

Object Storage Service

Parallel File System Feature Guide

Issue 07
Date 2024-03-19



Copyright © Huawei Cloud Computing Technologies Co., Ltd. 2024. All rights reserved.

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Cloud Computing Technologies Co., Ltd.

Trademarks and Permissions



HUAWEI and other Huawei trademarks are the property of Huawei Technologies Co., Ltd.

All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

The purchased products, services and features are stipulated by the contract made between Huawei Cloud and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Contents

1 Introduction.....	1
1.1 About PFS.....	1
1.2 Application Scenarios.....	1
1.3 Constraints.....	1
1.4 Using PFS.....	2
1.5 Billing.....	3
1.6 Features.....	3
1.6.1 Lifecycle Management.....	3
1.6.2 Permissions Configuration.....	4
2 Using PFS on a Console.....	6
2.1 Creating a Parallel File System.....	6
3 Using PFS with OBS APIs.....	8
3.1 Compatibility Between OBS APIs and PFS.....	8
4 FAQs.....	15
4.1 Why Is 256 TB Displayed After I Mounted a Parallel File System?.....	15
4.2 Can I Sort Files in the File List?.....	15
A Change History.....	16

1 Introduction

1.1 About PFS

Parallel File System (PFS) is a high-performance semantic file system provided by OBS. It features access latency in milliseconds, TB/s-level bandwidth, and millions of IOPS.

You can access data in a parallel file system via OBS APIs.

1.2 Application Scenarios

PFS is highly compatible, scalable, and reliable, and delivers amazing performance.

It is mainly used in the following scenarios:

Big data: log analysis, content recommendation, operation reports, user profiling, and interactive analysis

1.3 Constraints

Operations

- An existing OBS bucket cannot be changed to a parallel file system. For details about how to create a parallel file system, see [Creating a Parallel File System](#).
- Custom domain names of parallel file systems cannot be configured on OBS Console. You can configure them using APIs by referring to [Configuring a Custom Domain Name](#).

Functions

- [Image processing](#) currently cannot be used to process (such as downsize, resize, or watermark) images stored in parallel file systems.
- Server-side encryption is not supported.
- Cross-region replication is not supported.

- Versioning is not supported.
- Bucket inventory is not supported.
- Static website hosting is not supported.
- Configuration of default storage class for a parallel file system is not supported.
- Deep Archive storage is not supported.
- A parallel file system can be mounted to multiple Linux servers for concurrent reads, but this is not recommended for concurrent writes.

Performance

- A parallel file system provides a maximum bandwidth of 10 MB/s per TB by default.

Naming

- In a parallel file system, a file name cannot contain two consecutive slashes (/). For example, if you name a file as **test//123.txt**, an error will be reported.

1.4 Using PFS

You can use the console and REST APIs to process files stored in PFS anytime, anywhere and retrieve the processed files quickly. PFS supports both Portable Operating System Interface (POSIX) and OBS APIs, so you can process files the same way you process objects. This achieves interoperability between objects and files.

The table below describes the ways to use PFS in detail.

NOTE

Access permissions for OBS also apply to PFS. Before using PFS, make sure that you have the permissions required to access OBS resources.

Table 1-1 Ways to use PFS

Way	Function	Reference
PFS Console	On the console, you can create parallel file systems and manage them.	Creating a Parallel File System
OBS API	You can make API calls to use parallel file systems.	Compatibility Between OBS APIs and PFS
Other OBS tools	PFS allows for interoperability between objects and files, so you can also use other OBS tools (such as obsutil and OBS Browser+) to access PFS.	OBS Tools

1.5 Billing

Parallel file systems support both pay-per-use and yearly/monthly (resource packages) billing modes. For details about resource packages, see [Resource Package Overview](#).

For details about PFS pricing, see [Product Pricing Details](#).

1.6 Features

1.6.1 Lifecycle Management

The use cases and main functions of object lifecycle management also work on files in parallel file systems. For more information, see [Lifecycle Management](#).

To manage the lifecycle using SDKs, see [SDK Overview](#).

Differences Between File and Object Lifecycle Management

- A lifecycle rule can be used to manage files, but it cannot transition a folder to the Archive storage class. However, a lifecycle rule can delete an empty folder upon expiration.
- File deletion upon expiration and transition to the Archive storage class can be configured using either the API or console, while transition to Infrequent Access can only be configured using the API. PFS does not support versioning, so lifecycle actions (including deletion upon expiration and transition to the Archive or Infrequent Access storage class) that are related to versioning cannot be applied.
- If direct reading is enabled for a parallel file system, you can read files stored in the Archive storage class without restoring them first.
- A maximum of 20 lifecycle rules can be configured for a parallel file system.
- The time applied for a lifecycle rule to work on a file is when the data of the file was last updated.
- Lifecycle rules cannot transition files to the Deep Archive storage.
- After a lifecycle rule is configured for a parallel file system, there are limits on how many directories that the rule can be applied to. If the configuration exceeds the limit, the lifecycle rule execution will be prolonged.
 - a. There can be no more than 100,000 level-1 subdirectories in each directory.
 - b. There can be no more than 10 million subdirectories (folders) matching the prefix defined in the rule in total.
 - c. There can be no more than 30 million files matching the prefix defined in the rule in total.

Notes

- If renamed files or files in a renamed folder meet the conditions specified in a lifecycle rule, the time when the file data was most recently updated, not

when the files were renamed, will be applied for the lifecycle rule to take effect. In addition, the effective time of the lifecycle rule may be delayed for up to seven days.

- For a file copy on a client, the lifecycle rule determines when to expire the file copy or transition it to the Archive storage class based on the file copy creation time.
 - For example, if a file, **src.txt**, was created on January 1, 2019, and was then copied to the **des.txt** file by running the **cp -a src.txt des.txt** command on September 1, 2019. Then the lifecycle rule calculated when to perform the specified actions on the file copy based on September 1, 2019.
- In the lifecycle rule of a parallel file system, directories in the file system are periodically scanned, and then deleted if they meet the expiration conditions. Scan intervals (usually seven days) vary depending on cluster configurations. A scan starts from the deepest directory, and the scanned empty directories that meet the expiration conditions will be deleted, but those non-empty directories will not be processed. For this mechanism, a single-level directory that meets the expiration conditions will be deleted 0 to 7 days after it has been emptied. Accordingly, a two-level directory will be deleted 0 to 14 days after it has been emptied. Each time a directory level is added, the waiting time for deletion increases by seven days.

1.6.2 Permissions Configuration

The use cases and main functions of object access control also work on files in parallel file systems. For more information, see [Permissions Configuration Guide](#).

Differences Between File and Object Permission Configurations

To exactly match a specific directory, the resource path in the policy must end with a slash (/). When checking permissions, parallel file systems consider objects as directories. If the object identifier does not end with a slash (/), the system will add a slash (/) to the end of the object identifier and then performs a policy matching.

IAM Permission Configuration Examples

Example 1: Grant a user the permissions required to download **dir_1**, excluding its subdirectories.

In the following configuration, the resource path ends with a slash (/). In such case, a success response can be returned when **dir_1** or **dir_1/** is contained in the URL of a head request.

Note that this configuration is not applied to subdirectories or files in **dir_1**. Therefore, a failure response will be returned if a head request is sent to **dir_1/file1**.

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "obs:object:GetObject",
```

```
    ],  
    "Resource": [  
      "obs:*:*:object:examplebucket/dir_1/",  
    ]  
  }  
]  
}
```

Example 2: Grant a user the permissions required to download **dir_1** and its subdirectories.

In the following configuration, the resource path uses prefix-based matching and ends with a wildcard (*). In such case, a success response can be returned when a head request is sent to **dir_1/file1**.

```
{  
  "Version": "1.1",  
  "Statement": [  
    {  
      "Effect": "Allow",  
      "Action": [  
        "obs:object:GetObject",  
      ],  
      "Resource": [  
        "obs:*:*:object:examplebucket/dir_1/*",  
      ]  
    }  
  ]  
}
```

Bucket Policy Configuration Examples

Example: Grant a user the permissions required to download **dir_1**, excluding its subdirectories.

In the following configuration, the resource path ends with a slash (/). In such case, a success response can be returned when **dir_1** or **dir_1/** is contained in the URL of a head request.

```
{  
  "Statement": [  
    {  
      "Sid": "test",  
      "Effect": "Allow",  
      "Principal": {"ID": ["domain/b4bf1b36d9ca43d984fbc9491b6fce9:user/  
71f3901173514e6988115ea2c26d1999"]},  
      "Action": ["*"],  
      "Resource": [  
        "examplebucket/dir_1/",  
      ]  
    }  
  ]  
}
```


2 Using PFS on a Console

2.1 Creating a Parallel File System

You can create a parallel file system on OBS Console.

Procedure

- Step 1** On the console homepage, click **Service List** in the upper left corner and choose **Storage > Object Storage Service**.
- Step 2** In the navigation pane, choose **Parallel File Systems**.
- Step 3** In the upper right corner of the page, click **Create Parallel File System**.

Figure 2-1 Creating a parallel file system

The screenshot shows the 'Create Parallel File System' form in the OBS console. It is divided into several sections:

- Region:** A dropdown menu is set to 'CN East-Shanghai1'. Below it, there is explanatory text: '- Regions are geographic areas isolated from each other. Resources are region-specific and cannot be used across regions through internal network connections. For low network latency and quick resource access, select the nearest region. Once a parallel file system is created, the region cannot be changed.' and '- Parallel file systems are not available in Dedicated Cloud (DeC) scenarios.'
- Data Redundancy Policy:** Two buttons are present: 'Multi-AZ storage' (selected) and 'Single-AZ storage'. Below the buttons, text explains: 'Data is stored in multiple AZs in the same region, improving availability. Pricing details' and 'If a file system is created in the single-AZ mode, data in the file system is stored in only one AZ. The single-AZ mode applies to data storage that requires low latency. Multi-AZ storage uses relatively high billing standards.'
- File System Name:** A text input field contains 'postik-123456'. Below the field, there are three warning icons: 'ⓘ The file system name must be unique.', 'ⓘ The file system name must be unique.', and 'ⓘ The value cannot be modified after the creation.'
- Policy:** Three buttons are shown: 'Private' (selected), 'Public Read', and 'Public Read and Write'. Below the buttons, text states: 'Only you and users authorized by you are allowed to access the parallel file system.'
- Tags:** A section with the text: 'It is recommended that you use TMS's predefined tag function to add the same tag to different cloud resources. View predefined tags'. Below this, there are two input fields: 'Tag key' and 'Tag value'. At the bottom, it says 'You can add 10 more tags.'

- Step 4** Select a region and enter a name for the parallel file system.

 **NOTE**

- Once a parallel file system is created, its name cannot be changed.
- URLs do not support uppercase letters and cannot distinguish between names containing uppercase or lowercase letters. For example, if you attempt to access the parallel file system **MyFileSystem** using a URL, the file system name will be resolved to **myfilesystem**, causing an access error. For this reason, a parallel file system name can contain only lowercase letters, digits, periods (.), and hyphens (-).

Step 5 Configure a data redundancy policy. If you select multi-AZ storage, data will be stored in multiple AZs for a higher reliability.

 **NOTE**

Once a parallel file system is created, you cannot change the data redundancy policy, so take care when selecting a data redundancy policy.

Step 6 Configure a policy. You can select **Private**, **Public Read**, or **Public Read/Write** for the parallel file system.

Step 7 Configure direct reading. With direct reading enabled, you can directly download objects in the Archive storage class without restoring them first. Direct reading is a billable function. For details, see [Product Pricing Details](#).

Step 8 (Optional) Add tags. Tags are used to identify parallel file systems in OBS, for the purpose of classification. Each tag is represented by one key-value pair. For details about how to add a tag, see [Tags](#).

Step 9 (Optional) Buy storage packages. By default, you are billed on a pay-per-use basis for using parallel file systems. You can also purchase [storage packages](#) to save more. After selecting a required package, go to the **Confirm** page to complete the purchase.

Storage packages can also be purchased after the parallel file system is created.

Step 10 Confirm the settings at the bottom of the page and click **Create Now**.

Step 11 View the file system you created just now in the parallel file system list.

Then, you can use the parallel file system the same way you use a bucket. For details, see [Using PFS](#).

----End

3 Using PFS with OBS APIs

3.1 Compatibility Between OBS APIs and PFS

You can call some OBS APIs to use PFS. There may be additional requirements when you call these APIs.

For details about the OBS APIs, see [Object Storage Service API Reference](#).

APIs for Basic Bucket Operations

Table 3-1 APIs for basic bucket operations

API	PFS Compatible	Differences
Listing buckets	Yes	The x-obs-bucket-type:POSIX header is required for obtaining the list of parallel file systems.
Creating a bucket	Yes	The x-obs-fs-file-interface:Enabled header is required for creating a parallel file system.
Listing objects in a bucket	Yes	-
Obtaining bucket metadata	Yes	-
Obtaining bucket region locations	Yes	-
Deleting a bucket	Yes	-

APIs for Advanced Bucket Settings

Table 3-2 APIs for advanced bucket settings

API	PFS Compatible	Differences
Configuring a bucket policy	Yes	-
Obtaining bucket policy information	Yes	-
Deleting a bucket policy	Yes	-
Configuring a bucket ACL	Yes	-
Obtaining bucket ACL information	Yes	-
Configuring logging for a bucket	Yes	-
Obtaining a bucket logging configuration	Yes	-
Configuring bucket lifecycle rules	Yes	-
Obtaining bucket lifecycle configuration	Yes	-
Deleting bucket lifecycle rules	Yes	-
Configuring versioning for a bucket	No	-
Obtaining bucket versioning status	No	-
Configuring event notification for a bucket	Yes	-

API	PFS Compatible	Differences
Obtaining the event notification configuration of a bucket	Yes	-
Configuring storage class for a bucket	No	-
Obtaining bucket storage class information	No	-
Configuring cross-region replication for a bucket	No	-
Obtaining the cross-region replication configuration of a bucket	No	-
Deleting the cross-region replication configuration of a bucket	No	-
Configuring tags for a bucket	Yes	-
Obtaining bucket tags	Yes	-
Deleting bucket tags	Yes	-
Configuring bucket storage quota	Yes	-
Querying bucket storage quota	Yes	-
Querying information about used space in a bucket	Yes	-

API	PFS Compatible	Differences
Configuring bucket inventories	No	-
Obtaining bucket inventories	No	-
Listing bucket inventories	No	-
Deleting bucket inventories	No	-
Configuring a custom domain name for a bucket	Yes	-
Obtaining the custom domain name of a bucket	Yes	-
Deleting a custom domain name of a bucket	Yes	-
Configuring bucket encryption	No	-
Obtaining bucket encryption configuration	No	-
Deleting the encryption configuration of a bucket	No	-
Configuring the direct reading policy for Archive objects in a bucket	Yes	-
Obtaining the direct reading policy for Archive objects in a bucket	Yes	-

API	PFS Compatible	Differences
Deleting the direct reading policy for Archive objects in a bucket	Yes	-

APIs for Static Website Hosting

Table 3-3 APIs for static website hosting

API	PFS Compatible	Differences
Configuring static website hosting for a bucket	No	-
Obtaining the static website hosting configuration of a bucket	No	-
Deleting the static website hosting configuration of a bucket	No	-
Configuring bucket CORS	Yes	-
Obtaining the CORS configuration of a bucket	Yes	-
Deleting the CORS configuration of a bucket	Yes	-
OPTIONS buckets	No	-
OPTIONS objects	No	-

APIs for Object Operations

Table 3-4 APIs for object operations

API	PFS Compatible	Differences
PUT objects	Yes	<ul style="list-style-type: none"> • Headers not supported: x-obs-storage-class, x-obs-website-redirect-location, success-action-redirect, and x-obs-expires • Objects uploaded using this API cannot be directly stored in the Infrequent Access or Archive storage class and are stored in the Standard storage class by default. You can later change the storage class by using a lifecycle rule or modifying the metadata.
POST objects	Yes	Headers not supported: x-obs-storage-class , x-obs-website-redirect-location , success-action-redirect , and x-obs-expires
Copying objects	Yes	Data can be replicated only between parallel file systems or buckets that are in the same cluster.
Obtaining object content	Yes	-
Obtaining object metadata	Yes	-
Deleting an object	Yes	-
Batch deleting objects	Yes	-
Restoring Archive objects	Yes	-
Appending objects	No	-
Configuring object ACL	Yes	-
Obtaining object ACL information	Yes	-

API	PFS Compatible	Differences
Modifying object metadata	Yes	In a parallel file system, the storage class of a directory cannot be changed. To change the storage class of a file in the directory, modify the metadata of the file or use a lifecycle rule to change the storage class of files in batches.
Modifying an object	Yes	This is a PFS only API and is not supported for OBS buckets.
Truncating an object	Yes	This is a PFS only API and is not supported for OBS buckets.
Renaming an object	Yes	This is a PFS only API and is not supported for OBS buckets.

APIs for Multipart Uploads

Table 3-5 APIs for multipart uploads

API	PFS Compatible	Differences
Listing initialized multipart tasks in a bucket	Yes	-
Initiating multipart upload tasks	Yes	-
Uploading parts	Yes	-
Copying parts	Yes	Copying parts is not supported for an appended file.
Listing uploaded parts	Yes	-
Merging parts	Yes	-
Canceling multipart tasks	Yes	-

4 FAQs

4.1 Why Is 256 TB Displayed After I Mounted a Parallel File System?

There is no limit on the capacity of a parallel file system. When you run the `df -h` command on a client, the system returns 256 TB just for display purposes. 256 TB itself does not have any meaning.

4.2 Can I Sort Files in the File List?


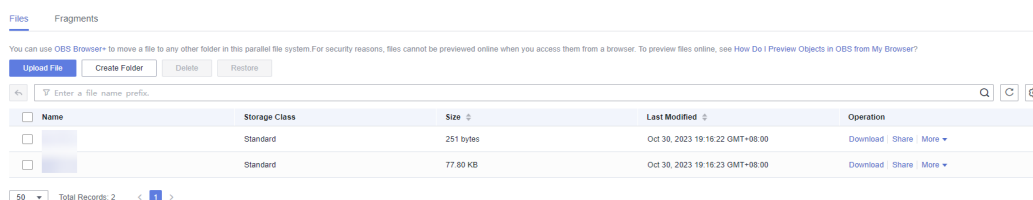
In the file list, you can click  next to the size or last modification time to sort files. File search by last modification time can only display the first 1,000 records. If there are more than 5,000 files in a parallel file system, the files are sorted in alphabetical order and can be searched only by file name prefix.

Figure 4-1 File list



The screenshot shows a file list interface with the following columns: Name, Storage Class, Size, Last Modified, and Operation. There are two rows of data. The first row has a size of 251 bytes and a last modified date of Oct 30, 2023 19:16:22 GMT+08:00. The second row has a size of 77.80 KB and a last modified date of Oct 30, 2023 19:16:23 GMT+08:00. The interface includes buttons for Upload File, Create Folder, Delete, and Restore, and a search bar with the placeholder text 'Enter a file name prefix'.

Name	Storage Class	Size	Last Modified	Operation
	Standard	251 bytes	Oct 30, 2023 19:16:22 GMT+08:00	Download Share More
	Standard	77.80 KB	Oct 30, 2023 19:16:23 GMT+08:00	Download Share More

A Change History

Released On	Description
2024-03-19	This is the seventh official release. This issue incorporates the following change: <ul style="list-style-type: none">Added naming restrictions in Constraints.
2024-02-28	This is the sixth official release. This issue incorporates the following change: <ul style="list-style-type: none">Added the content related to Deep Archive storage (under limited beta testing).
2023-10-30	This is the fifth official release. This issue incorporates the following change: <ul style="list-style-type: none">Added Can I Sort Files in the File List?
2023-07-10	This is the fourth official release. This issue incorporates the following change: <ul style="list-style-type: none">Updated the content under Notes in section "Lifecycle Management."
2023-06-16	This is the third official release. This issue incorporates the following change: <ul style="list-style-type: none">Added Why Is 256 TB Displayed After I Mounted a Parallel File System?
2020-07-30	This is the second official release. This issue incorporates the following change: <ul style="list-style-type: none">Added the lifecycle management feature.
2019-11-22	This is the first official release.