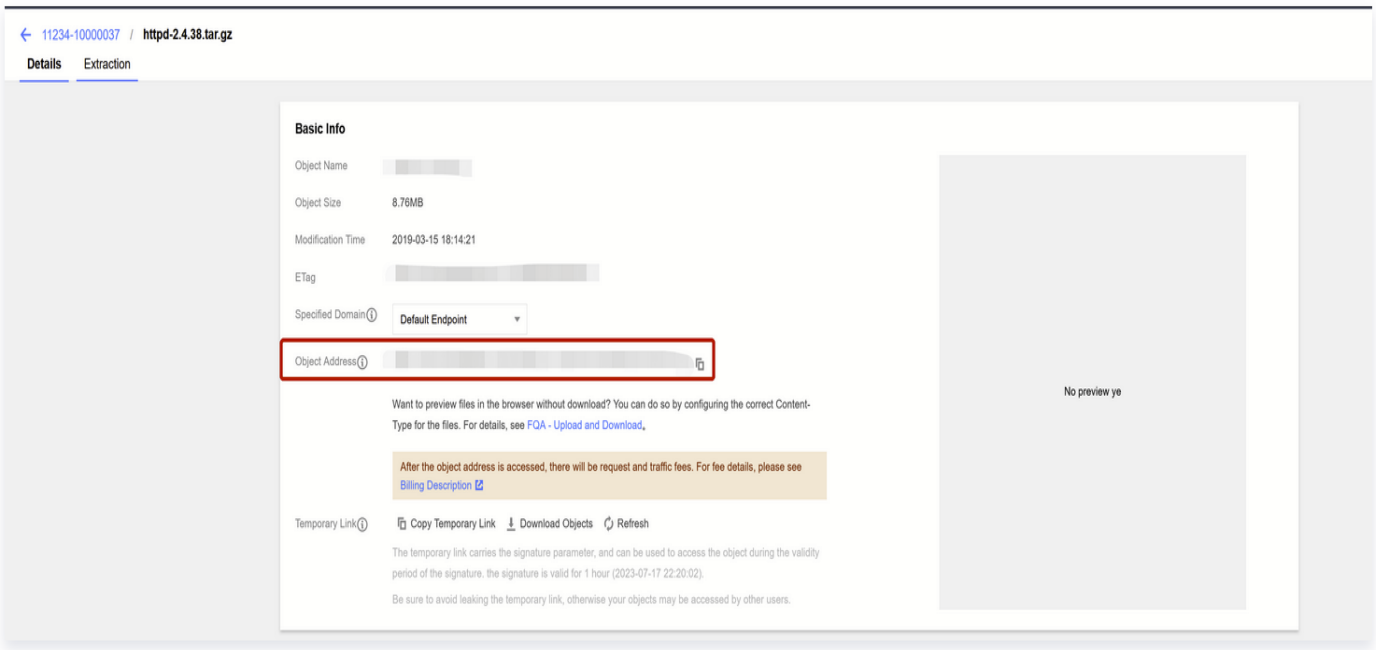

Step 9. Select the destination file system

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/`.

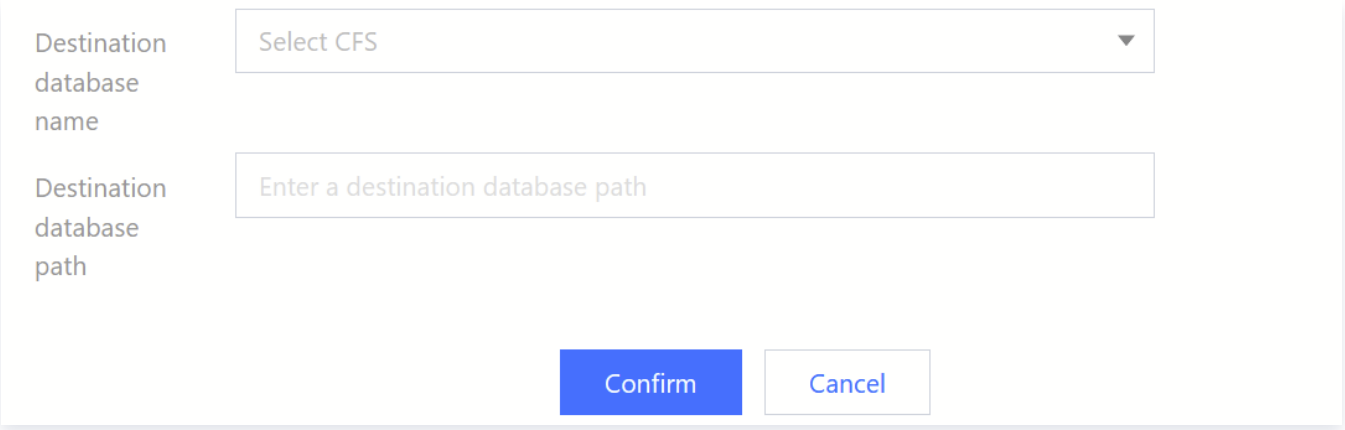

Viewing Migration Task Result

Last updated：2023-08-29 10:28:28

Scenario

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

Preparations

You have logged in to the CFS console.

Instructions

Viewing results in the data migration service panel

On the Data Migration Service page, you can select the corresponding region to view basic migration task information, such as migration settings, progress, etc.

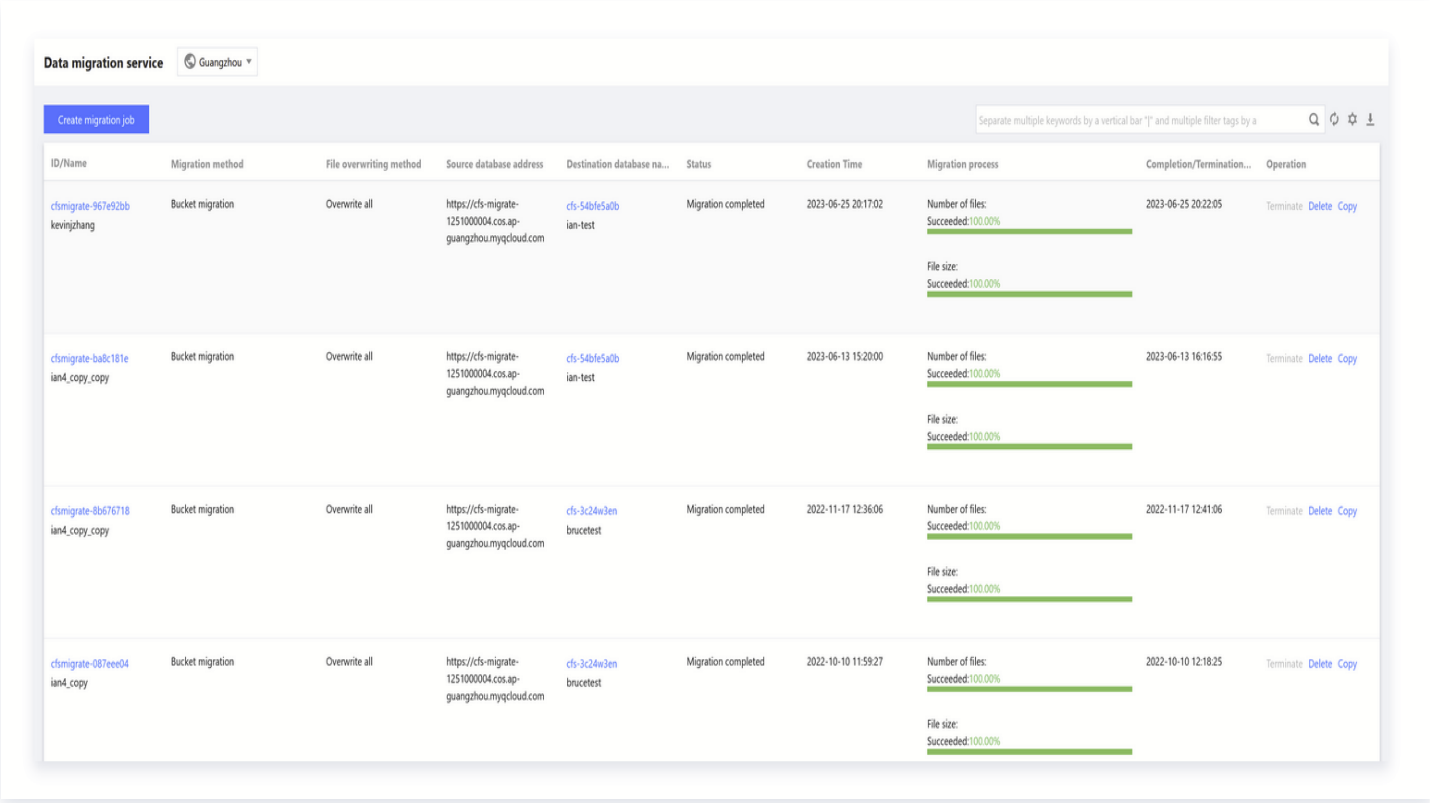

Viewing detailed information

On the Data Migration Service page, click on the ID/name of the migration task to view detailed migration information, such as total number/capacity of files to be migrated,

number/capacity of files awaiting migration, number/capacity of files already migrated, and number/capacity of files where migration failed.

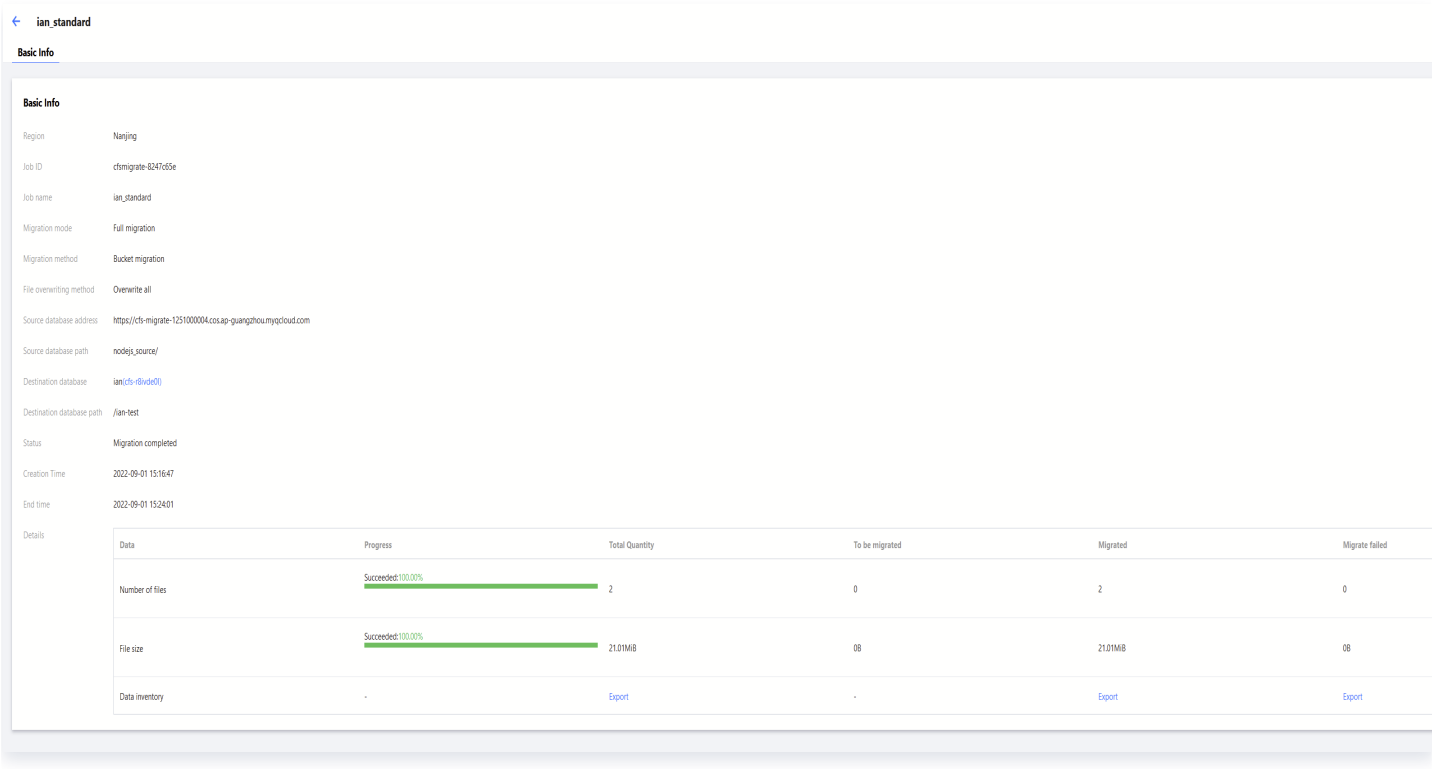

User Permission Management

Last updated: 2023-08-29 10:29:20

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.

Preparations

The file system has been mounted using either the Turbo protocol or NFS V3. For detailed instructions, please refer to Using CFS Turbo File System on Linux Clients and Using CFS File System on Linux Clients.

Command description

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

Sample

```
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 to read and execute file1
sudo setfacl -m u:otheruser:rwx file1 # Grant the otheruser user the
permission to read/write and execute file1
```

User Quotas
Features

Last updated: 2023-08-29 10:29:37

Overview

CFS user quota is a resource management feature for data management involving multiple users. It enables you to set capacity quotas and quotas of files 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.

Scenarios

Storage Space Allocation for Multiple Departments/Tenants
Different capacity usage limits need to be set for different final 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, making storage resources used more efficiently.

Strengths

User-specific capacity quota and quota of files

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

- 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 quota of files 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.

Operation Guide

Last updated: 2023-08-29 10:32:09

Setting quotas

1. Access the quota settings interface: Log in to the File System Console and click on **Create Quota**.

> **Note**
> This feature is currently available to allowlist users only. If necessary, you can submit a ticket to contact us.

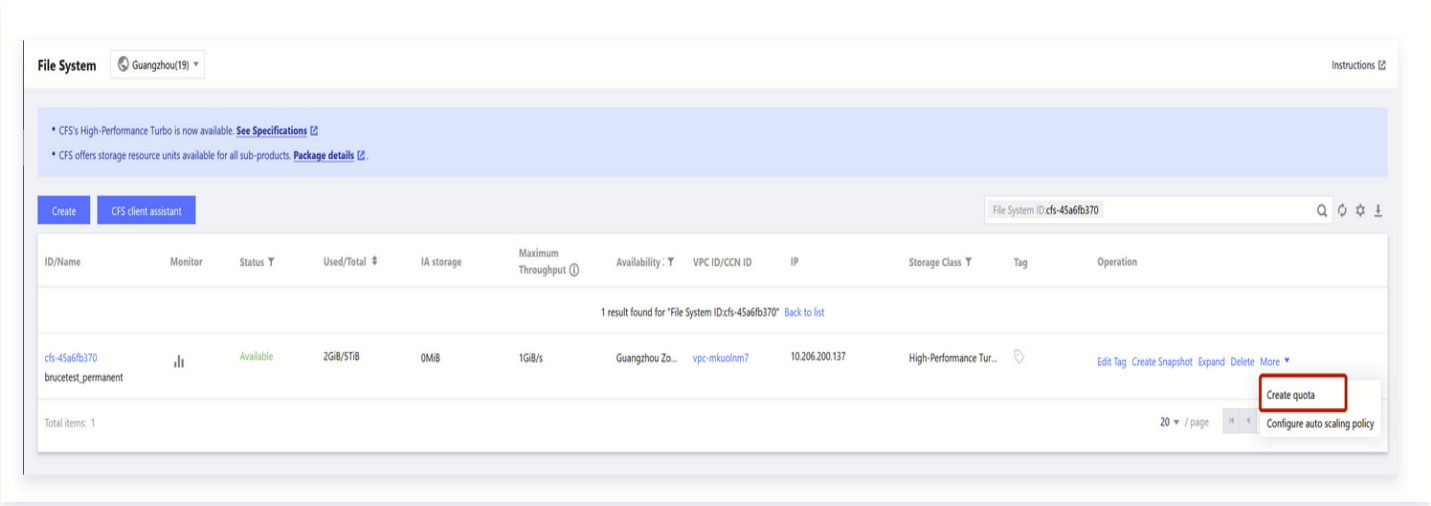

2. To create a user quota: Select the quota type as Uid (User ID), enter the Uid number and the corresponding capacity and file count quota values, then click **Confirm**.

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

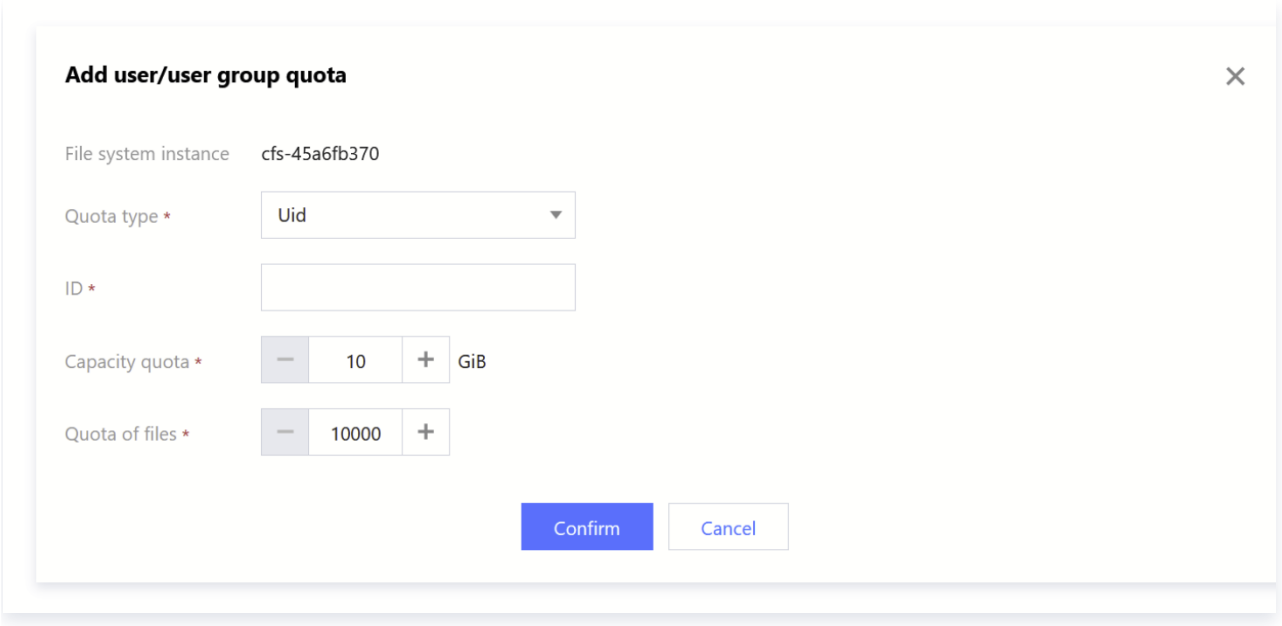

3. To create a user group quota: Select the quota type as Gid (Group ID), enter the Gid number and the corresponding capacity and file count quota values, then click **Confirm**.

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

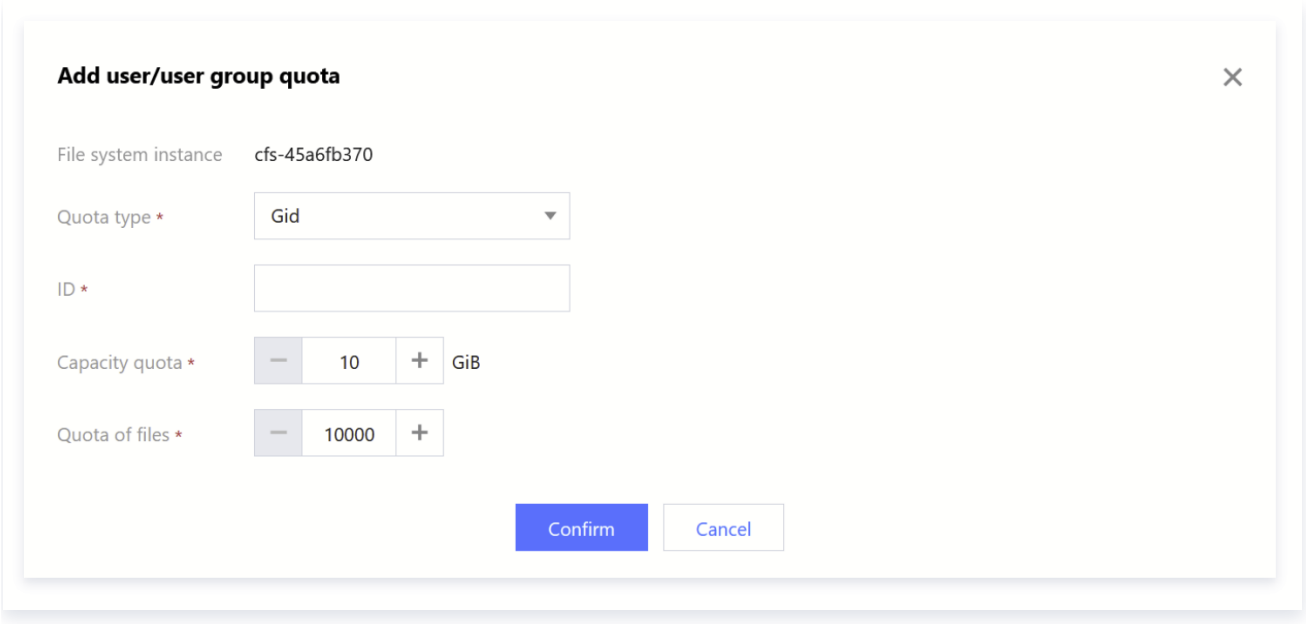

Viewing quotas

Click on **File System** to enter the details page. Select **Quota Information** at the top to view.

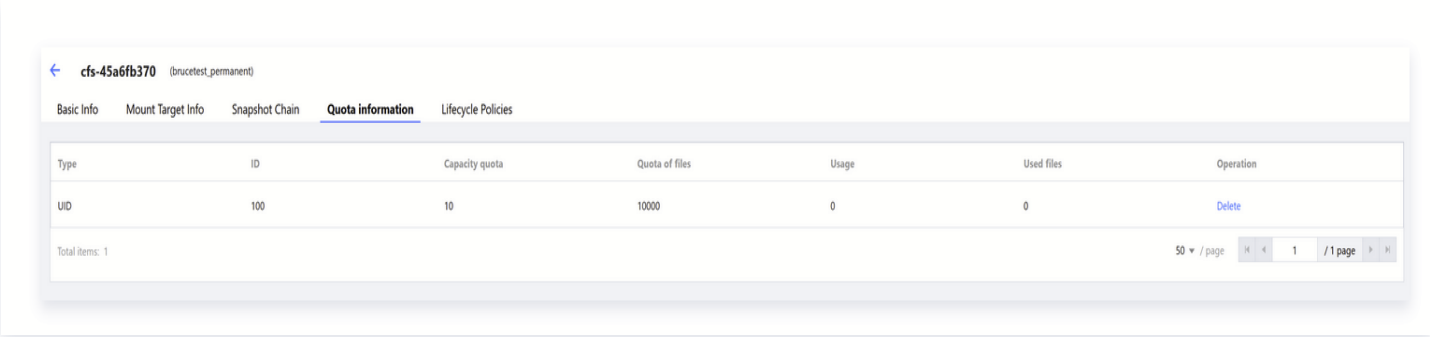

Modifying quotas

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

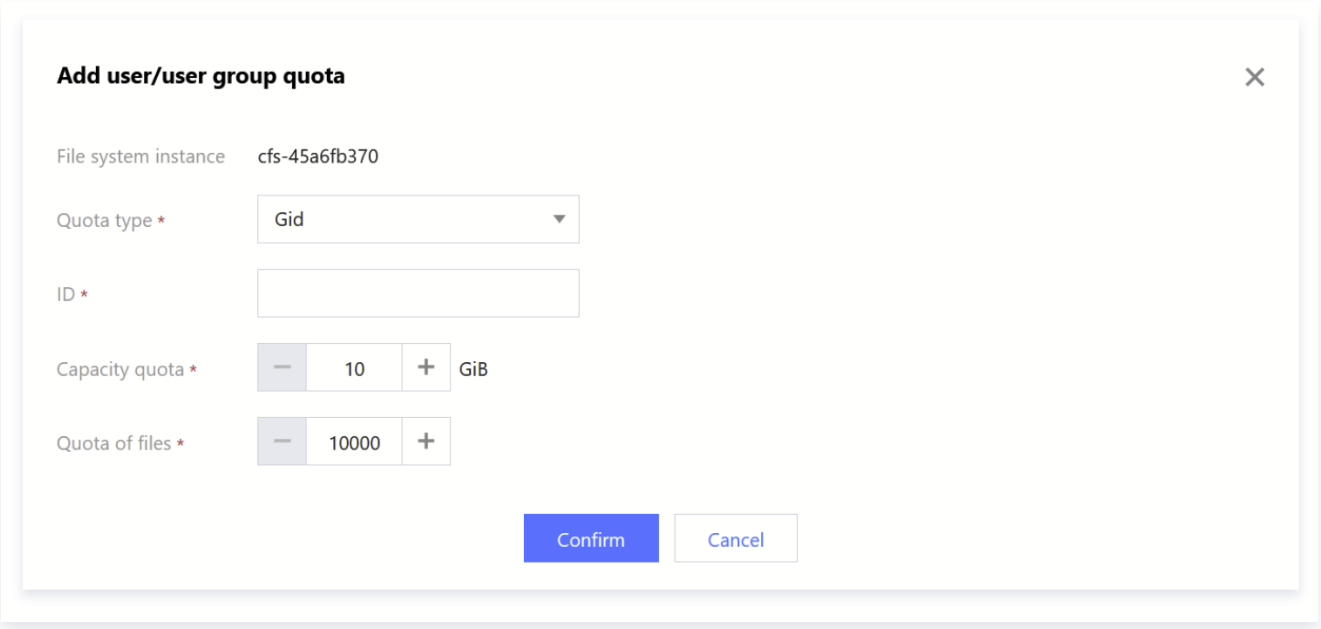

Deleting quotas

Click on **File System** to access the details page, then click **Delete** on the right to proceed.

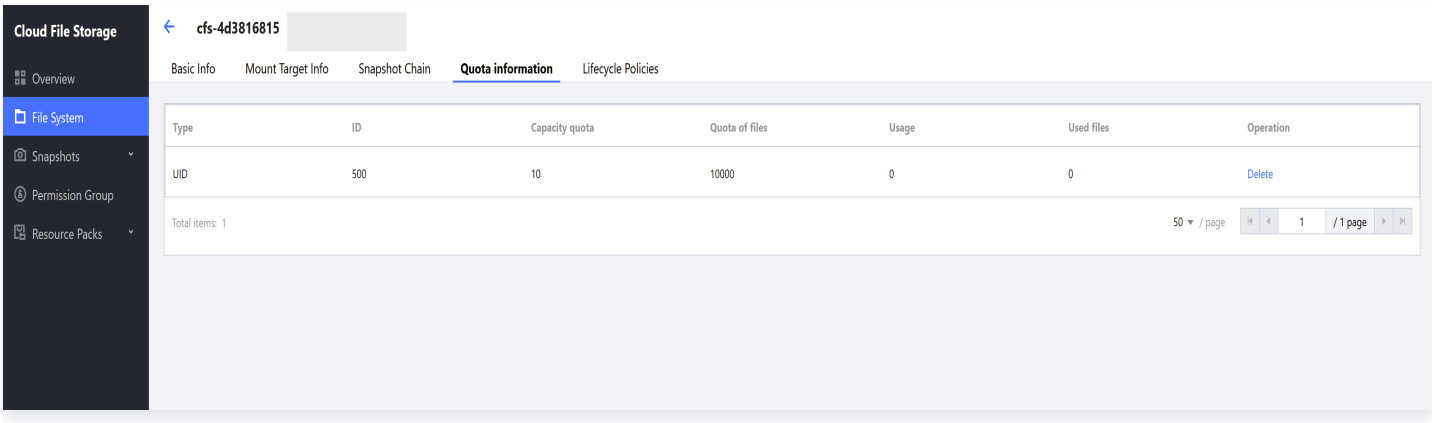

Limits

Last updated: 2023-08-29 10:29:48

> **Note**
> This feature is exclusively available for Turbo file systems and is currently in a popular public beta. To use it, please submit a ticket to get in touch with us.

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

Item	Description
Capacity quota	10 GiB to 1,000 TiB, with a minimum step of 1 GiB.
Quota of files	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 write-in exceeding the limit.
> - If both capacity quota and quota of files are small, the capacity actually available for write-in and number of files allowed will be greater than the quota values.

Data Encryption Features

Last updated: 2023-08-29 10:32:28

Overview

You can enable encryption for your 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. Keys are stored in CFS 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 plainly stored in any permanent storage medium.

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 the Turbo file systems only.
- The encryption feature involves data encryption and decryption and will cause 10%-15% performance loss to the file system. We recommend you use it based on your actual needs.

Operation Guide

Last updated：2023-08-29 11:44:45

Setting quotas

Step 1: Accessing the Console

Log in to the File System Console and click **Create** to establish a new file system.

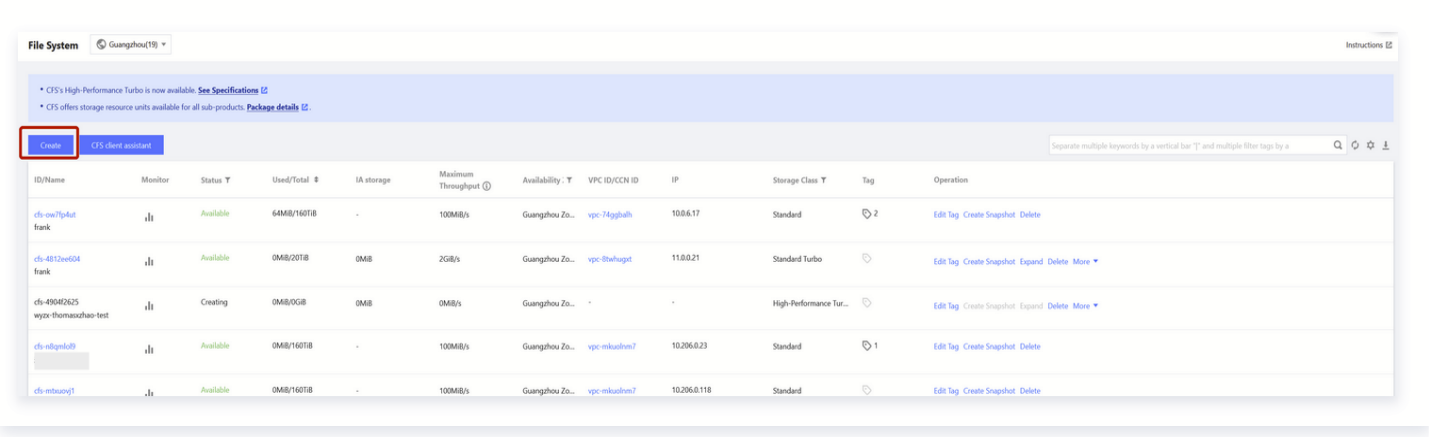

Step 2: Select the file system type (Turbo)

Choose the Turbo file system, then click **Next:Set Up Details.**

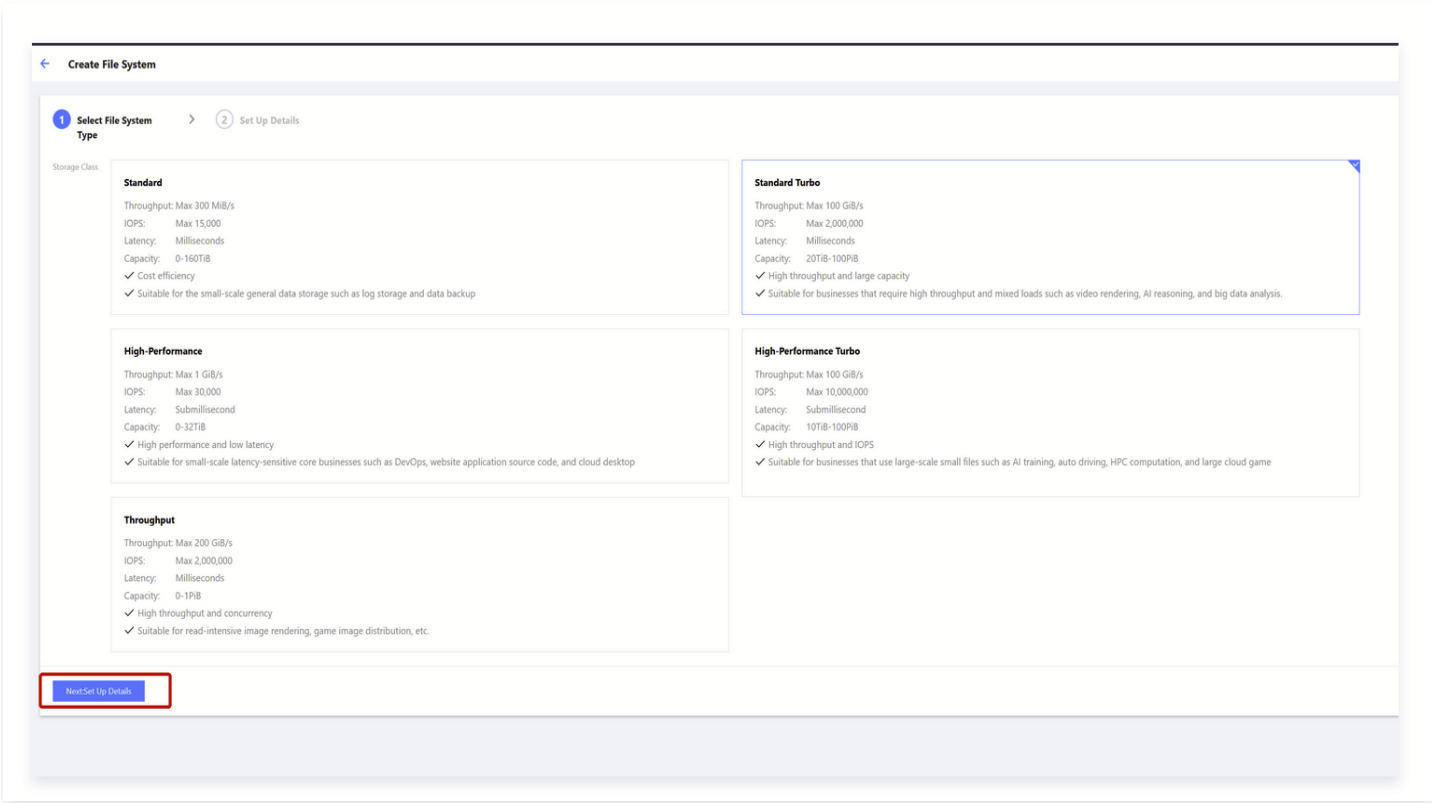

Step 3: Enable encryption

1. Select the Encryption Option

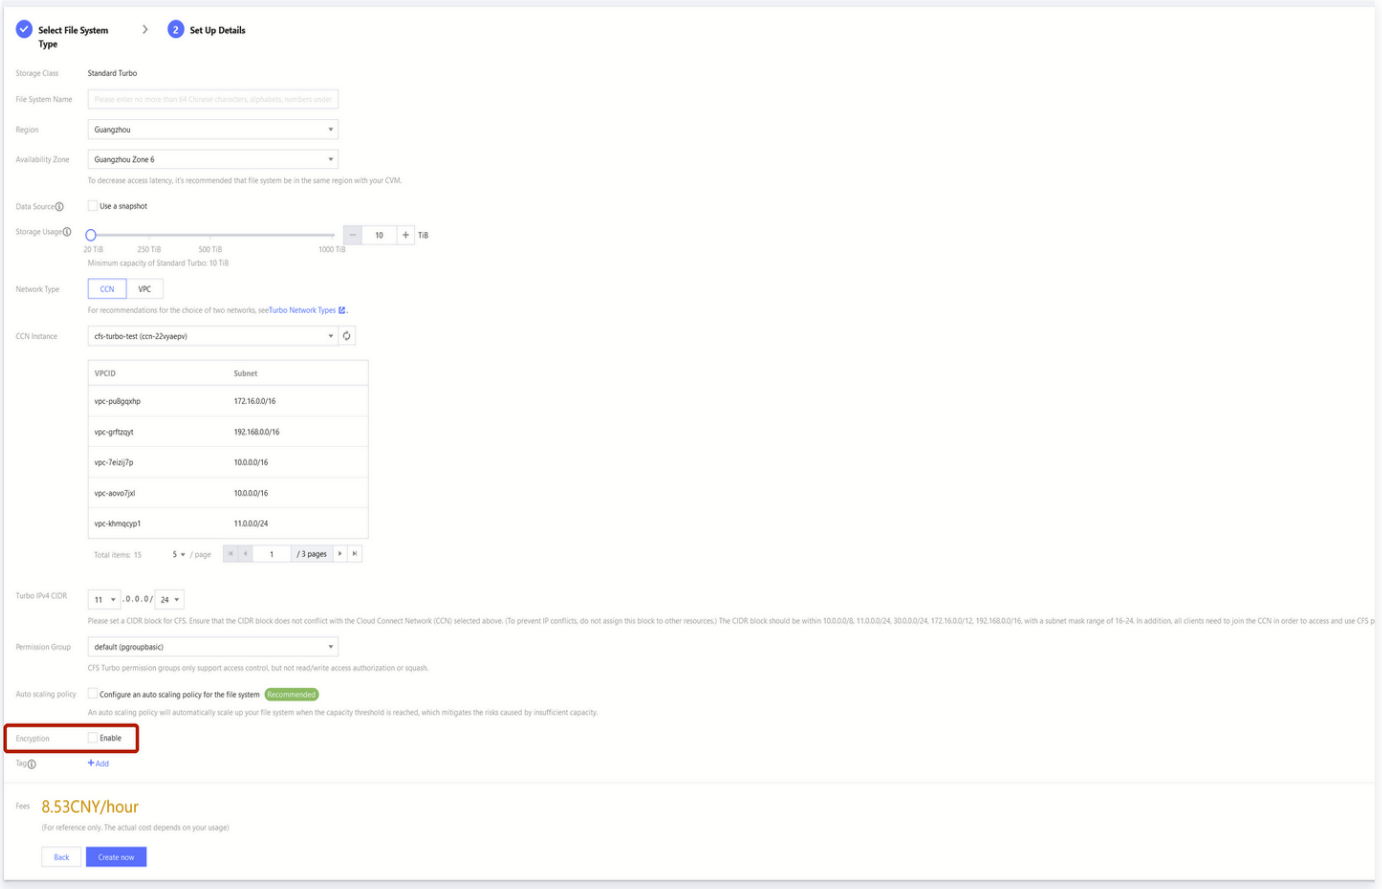

2. **Encryption Status**

   You can check whether the file system is encrypted on the file system details page.

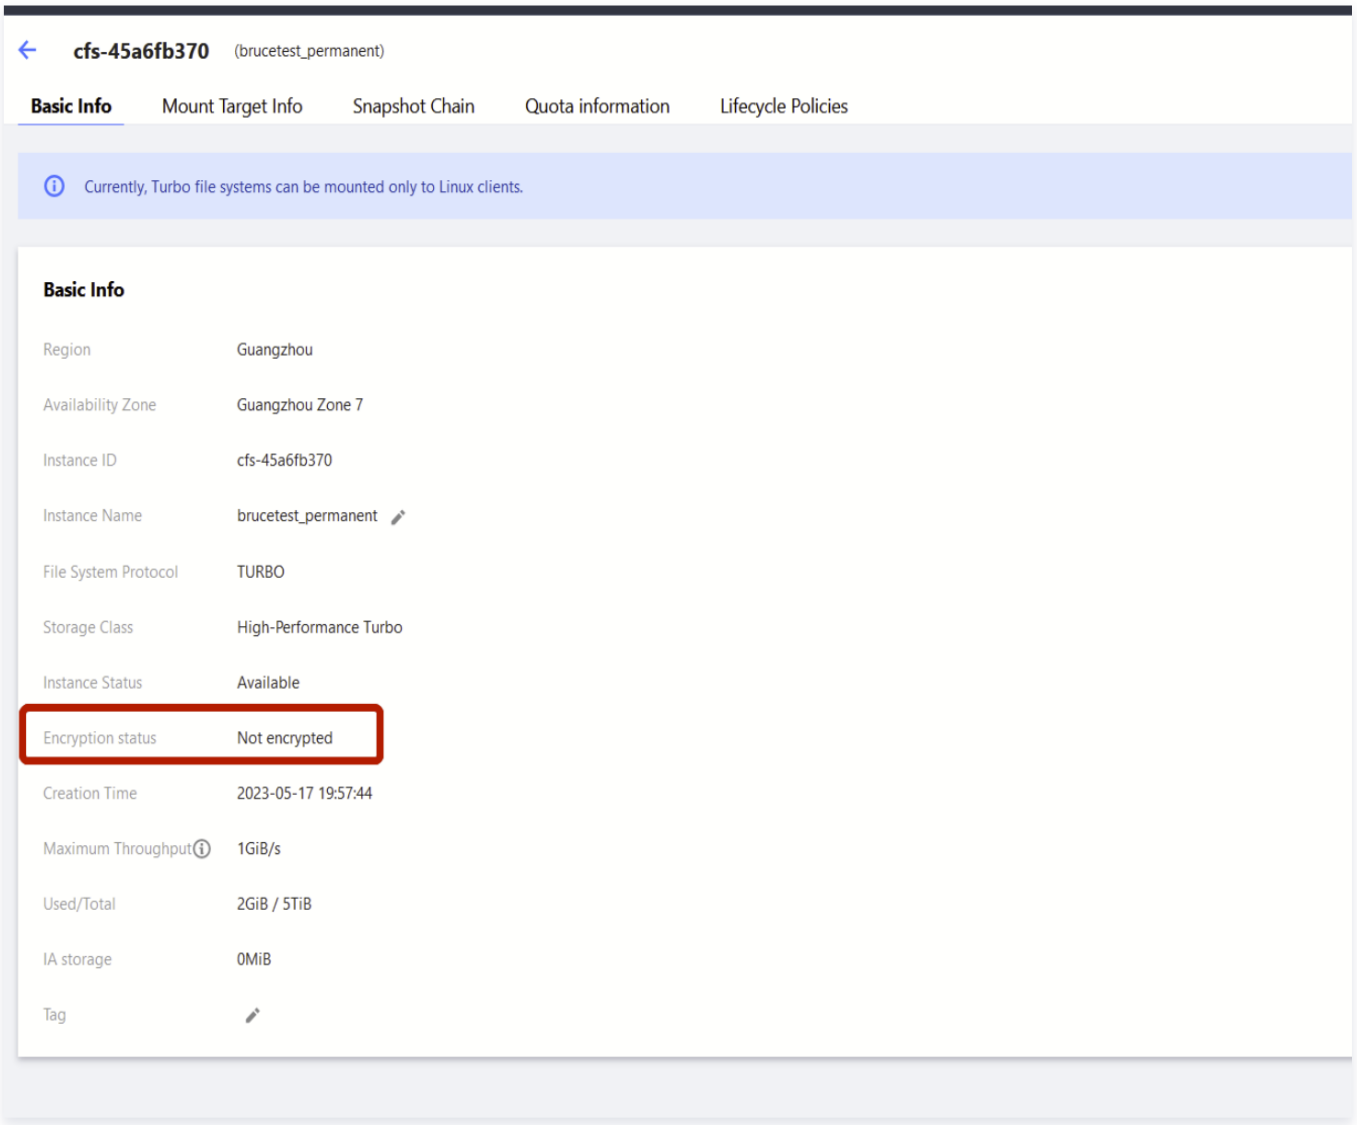

Limits

Last updated: 2023-08-29 10:32:48

The CFS encryption feature is subject to the following limitations:

Item	Description
CFS limitations	• 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 Limitations	After the encryption option is selected, the system will use CFS encryption keys that cannot be specified, deleted, or changed.

Data Lifecycle Features

Last updated: 2023-08-29 10:36:59

Overview

Data lifecycle management is an advanced data management feature provided by Tencent Cloud File System (CFS) to address the problem of balancing 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 storage with lower access frequency 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 lower storage costs.

Scenarios

Support for storage of massive amounts of data

In scenarios involving massive data storage, traditional cloud file storage was 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, to reduce the impact of frequent transitioning and restoration on your business, and make 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 hot/cold tiered data storage, 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 convert files that have not been accessed within 14, 30, 60, or 90 days to infrequenctly accessed (IA) storage files. Lifecycle management will determine whether to convert a file based on its access time (atime).

- The following operations will update the access time:
  - Reading a file
  - Writing file
- The following operations will not update the access time:
  - Querying file metadata (such as by performing the ls or state 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.

Operation Guide

Last updated: 2023-08-29 11:06:41

Setting a data lifecycle policy

Step 1. Create a lifecycle policy

1. Log in to the Lifecycle Policy Page and click **Create Policy**.

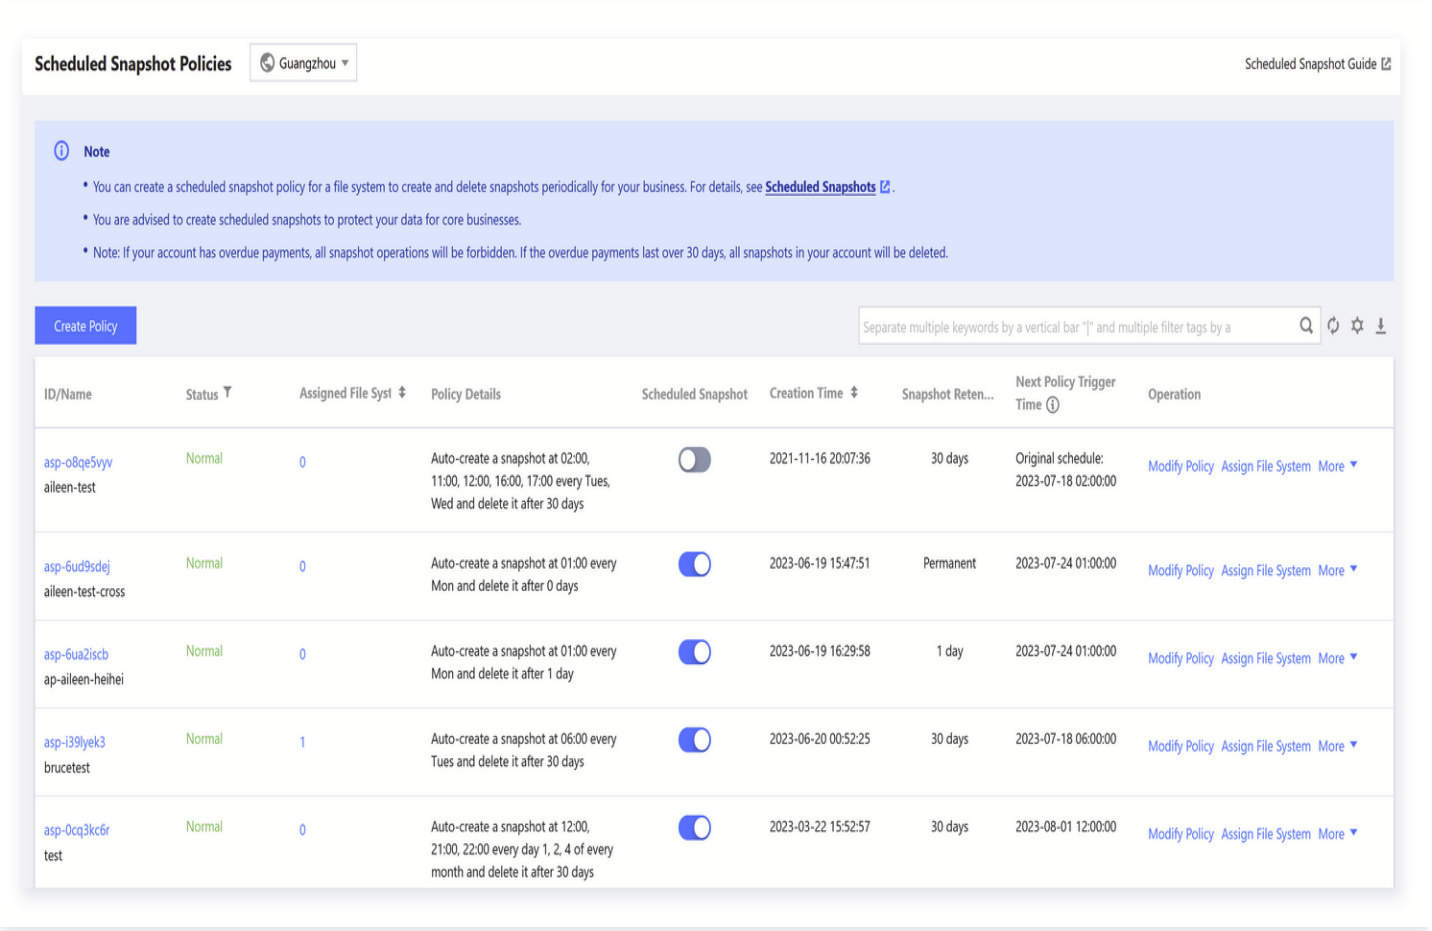

> **Note**
> This feature is exclusively supported by Turbo File System. To use it, please submit a ticket to contact us.

2. Specify a lifecycle management policy based on actual business needs, including time period and file size, then click **Next**.

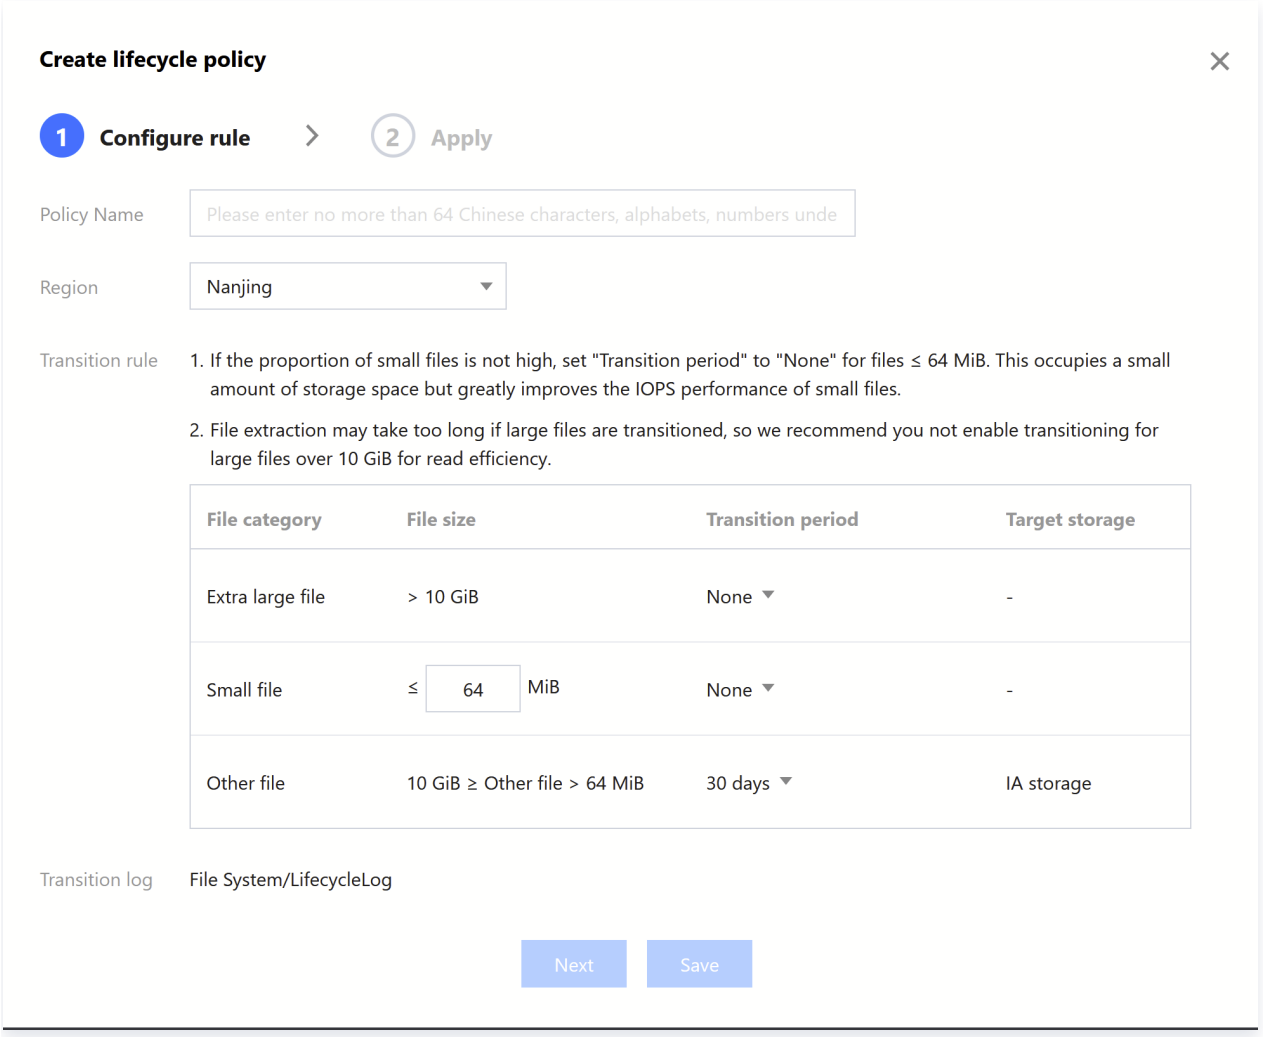

Note

A file less than or equal to 64 MB is defined as a small file.

Step 2. Configure a lifecycle policy

Apply the lifecycle policy to the specified path of the Turbo File System and click **Submit**.

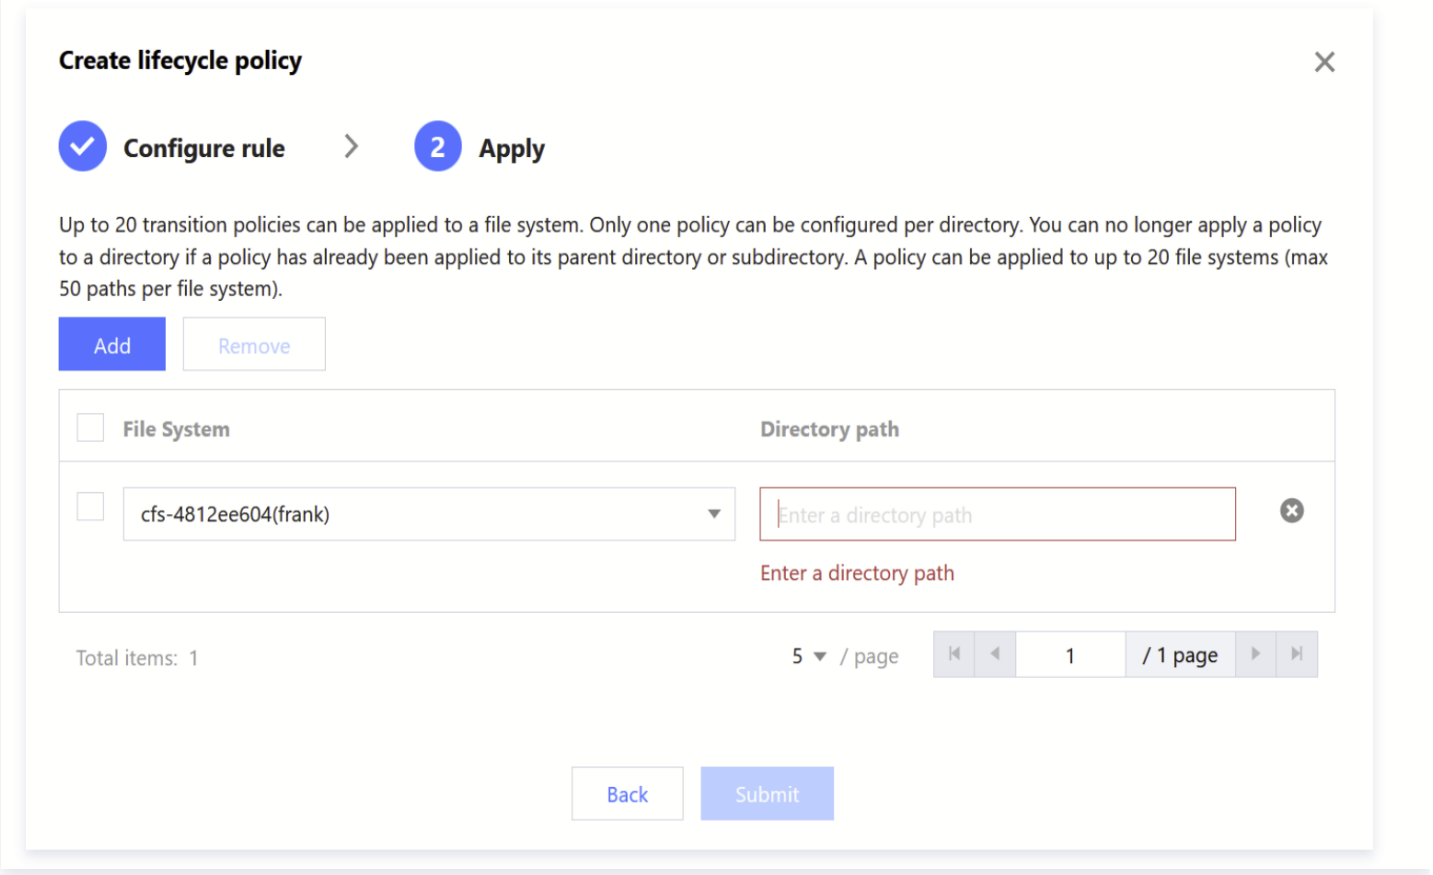

Note

- 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. Navigate to the Lifecycle Management Page.

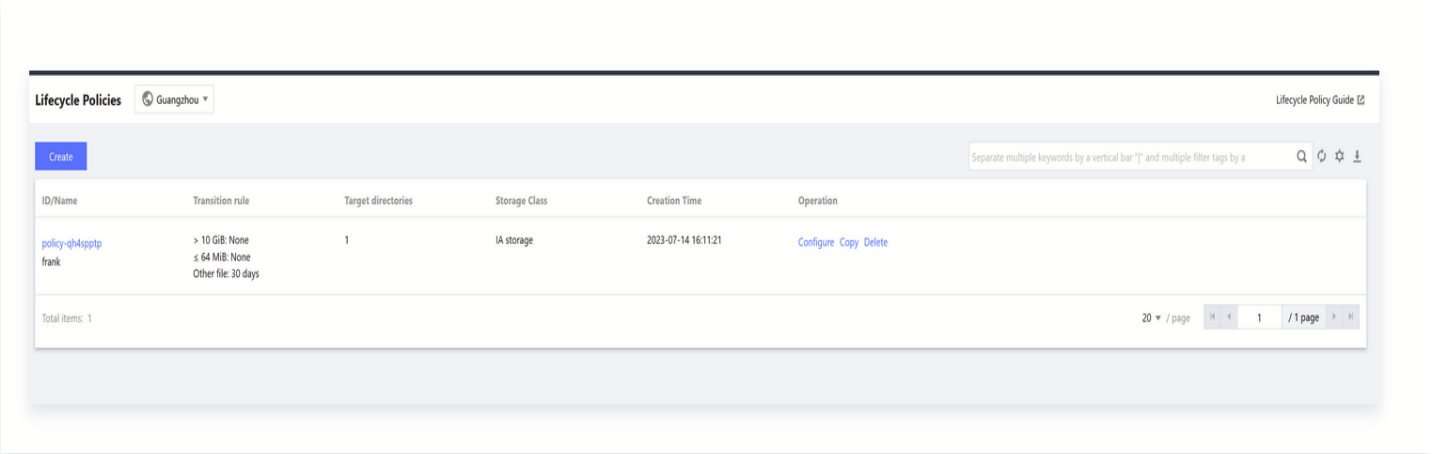

2. Click on the **Policy ID** to view the lifecycle management policy.

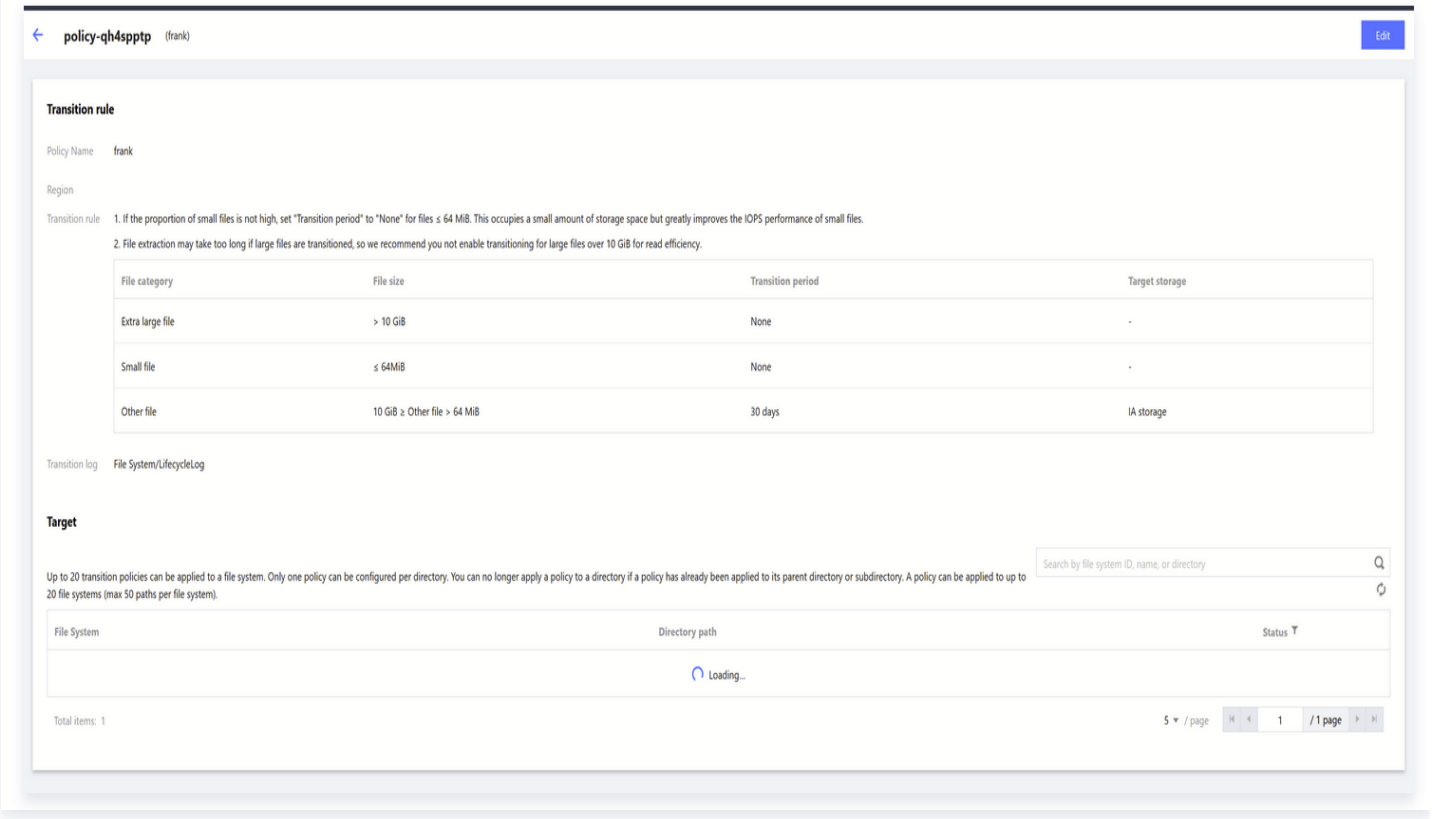

Step 2. View the active lifecycle policy for a specified file system

1. Navigate to the File System List Page and click on the ID of the file system instance you wish to view.

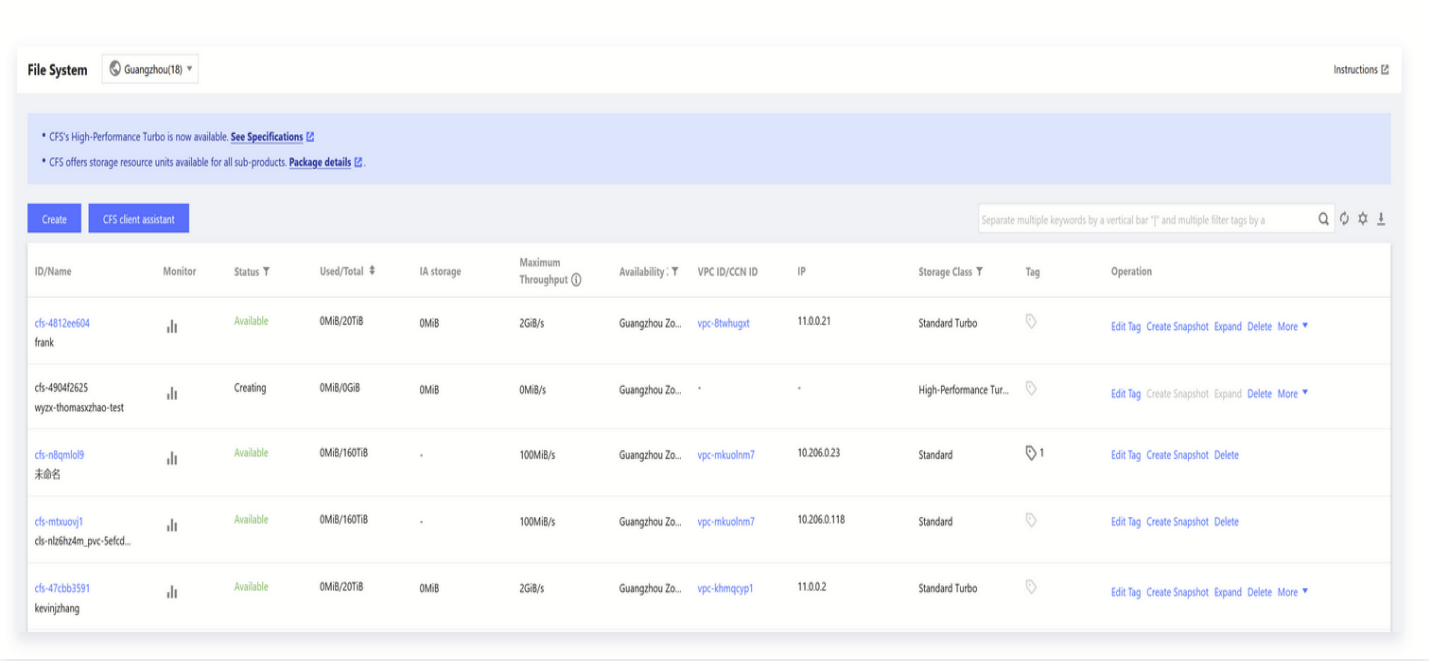

2. Navigate to the File Lifecycle Policy subpage to view the currently active lifecycle policies.

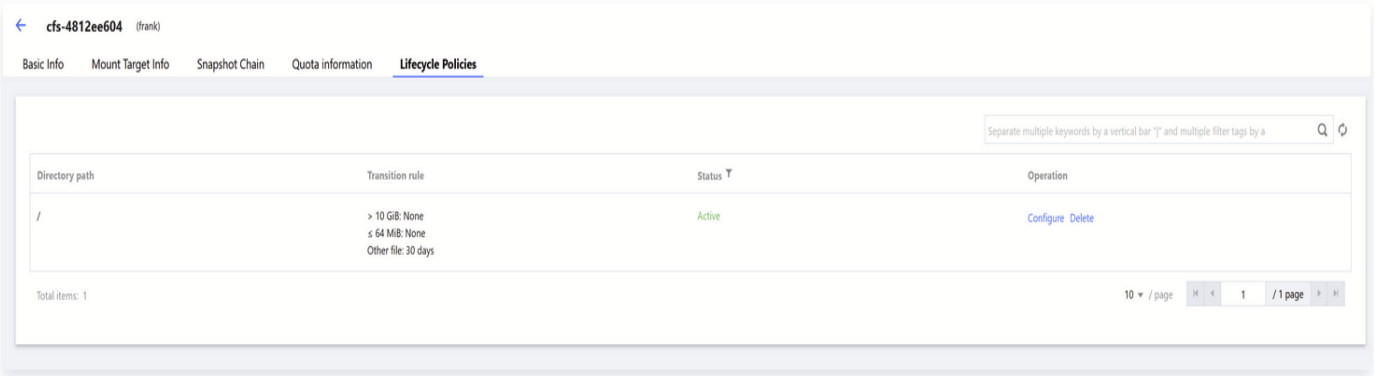

Modify/Delete Lifecycle Policy

1. After entering the lifecycle policy you wish to modify, click **Edit** in the upper right corner to make changes or delete it.

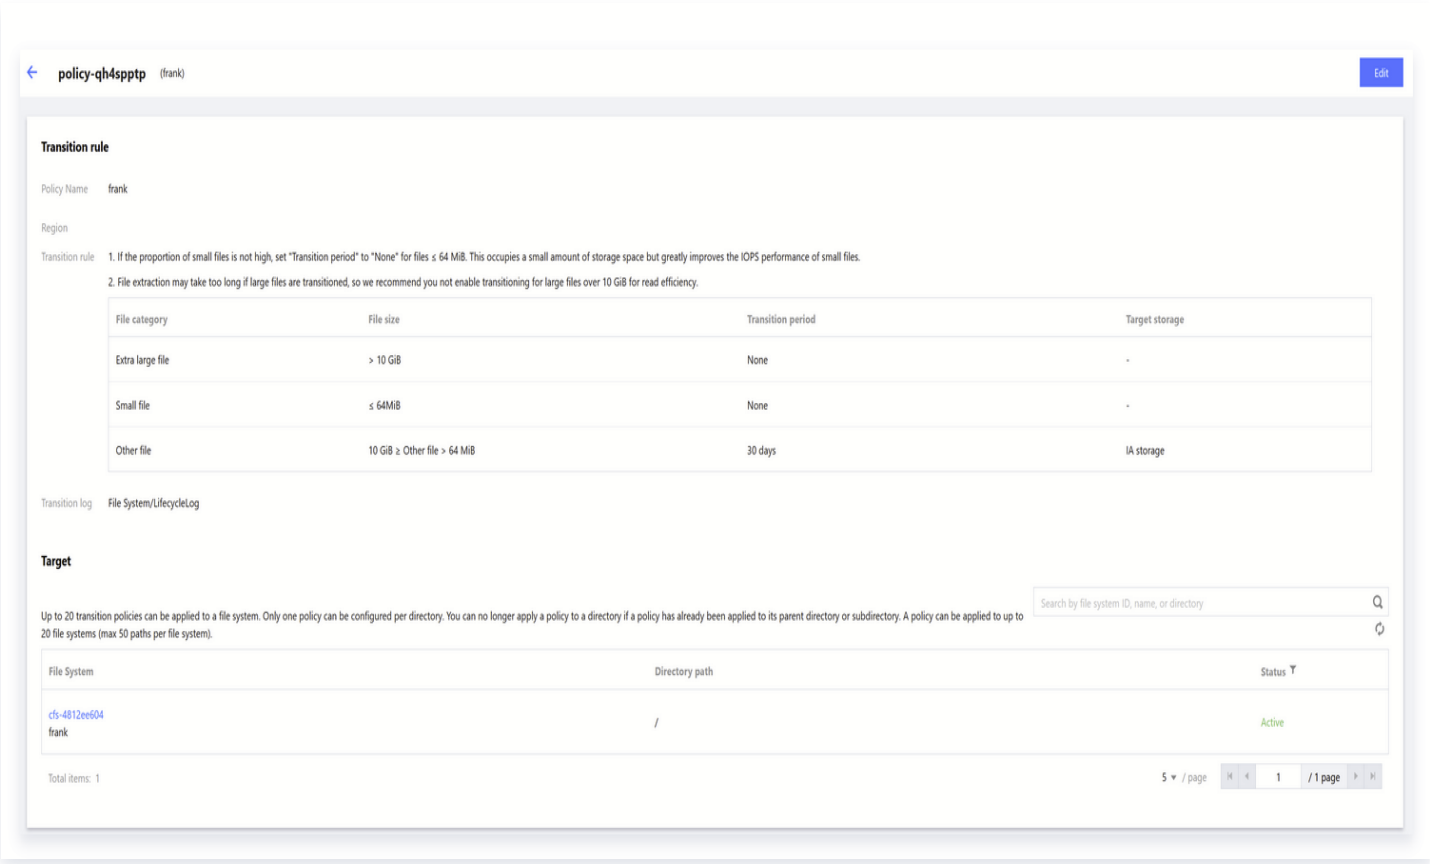

2. After making changes, click **Confirm**.

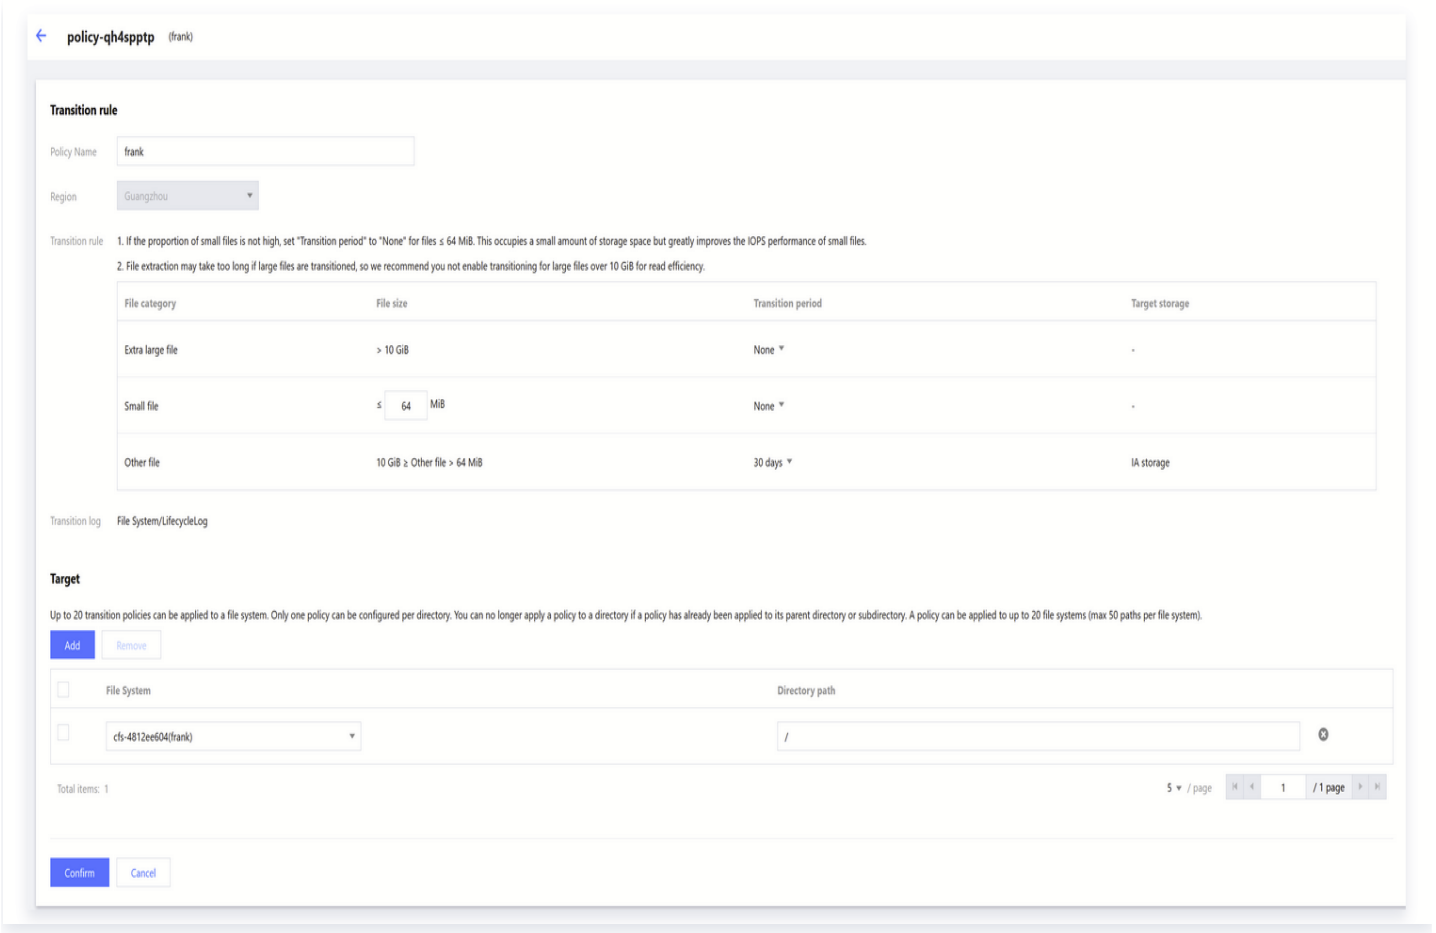

Viewing the settled data capacity

Click on the File System List Page to view the settled data capacity.

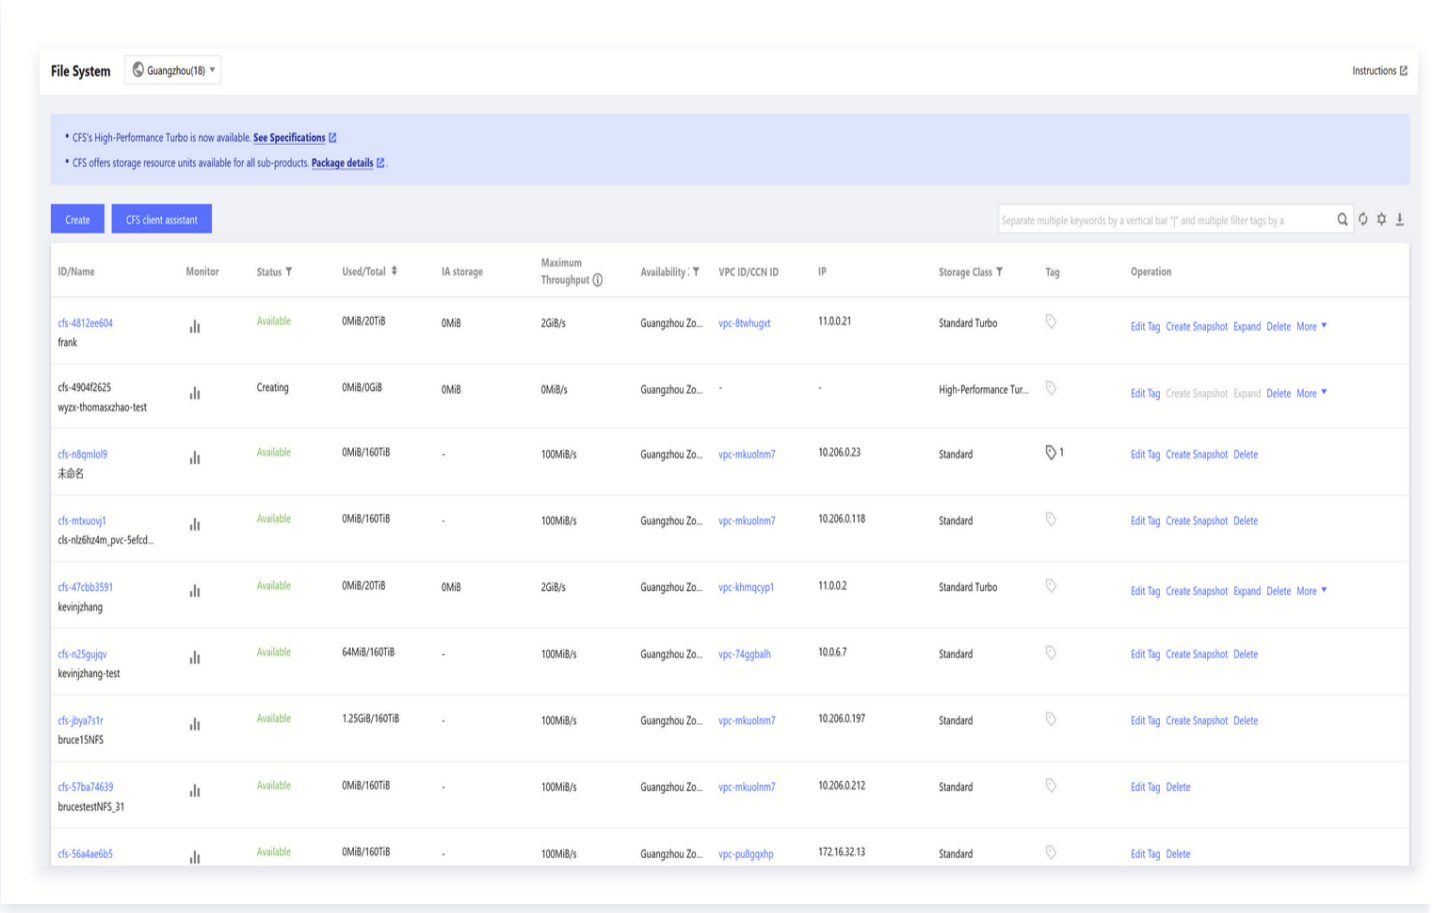

Viewing File Settlement Status

If you need to view the settlement status of a file, you can refer to the following instructions:

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

Note

- If 'archived' is returned, it indicates that the file has been transitioned to infrequent access storage but has not been released from the file system.
- If marked as 'released', it indicates that the file has been transitioned to infrequent access storage and has been released in the file system.
- By default, the system automatically performs a release operation one hour after archiving, freeing up file system space.

Load data

You can perform a batch prefetching action on the files in a specific directory using the following command:

- Single file:

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

- All files in the directory:

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

Proactively release data

After loading data through active prefetching, if you need to actively release it, refer to the following command:

- Single file:

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

- All files in the directory:

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

> **Note**
> Unsettled files cannot execute the hsm_release operation.

Limits

Last updated: 2023-08-29 10:40:14

Description

The limits are as described below:

Limit Type	Limit Description
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 system.
Lifecycle management policies associated with a file system	A file system can be associated with up to 20 life cycle management policy.
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 and child directories can be associated with up to 1 lifecycle management policy.

Note
This feature is exclusively available for Turbo File System. To utilize it, please submit a ticket to get in touch with us.

Supports and Limits

- 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 AI storage, and those in the same directory that have been in AI storage will remain unchanged in their storage status.
- Clusters created before March 15, 2023, do not currently support data lifecycle management. If you need to use this feature, please submit a ticket to contact us. Once the cluster is upgraded, it will support this feature.